THE ECONOMIC JOURNAL

The Economic Journal, **127** (*October*), F24–F49. Doi: 10.1111/ecoj.12334 © 2017 The Authors. The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

WHY ARE HOUSEHOLDS THAT REPORT THE LOWEST INCOMES SO WELL-OFF?*

Mike Brewer, Ben Etheridge and Cormac O'Dea

We document that households in the UK with extremely low measured income tend to spend much more than those with merely moderately low income. This phenomenon is evident throughout three decades worth of microdata and across different employment states, levels of education and marital statuses. Of the likely explanations, we provide several arguments that discount over-reporting of expenditure and argue that under-reporting of income plays the major role. In particular, by using a dynamic model of consumption and saving, and paying special attention to poverty dynamics, we show that consumption smoothing cannot explain all the apparent dissaving.

Economists have consistently argued, on both theoretical and practical grounds, that having a low level of consumption provides a better measure of poverty than having a low level of income. On a theoretical level, consumption in any one time period is more closely aligned with life-time utility than is income, as households with high life-time resources can borrow to smooth consumption when income is temporarily low (Poterba, 1989; Cutler and Katz, 1992; Slesnick, 1993). On a practical level, Meyer and Sullivan (2003) argue that, in US surveys at least, income is more likely than expenditure (which can proxy for consumption) to be under-reported by those with the least resources.

Nevertheless, there are obvious reasons why poverty is generally measured using income. Principally, governments dispense benefits according to income rather than expenditure, as the latter is more manipulable and harder to monitor. Income measures are therefore more closely aligned to the mechanisms of government intervention, and, for good or ill, income measures therefore often guide policy. In the UK, for example, there are four statutory measures of child poverty against which the government of the day has to report progress annually (and, ideally, 'eradicate' by

* Corresponding author: Mike Brewer, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ, UK. Email: mbrewer@essex.ac.uk.

We are very grateful to Thomas Crossley for advice and encouragement and to Robert Joyce, Laura Keyse, Barra Roantree, Richard Tonkin and Karen Watkins for a number of useful conversations and to seminar participants at the Institute for Fiscal Studies, the Institute for Social and Economic Research at the University of Essex, and the Department for Economics at the University of Sheffield for useful comments. The authors gratefully acknowledge funding from the ESRC Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (Reference RES-544-28-5001) (Brewer, Etheridge and O'Dea) and from the ESRC Research Centre on Micro-Social Change at the University of Essex (ES/L009153/1) (Brewer). The Living Cost and Food Survey and its predecessors, and the Family Resources Survey, are Crown copyright and are reproduced with the permission of the Controller of HMSO and the Queen's Printer for Scotland, and are available from the Economic and Social Data Service (http://www.esds.ac.uk). The ESDS, the original owners of the data (the Office for National Statistics and the Department for Work and Pensions respectively) and the copyright holders bear no responsibility for their further analysis or interpretation. Any errors remain the responsibility of the authors.

2020–1; see Brewer *et al.*, 2011), all of which define poverty in terms of a low household income. Similar targets exist across the European Union.

Given that both expenditure-based and income-based poverty measures are important, it is crucial, therefore, to understand what causes them to diverge (which they often do): is it the theoretical reasons mentioned above, the practical, measurement, issues, or some combination of the two? This question is important because the competing hypotheses have very different implications for economic analysis. If under-reporting of income is more important, then this implies that poverty measurement is subject to the vagaries of the survey instrument and that more effort should be put into survey design and implementation, and that the analyst should employ statistical techniques to limit the effect of contamination. On the other hand, if consumption smoothing is more important, then the bottom tail of income recorded in surveys contains vital information on economic behaviour and shows how important are stocks of assets, credit markets and other insurance mechanisms to household welfare.

This article attacks this question by analysing a puzzling and, as we show, pervasive feature of the distribution of expenditure conditional on income at the bottom of the income distribution. Figure 1 illustrates the puzzle, using data from the UK Living Costs and Food Survey (LCFS) to plot non-parametrically the relationship between median spending given income and median income given spending. Median income is monotonically increasing in household expenditure but median expenditure for those with very low income (below £75 per week, or those in the bottom 1% of the income distribution) is very high relative to income. In fact, those at the very bottom of the income distribution have a median expenditure equal to the population-median level of expenditure (roughly £400). From approximately £110 per week (around the second percentile), median expenditure rises with income. In short, the graph of median expenditure by reported income contains a sharp non-monotonicity; that is it maps out a 'tick' (for North Americans: a 'check'). As we show later, this tick has existed to a greater or lesser extent over time and is evident among different sub-groups of households.

This pattern could be explained by any combination of under-reporting of income, over-reporting of expenditure, or that households smooth consumption over time. This article investigates what role each of these candidates has to play. We argue that over-reporting of expenditure is unlikely to be playing a large role here, in part because of evidence from surveys that use a 'balance edit' (in which households who report expenditure far in excess of income are asked to account for the source of their spending) that low-income households tend to report expenditure levels more accurately than income, but also because the tick pattern has been shown to exist for measures of living standards other than consumption. Of under-reporting of income and consumption smoothing, we provide several pieces of evidence that the former plays the major role. First, we show that a similar tick exists for households of very different composition, and remains when we condition on employment status. Given that most transitions into poverty are between rather than within work status (Altonji *et al.*, 2013), then conditioning on labour market status should remove most of the effect of the consumption smoothing channel. Second, a comparison of our microdata and administrative data strongly suggests that some specific components of income are

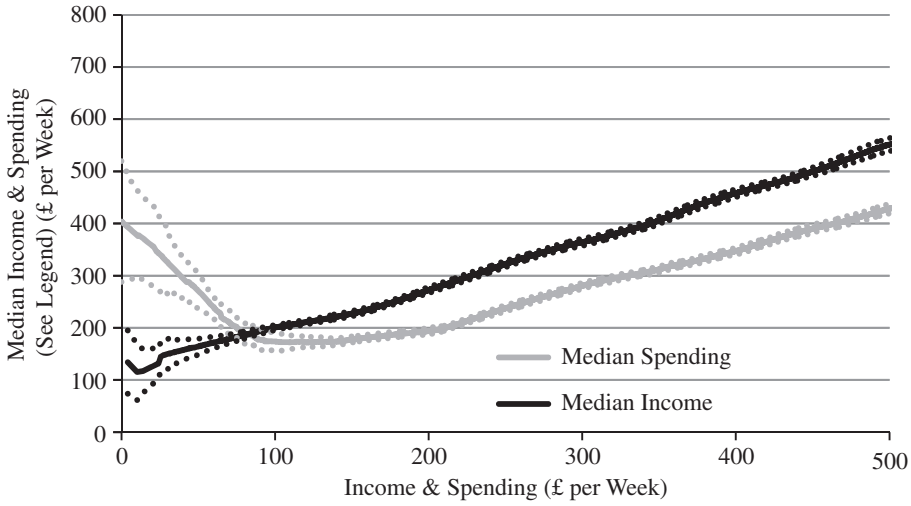

Fig. 1. *Median Expenditure by Income, and Median Income by Expenditure*
Notes. Data are from 2005 to 2009. Values on the axes are given in equivalised real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale. The solid lines in the graph represent smoothed conditional medians and are generated by performing locally-weighted median regressions. Further details are given in Section 2. The dotted lines represent the 95% confidence intervals for the smoothed median.

under-recorded. In particular, we find that, as Meyer *et al.* (2009) show for several surveys in the US, UK household surveys do not capture anything like the amount of cash transfer payments which the government reports paying out; we also show that this under-recording has grown in recent years as a proportion of household income. Third, we provide indirect evidence that income is under-recorded by analysing the data through the lens of a dynamic model of consumption and saving. The methodology here is to specify a model which favours the consumption-smoothing hypothesis and then to show that it cannot reproduce the tick. The main result of the simulations is that, even with the most 'generous' model (one without borrowing constraints) only around 60% of the tick can be explained. If we choose less generous specifications, the fraction comes down: in particular, when borrowing is constrained (a more realistic setting than the model without borrowing constraints), then only a small fraction of the tick can be explained.

We are, of course, not the first to highlight the nature of the income-expenditure relationship at the bottom of the income distribution, nor to discuss the relative merits of consumption and income for measuring poverty. Many papers (including Sabelhaus and Groen, 2000; Meyer and Sullivan, 2003, 2008, 2011*a*,*b*) examine the income-expenditure relationship for US data (specifically, for the US Consumer Expenditure Survey (CEX) survey) but it is reasonable to wonder whether their findings might reflect some feature of that particular survey instrument. Our first contribution, therefore, is to provide comprehensive evidence on the existence, size and nature of the tick using a different survey instrument; one that collects information on both income and spending in different ways from the CEX. Our second contribution is to assess whether the observed data is consistent with a model of consumption and saving.

This is a similar aim to that pursued by Sabelhaus and Groen (2000): after observing that the bottom decile of the income distribution in the CEX records expenditure far above income, they compute average permanent income across decile groups and argue that the joint distributions of expenditure, current income and permanent income cannot be reconciled with the permanent income model. We build on this work by solving explicitly for optimal consumption and specifying an income process that specifically models low-income shocks. Finally, we also argue that the mismatch between low income and low spending in household surveys is especially important in the UK given the continuing use of income-based measures of poverty in national and European policy debates. Although some mismatch between income and spending at the bottom of the distribution has been documented before (Saunders *et al.*, 2002; Attanasio *et al.*, 2006; Brewer *et al.*, 2006), we provide a comprehensive assessment across all groups on society and using over three decades of micro-data.

The article proceeds as follows: after briefly describing the data in Section 1, we examine the expenditure tick in detail in Section 2. Section 3 examines the direct evidence that income is under-recorded at the bottom end of the distribution. In Section 4 we examine whether consumption smoothing can explain the tick using the dynamic model. Section 5 concludes.

1. Data

Our main data set is the LCFS (known between 2001 and 2007 as the Expenditure and Food Survey, and the Family Expenditure Survey before that; we refer to it as the LCFS). We use data from 1978 to 2009. The LCFS is an annual, nationally-representative, cross-sectional survey with an annual sample of between 5,000 and 7,000 households. Its primary purpose is to provide data for the calculation of the commodity weights for the UK's price indices. It aims to collect a comprehensive measure of household spending with a two-week diary, in which respondents are asked to record everything they purchase, supplemented by a questionnaire in which respondents are asked about any spending on infrequently purchased items over the past number of months.[1] The LCFS is the only nationally-representative UK survey which contains comprehensive measures of both income and expenditure.

Our aim in this article is to assess the causes of differences between individual households' income and expenditure over a particular period and, therefore, it is important that our measures of both income and expenditure are as broad and consistent as possible.

Our measure of income is comprehensive: it includes income from employment and self-employment, state and private pensions, state benefits, assets and allowances from outside the household. It includes all those cash components included in the measure of income typically used to measure income poverty in the UK – known as the 'households below average income' (HBAI) definition of income. The components that are included in the HBAI series but not included in our income measure are imputed or non-cash receipts (such as the value of free TV licences that the elderly

[1] The number of months differs by item. For example, respondents are asked to record any spending on motor vehicles in the past 12 months but any spending on household fuel in the past three months.

receive, the value of housing benefit paid directly to landlords and the value of free-school meals given to children). As we are interested in comparing flows of cash into a household with flows of cash out (a quantity that has to equal net saving plus net measurement error in the two series), we do not include the value of these in-kind receipts in our measure of income. The horizon over which income receipts are measured differs by the type of receipt. For earnings, 'usual' earnings are recorded. For most respondents this is their most recent wage or salary payment but is replaced with the 'usual' wage or salary payment if the last payment was deemed by the respondent to be 'unusual'. For state pensions, private pensions, and benefits, the most recent receipt is recorded. For asset income (i.e. interest and dividends), where receipts are likely to be less frequent, all payments over the previous 12 months are recorded; we convert all these receipts to weekly quantities.

We derive a measure of spending that records all cash outlays by a household in a given period. Spending on items that tend to be purchased frequently (e.g. food, drink, household consumables, petrol) is taken from the diary and spending on items that are purchased less frequently (e.g. household furnishings and appliances, household bills, mortgage and rental payments) is taken from responses to the questionnaire. Our expenditure series can therefore be interpreted as average (weekly) expenditure in recent months.

Both income and spending are therefore measured over shorter periods than in the US Consumer Expenditure Survey (which measures income over the previous year and spending over the previous quarter) and over periods which are more similar to each other than in that survey but it is not the case that income and spending are collected for exactly the same period of calendar time, as occurs, for example, in the Canadian FAMEX/SHS surveys, where survey respondents are asked to recall all income and spending over the previous 12 months.

Our measures of both income and spending are at the level of the household and are adjusted for household composition using the modified OECD equivalence scale.[2] Using the Retail Prices Index, we express income and expenditure quantities in terms of real (December 2009) UK pounds per week. The LCFS contains some imputed data where satisfactory responses were not obtained from respondents. We do not include households that we know contain imputed data in our sample.

The quality of the LCFS data deserves some discussion. The response rate in the LCFS at the start of our data period was approximately 70% and has shown a steady decline to 52% in the most recent year of data (Office for National Statistics, 2014). These response rates are slightly lower than in the Family Resources Survey (62% in the most recent year of data: see Department for Work and Pensions, 2013*b*) – the survey which is used to calculate the official measures of income poverty in the UK. However, this difference comes not from large differences in the proportion refusing to participate but from a greater proportion of households in the intended sample with whom the survey team for the LCFS fail to make contact.

[2] This is usually expressed as giving a weight of 1 to the household head, of 0.5 to each additional adult member or child over the age of 14 and of 0.3 to each younger child. We follow the usual UK practice and re-base so that a two-adult childless household has a weight of 1, meaning that the scale becomes 0.67 for a single adult, 0.33 for each extra adult or child aged 14 or more and 0.2 for every child aged under 14.

One potential concern with the quality of the LCFS data is that, when implied economy-wide expenditure is calculated and compared to that reported in the UK National Accounts, the microdata accounts for only approximately two-thirds of aggregate expenditure totals (Deaton, 2005; Attanasio *et al.*, 2006; Crossley and O'Dea, 2010).[3] We therefore might be concerned that expenditure data suffer from their own under-reporting issues. The LCFS is the only nationally-representative survey that collects comprehensive data on household budgets and, as such, it is hard to validate household reports. However, Leicester and Oldfield (2009) compare the expenditure patterns reported in the LCFS with those observed in a commercial data set (the Kantar UK Worldpanel, previously known as the TNS UK Worldpanel). The latter data cover spending on food and groceries and are widely used in industry for market research. Leicester and Oldfield (2009) find that, for each type of family, levels of reported spending in the LCFS are higher than those in the Worldpanel.

We discuss the under-recording of expenditure in the microdata relative to the UK National Accounts further in Section 3. Although this phenomenon does not undermine the conclusion that we draw in this article (that the very high spending of those with the lowest incomes represents under-reporting of income to a greater extent than either over-reporting of spending or consumption smoothing), it is one of the reasons why, in our conclusion, we do not recommend a complete switch from using income to expenditure (or consumption) in the measurement of poverty in the UK.

2. Examining the 'Tick'

Figure 1 documented the clear pattern that households with the very lowest incomes have median reported spending associated with households with much higher incomes: that is, the relationship between reported income and median reported expenditure has a 'tick' pattern. But this pattern does not hold when the series are reversed: households with very low spending do seem to have very low income. As Meyer and Sullivan (2011*b*) argue, the sample of households with a low reported income in such an analysis has been, in part, selected on having negative measurement error, so we would expect average dissaving to be biased upwards. However, the levels of income for those with low recorded spending look more plausible than the levels of spending for those with low recorded income. This suggests that, amongst these households, there is more measurement error in income than spending.

Figure 2 shows, by plotting selected quantiles of expenditure by income for households in the bottom 10% of income, that the tick is not confined to the median (these profiles and those in the Figures that follow are generated using locally weighted quantile regressions). The graph also includes a 45-degree line, along which income is equal to expenditure.[4] A large majority of those in the bottom 10% of households are either dissaving or misreporting either of income or spending; among those reporting less than £75 per week (the bottom 1%), over 90% of households report values of

[3] This is also true of the US and its expenditure survey: see Bee *et al.* (2012) and Barrett *et al.* (2013).

[4] Tables A1 and A2 in the online Appendix show 95% confidence intervals around the estimates of quantiles shown in the Figures of this Section at each of the first, second and fifth percentiles of income.

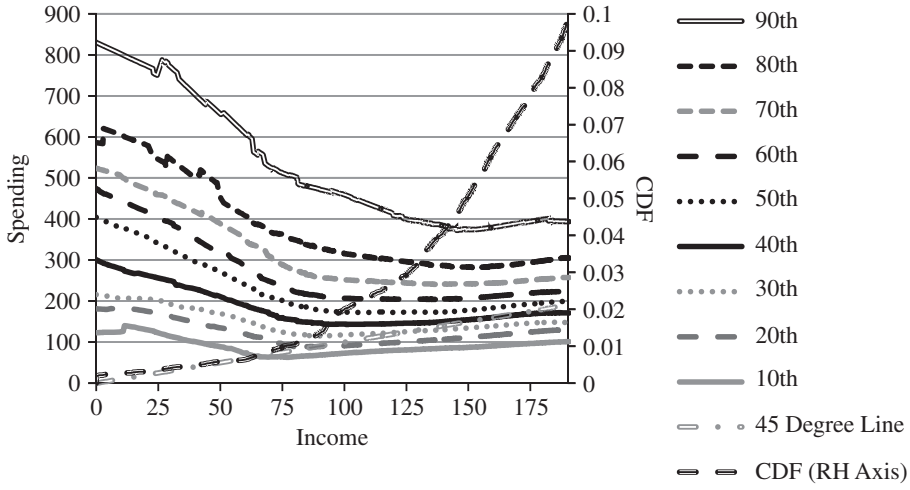

Fig. 2. *Quantiles of Expenditure by Income*
Notes. The lines in the graphs represent smoothed conditional quantiles and are generated by performing locally-weighted quantile regressions. Data described in note to Figure 1.

income and spending that are either misreported or consistent with dissaving. This graph also includes a Cumulative Distribution Function (CDF) of income (to be read off the right-hand axis). The profiles turn, so that the relationship between a particular quantile of expenditure and income becomes positive at points ranging from the first to the fifth percentile of income. The phenomenon we are discussing, therefore, is something particular to the very bottom of the income distribution. However, although those in that part of the income distribution on which we focus represent only a small share of the whole population, they represent a larger share of those in poverty (for example, on the relative income measure most commonly used in the UK, 16% of individuals were deemed to be in poverty in 2011–2: Department for Work and Pensions, 2013*c*).

The tick is a robust feature of the data and cannot easily be explained by demographic factors. Figure 3 divides the sample according to the age at which full-time education was completed (the household's education level is taken to be that of the most educated person in the household) and shows the profile of median reported spending conditional on reported income: the tick is apparent in each of three education categories. Figure 4 shows that the pattern appears when we condition on working-age family types (though it is not evident for pensioners, whose profiles we do not show in the Figure). Figure 5 shows median real expenditure conditional on real reported income for five-year periods starting in each of 1980, 1985, 1990, 1995, 2000 and 2005. In each of the periods, the tick is evident and its magnitude has been growing (i.e. median expenditure at the bottom of the reported income distribution has been growing). This implies some combination of worsening measurement error (perhaps related to the decline in survey participation) or an increasing ability to smooth consumption (and it is certainly the case that access to credit has become more widespread over this over 30-year period).

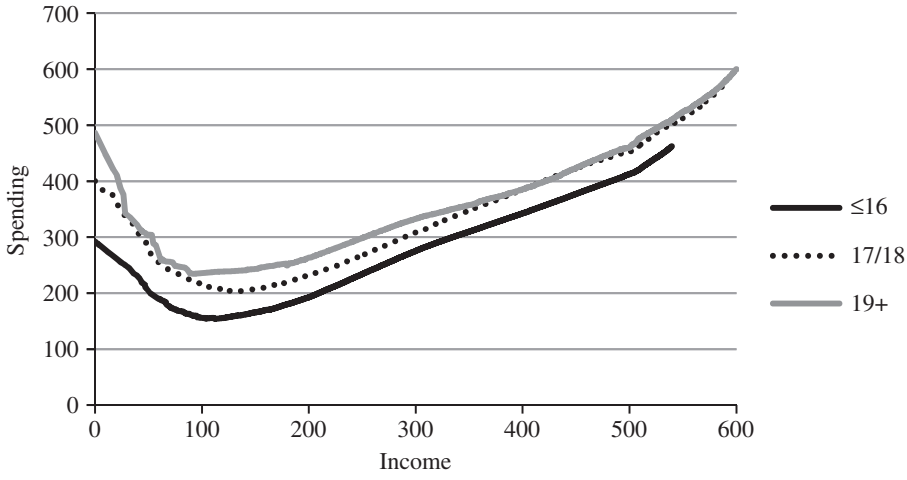

Fig. 3. *Median Expenditure by Age Completed Full-time Education*
Note. See notes to Figure 1.

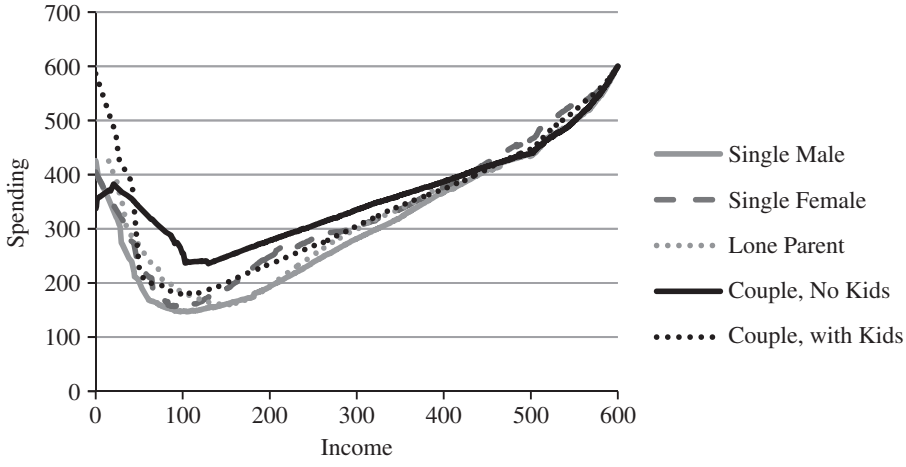

Fig. 4. *Median Expenditure by Family Type (Working-age Households)*
Note. See notes to Figure 1.

Figure 6 shows median expenditure across the income distribution split by work status. The Figure shows that the tick is evident whether the head of household is working, is out of work, or is self-employed. Very low-income employees are either reporting earnings below the minimum wage (and we show in the next Section that this is not uncommon) or are working few hours per week.[5] This Figure provides our

[5] Among individuals in our sample aged over 23 (lower minimum wages apply to younger individuals), aproximately 60% of employees reporting very low earnings (below £75 per week in gross earnings) are reporting below-minimum wage levels of earnings, with the remainder all reporting usual hours of work of not more than 20.

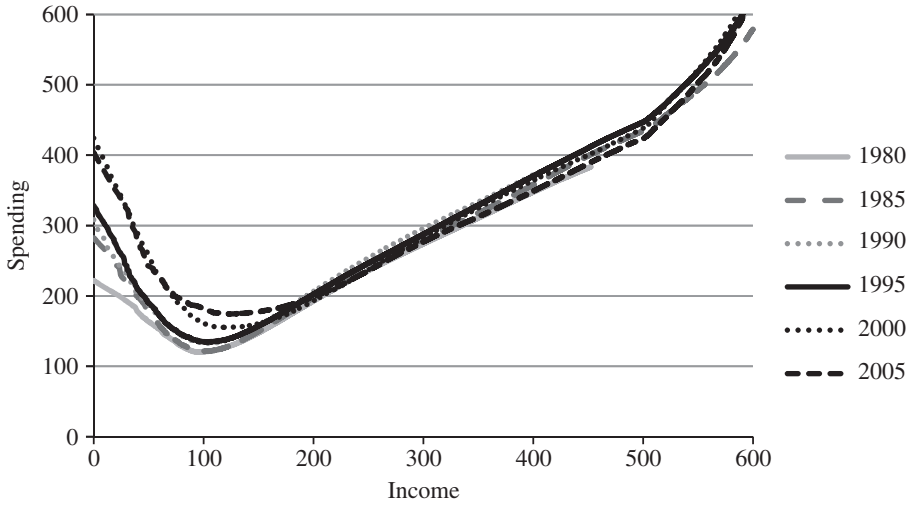

Fig. 5. *Median Expenditure by Period (Five Year Averages)*
Notes. Data are for the five years from the year shown in the legend. For other notes, see Figure 1.

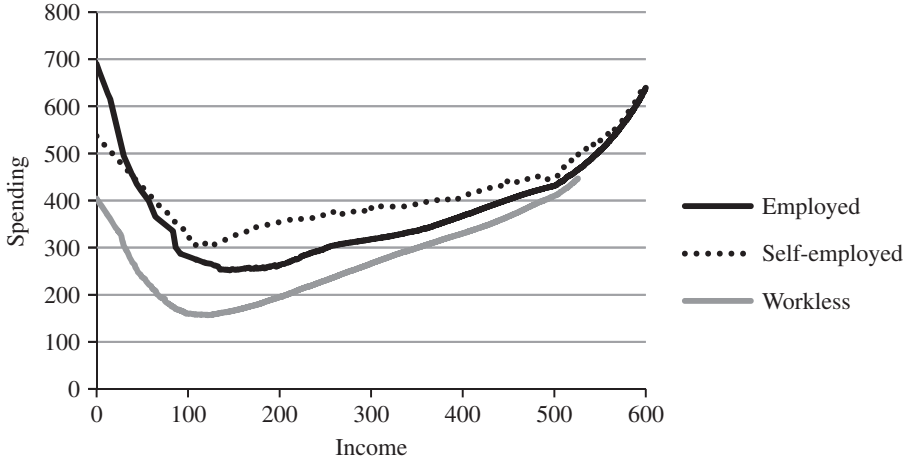

Fig. 6. *Median Expenditure by Employment Status*
Notes. See notes to Figure 1.

first piece of evidence that the tick is not explained solely by consumption smoothing: given that most large temporary fluctuations in income represent moves between employment states (Altonji *et al.*, 2013), we would expect that, if consumption smoothing explains all of the tick, it should be substantially less obvious in this Figure and median expenditure here would be monotonic in income. That the tick remains implies there is a role for measurement error. As an aside, the Figure also shows that the spending of very low-income workless households is lower than that of employed

and self-employed households with similarly low incomes, with the latter reporting substantially higher spending than either of the other two work groups even at middle incomes. The distinctive difference between the spending patterns of the employed and the self-employed is highly suggestive of a greater propensity for the self-employed to under-report their income; this difference has been used directly by Pissarides and Weber (1989) and Hurst *et al.* (2014) in the UK and US respectively to estimate the extent of that under-reporting.

Finally, we have checked and confirmed that where shown, the tick is statistically significant. We do this using the framework of Chetverikov (2012) which involves testing, non-parametrically, the monotonicity of the relationship between income and spending: we reject monotonicity of the relationship at all conventional levels, with p values always <0.01; further details of the test are given in the online Appendix. We have also checked the significance of using a simpler parametric test by running a median regression of reported spending on a set of indicator variables for £50 income bands: the reported expenditure of those in the lowest income group (£0–£50) is statistically significantly higher than each of the next seven income groups, representing households with equivalised reported income of less than £400.

3. Direct Evidence on Mismeasurement of Income and Expenditure at the Bottom End of the Distribution

In this Section, we provide direct evidence for whether the tick is consistent with under-recording of income, over-recording of spending, or consumption smoothing (in this case, running down assets or increasing debt). We present some evidence that can be used to support or refute these three (not-mutually-exclusive) hypotheses by comparing data from the LCFS to other sources and reviewing what is known about the asset holdings of those with low cash incomes and/or outlays in the UK.

3.1. *What Evidence is There that Income is Mismeasured?*

Table 1 shows the fraction of the amount that the UK government reported that it paid out in various state benefits (in 2009/10) that is recorded in both the LCFS and the other main survey used to measure UK's household income distribution, the Family Resources Survey. These 'coverage rates' are high for the two benefits which are universal or close to it: child benefit and the basic state pension.[6] However, coverage rates are much lower for the two main means-tested cash benefit programmes (68% for income support and pension credit and 50% for tax credits) and also low (52%) for the large category known as 'other non-contributory benefits', which mostly comprises benefits paid to disabled people or those requiring care in their own homes. We cannot tell with certainty which households in the survey are under-reporting income from state benefits: although we know which households are reporting receipt of income from cash benefits, we cannot know for sure which households should be

[6] It should be noted that there are legitimate reasons why the fraction captured should be below 100%, as some benefits are paid to people outside the sampling frame of the LCFS, either because they are outside the UK, or in the UK but not living in private households.

Table 1

Coverage in Living Costs and Food Survey (LCFS) and FRS of Spend on Cash Benefit Programmes in UK, 2009–10

	LCFS		FRS	
Tax credits	Coverage 50%	Total spend (£ bn/year) 21,270	Coverage 66%	Total spend (£ bn/year) 26,652
Non-means-tested benefits of which:	85%	116,890	86%	112,673
Retirement pension	95%	66,480		
Other non-contrib.	52%	27,970		
Child benefit	96%	11,880		
Incapacity benefit and ESA	74%	6,670		
Maternity pay	119%	1,900		
War pensions	33%	1,020		
Student support	236%	970		
Means-tested benefits of which:	76%	36,710	69%	46,355
Rent rebates and allowances	83%	18,930		
Income support	68%	16,580		
Jobseekers allowance	80%	1,200		

Notes. Total spend reports official estimates of spending on specific programmes, and coverage measures the fraction of this captured in the survey. Authors' calculation using: Tables 13 and 14 of Barnard *et al.* (2011) and previous versions, Department for Work and Pensions (2013*a*) and authors' analysis of Department for Work and Pensions et al. (2012) and previous years. Population grossing weights are used to give aggregate estimates. See online Appendix for details.

reporting income from cash benefits. But it is the means-tested benefits and tax credits, and benefits linked to ill health or disability, which have especially high under-reporting rates; better-off households should not, by design, be entitled to these means-tested benefits or income-related tax credits.

Figure 7 plots trends in the amount of 'missing' income from state benefits as a fraction of total household income in both surveys and shows that the importance of this missing benefit income has been rising gradually over the past decade even though the importance of benefit income overall as a share of household income has hardly changed.[7]

This issue of under-reporting has received more attention in the US than in the UK: see, for example, Marquis and Moore (1990) for an early example, Bound *et al.* (2001) for a general discussion of this phenomenon and Meyer *et al.* (2009) for a thorough investigation across 10 programmes and five US data sets. Marquis and Moore (1990) investigate the role played by forgetfulness, memory decay, confusion over names, telescoping, learning to under-report, interviewer effects and proxy responses, concluding that 'while these hypotheses give occasional insights into isolated error problems, they do not provide a fundamental understanding of the error dynamics [in the Survey of Income and Program Participation]'. Less is known about why UK surveys also show under-reporting: Lynn *et al.* (2012) was the first study of linked

[7] Manipulation of the same data (Tables 13 and 14 in Barnard *et al.*, 2011 and earlier editions) shows that benefit spending as a share of total household income has risen only very slightly, from 20% to 21%, over the same period.

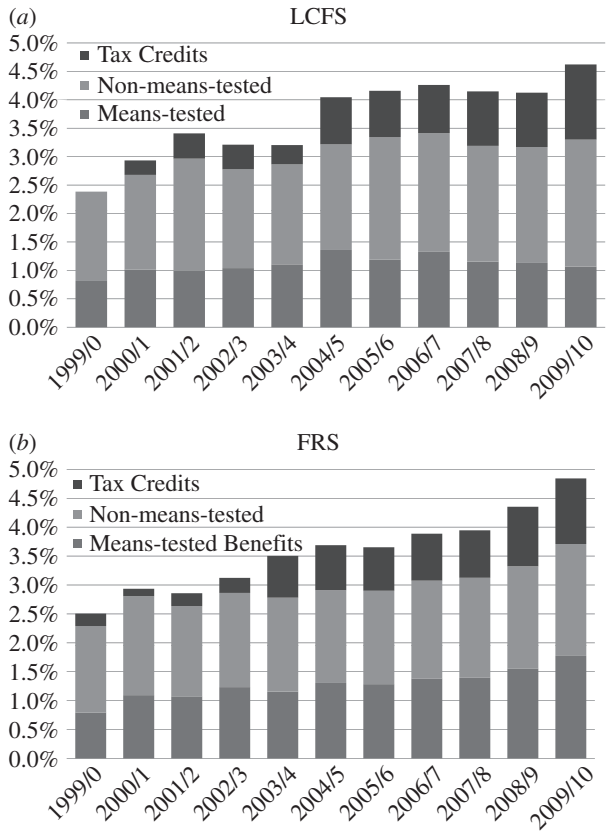

Fig. 7. *'Missing' Income from Cash Benefits as a Percentage of Total Household Income*
Notes. Authors' calculation using: Tables 13 and 14 of Barnard *et al.* (2011) and previous versions, Department for Work and Pensions (2013*a*) and authors' analysis of Department for Work and Pensions *et al.* (2012) and previous years. Confidence intervals for the estimates derived from the Family Resources Survey can be found in the online Appendix; the estimates derived from the Living Costs and Food Survey were taken from published aggregates, and no standard errors or confidence intervals are available.

survey-administrative data but it looks at a small number of cases in a different survey from those used in this article. Balarajan and Collins (2013) summarise the results of an unpublished study by UK government statisticians that compared administrative data linked to the Family Resources Survey; they conclude, based on interviews with survey interviewers and respondents who reported benefit income substantially at odds with administrative data, that part of the under-reporting is due to respondents' confusion over the names or purpose of benefit programmes and part is due to respondents simply not having the information needed to answer questions accurately. Given the aggregate data available to us, it is not clear whether this low (and declining) microdata coverage of benefit income is due to differential patterns of non-response to the surveys that is going uncorrected when grossing weights are calculated, incorrect item response amongst households who are genuinely receiving benefits, or due to

Table 2

Proportion of Employees Aged At Least 23 Reporting Implied Hourly Earnings Below the Minimum Wage

Year	Proportion
2001	6.2
2002	6.3
2003	6.5
2004	7.8
2005	6.8
2006	8.2
2007	8.2
2008	8.7
2009	10.5

Source. Authors' calculation from the LCFS Implied hourly earnings calculated by usual gross pay/usual hours.

households correctly reporting that they receive benefits but under-reporting the amounts.

We can also assess whether individuals are reporting wages that are below legal minima, as there has been a nationally-set, legally-binding, minimum wage in the UK since 1999. Table 2 shows the fraction of employees aged 23 and over (there are lower rates for younger workers) whose implicit hourly wage – usual weekly earnings divided by usual hours – is below the national minimum wage. The proportions (between 6% and 11%) are substantial, although we cannot tell whether earnings have been under-reported or hours worked over-reported (in reality, there are some exemptions but these apply to very few workers, and employers can count the value of some employer-provided benefits, although employer-provided health insurance is much less common in the UK than in the US.)

3.2. *What Evidence is There that Spending is Mismeasured?*

The phenomenon we investigate in this article – the divergence between income and expenditure for those with the lowest reported income – could, in principle, be driven by over-reporting of expenditure. We discussed the results of some studies in Section 1 that show that, at an aggregate level expenditure is under-reported relative to quantities recorded in the National Accounts. However, these analyses of aggregate spending do not rule out that expenditure is over-reported by some at the bottom of the income distribution. Direct evidence of misreporting at the bottom of the distribution has been obtained using the Canadian Survey of Family Expenditures (FAMEX/SHS). Until 2010, this survey made use of a 'balance edit', where participants are probed for more information if they report annual spending that is more than 20% different from the sum of net income and net dissaving (the latter of which is directly recorded in the Canadian data, unlike in the UK). Brzozowski and Crossley (2011) show that this balance edit substantially reduces apparent dissaving at the bottom of the reported income distribution; crucially, a comparison of responses before and after

the additional 'balance-edit' questions show that it corrects under-reporting of income to a greater extent than under-reporting of expenditure.

More tellingly, Brewer *et al.* (2009) show, when looking at households with children in four different household UK surveys, that the tick-shaped pattern between income and expenditure also exists when using many other measures of, or proxies for, living standards. This set of measures and proxies are collected in a variety of different ways and have different conceptual properties; this strongly suggests that over-recording of spending in the LCFS cannot be the sole explanation for the expenditure tick.

3.3. *What Do We Know about Asset Holdings, Debt and Saving Flows Amongst Households with Low Reported Income or Spending?*

The divergence between income and expenditure could, of course, be explained by net dissaving (i.e. running down assets or borrowing) among those with the lowest income. Here we look for direct evidence on whether such dissaving could be happening.

Unlike the Consumer Expenditure Survey in the US, the LCFS has no direct measure of net saving flows, so we cannot investigate in detail the extent to which households whose reported spending far exceeds their reported income are dissaving or borrowing.[8] Indeed, even considering other UK data sources, we have a remarkably imprecise impression of the stock of net assets held by households with a low reported income or spending, let alone the net saving flows. Information on net assets is available in two nationally-representative surveys. One of these, the relatively new Wealth and Assets Survey, unfortunately lacks comprehensive good quality data on household income (see Income Annex to Daffin (2009)) and therefore we do not use it (though more comprehensive income data is to be collected in future waves). The other of these, the British Household Panel Survey (BHPS) collects comprehensive data on income and also has some limited data on assets at regular intervals – most recently in 2005. Figure 8 plots the relationship between net liquid assets and reported household income in 2005 and shows that most of those with the lowest reported incomes have low levels of net assets. We infer from this that the majority of households whose low reported income is less than their cash outlays will not be running down savings (as they do not have any) but we cannot rule out that a minority could be.

If households have access to credit markets, they could, of course, be borrowing to cover excess of their expenditure over their income. Brewer *et al.* (2009) investigate the relationship between a number of proxies for living standards using two longitudinal surveys (the BHPS and the Families and Children Survey) and show that, even when income is measured using an average over three consecutive waves, those with the lowest reported income have living standards that are higher than might be expected. For the higher living standards to be supported by borrowing over the

[8] These data in the US allows Aguiar and Bils (2011), for example, to construct a measure of spending from the Consumer Expenditure Survey equal to reported income less reported saving.

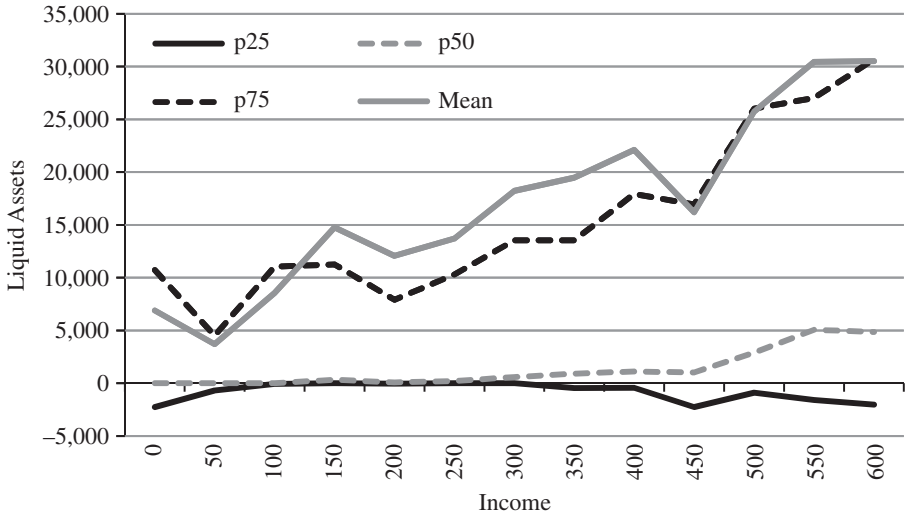

Fig. 8. *Distribution of Financial Assets by £50 Weekly Equivalised Income Bands*
Notes. Authors' calculations using British Household Panel Survey 2005. Financial assets comprise all liquid financial wealth (cash savings, stocks and shares) and are net of non-mortgage debt.

medium term (which is our interpretation of this three-year period), the level of access to credit markets would have to be implausible.

4. Can Consumption Smoothing Explain the Tick?

We now examine the tick using panel data on incomes and a theoretical model of consumption behaviour. To show that consumption smoothing cannot explain all the tick we formulate a model that both fits the panel income data *prima facie*, and has two features that induce agents at the bottom of the income distribution to have particularly high consumption: an income process that ascribes all ultra-low income realisations (i.e. those in the bottom 1%) to be temporary; and very lax borrowing constraints. We make these choices in order to bias the model in the direction of being able to reproduce a tick similar to that observed in the data without needing to appeal to measurement error, and we then show that the model cannot generate a tick of the size that we showed in Section 2.

Our focus here is the consumption of those with incomes in the bottom 1% of the population. The rationale for this is illustrated in Figure 9, which replots the data shown in Figure 5 with differently-scaled axes. This Figure shows, for each five-year period since 1980, median expenditure conditional on income as a proportion of median expenditure in the population graphed against income centile for centiles 1 to 8 of the income distribution. It underlines that the particularly high levels of median expenditure given income are most clearly evident among the bottom 1%, and that the median consumption of this group is as high or higher than the population – level median consumption. Although the phenomenon we are investigating – expenditure being far in excess of reported income – is also evident further up the income distribution, it is the consumption behaviour of those in the bottom 1% that is a

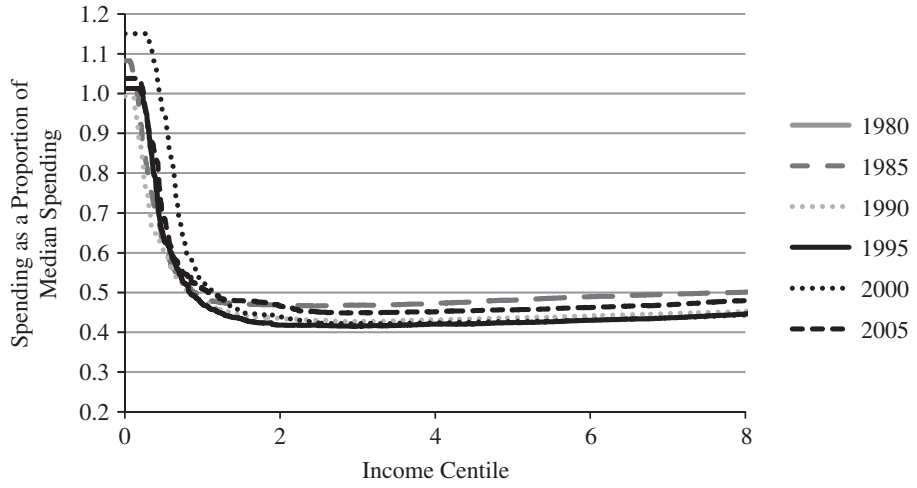

Fig. 9. *Median Normalised Expenditure by Centiles of Equivalised Real Income*
Notes. See notes to Figure 5.

particularly salient (and pervasive over time) feature of the data, so in this Section we focus on this bottom 1% as being 'unusual'.

4.1. *Empirical Evidence on Transitions Into and Out of the Bottom* 1%

Our income panel data come from the BHPS and our income concept is the same as that we use from the LCFS: net equivalised household income. The online Appendix gives full details of the treatment, including sample selection.

Our first task is to examine the empirical income transitions into and out of the bottom 1%. Table 3 shows these transitions to and from the remainder of the bottom decile and the other deciles. The data are split into yearly quantile bands conditional on decade-of-birth cohort and year. The Table shows that few households move far up and down the income distribution from year to year. However, it is clear that those with the very lowest measured incomes show different transitions from the next lowest-income groups: for example, those in the bottom 1% are much more likely to transit to each of the top five deciles than those in the 1–10% group. Therefore, on the working hypothesis that incomes are measured correctly, it seems *prima facie* possible that those in the bottom 1% could have higher permanent income than those in the 1–10% range. This result would yield a tick.

As a first simple way to interpret the data, therefore, we compute permanent incomes, a proxy for consumption, according to some simple assumptions: that income evolves as a Markov process according to this transition matrix, that households are infinitely lived and that they discount future income at 3% per year.[9] In addition, we ignore the contribution of wealth in this simple computation of permanent income. Finally, to fix the levels of income, we take current incomes in

[9] Results are robust to any sensible choice of discount rate.

Table 3

Empirical Transition Matrix for Income by Decile and for the Bottom Centile

Time $t-1$ \ t	Quantile Group										
	0–1	1–10	10–20	20–30	30–40	40–50	50–60	60–70	70–80	80–90	90–100
0–1	24.84	35.95	9.15	6.54	5.56	5.88	2.94	2.29	0.98	2.29	3.59
	(2.47)	(2.74)	(1.65)	(1.41)	(1.31)	(1.35)	(0.97)	(0.85)	(0.56)	(0.85)	(1.06)
1–10	3.72	47.98	23.91	9.75	5.34	3.16	1.82	1.17	1.25	0.89	1.01
	(0.38)	(1.00)	(0.86)	(0.60)	(0.45)	(0.35)	(0.27)	(0.22)	(0.22)	(0.19)	(0.20)
10–20	1.20	21.50	39.00	19.41	8.37	4.27	2.75	1.24	1.09	0.74	0.42
	(0.20)	(0.77)	(0.92)	(0.74)	(0.52)	(0.38)	(0.31)	(0.21)	(0.20)	(0.16)	(0.12)
20–30	0.63	8.08	18.69	34.14	20.44	8.61	4.25	1.97	1.48	0.91	0.81
	(0.15)	(0.51)	(0.73)	(0.89)	(0.76)	(0.53)	(0.38)	(0.26)	(0.23)	(0.18)	(0.17)
30–40	0.77	4.95	8.29	18.12	30.31	17.84	9.45	4.88	2.92	1.51	0.95
	(0.16)	(0.41)	(0.52)	(0.72)	(0.86)	(0.72)	(0.55)	(0.40)	(0.32)	(0.23)	(0.18)
40–50	0.53	2.91	5.15	8.27	17.29	31.38	18.64	8.41	4.08	2.20	1.14
	(0.14)	(0.32)	(0.42)	(0.52)	(0.71)	(0.87)	(0.73)	(0.52)	(0.37)	(0.28)	(0.20)
50–60	0.38	2.24	2.55	4.82	8.12	18.78	29.45	19.47	8.91	3.23	2.06
	(0.11)	(0.27)	(0.29)	(0.40)	(0.51)	(0.72)	(0.85)	(0.73)	(0.53)	(0.33)	(0.26)
60–70	0.48	1.58	1.75	2.77	4.62	7.74	19.18	32.88	18.77	7.60	2.64
	(0.13)	(0.23)	(0.24)	(0.30)	(0.39)	(0.49)	(0.73)	(0.87)	(0.72)	(0.49)	(0.30)
70–80	0.38	1.36	1.39	1.64	2.72	3.80	8.01	18.99	36.97	20.07	4.67
	(0.12)	(0.22)	(0.22)	(0.24)	(0.30)	(0.36)	(0.51)	(0.73)	(0.90)	(0.75)	(0.39)
80–90	0.37	0.92	0.71	0.88	1.77	2.55	3.84	6.53	18.44	45.83	18.14
	(0.11)	(0.18)	(0.16)	(0.17)	(0.24)	(0.29)	(0.35)	(0.46)	(0.72)	(0.92)	(0.71)
90–100	0.43	1.22	1.04	1.04	1.25	1.22	2.08	3.44	4.73	15.76	67.80
	(0.12)	(0.21)	(0.19)	(0.19)	(0.21)	(0.21)	(0.27)	(0.34)	(0.40)	(0.69)	(0.88)

Notes. Rows sum to 100%. Standard errors in parentheses and computed using $\sqrt{p(1-p)/N}$, where p is the relevant probability. See text and online Appendix A.4.1. for details on the computation. *Source*. BHPS.

2005 and 2006 for households with heads born between 1950 and 1970. With these assumptions, we can compute permanent income as the annuity income from total discounted life-time income. We plot this relationship in Figure 10. The Figure does display a tick but it is very small; on this first simple interpretation, therefore, the tick in the data cannot be reconciled with consumption-smoothing behaviour.

4.2. *An Income and Consumption Model*

In order to interpret the data more credibly, we now build a more realistic model of the income process. We start with a standard model, in which log income is given by the sum of a persistent component and a transitory component. Such an income process is known to describe income dynamics well across the bulk of the income distribution (see Meghir and Pistaferri (2011) for a review of the literature in this area). The model is given by the following:

$$\begin{aligned} v_t &= \rho v_{t-1} + \zeta_t \\ \tilde{y}_t &= v_t + \epsilon_t \end{aligned} \tag{1}$$

$$\begin{aligned} \zeta_t &\sim \mathrm{N}(0, \sigma_\zeta^2) \\ \epsilon_t &\sim \mathrm{N}(0, \sigma_\epsilon^2) \end{aligned} \tag{2}$$

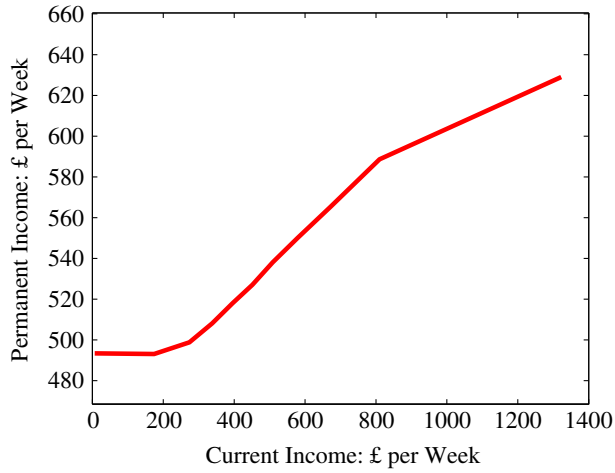

Fig. 10. *Relationship Between Current and Permanent Income Implied by the First-order Transition Matrix*
Notes. Permanent income here denotes discounted future income and excludes assets. See text for details.

such that $\tilde{y}_t$ is 'latent' log income for the household at time t, composed of a highly persistent component, v_t, and a transitory shock ϵ_t. The persistent component follows an AR(1) process with innovation ζ_t which is distributed normally. The parameter of persistence is $\rho \in [0,1)$. The transitory shock is *iid* and is also distributed normally.

This model as it stands is unsatisfactory because it does not pick up the salient transitions for the bottom 1%. In fact, it is easy to show that the average life-time income is monotonically increasing in observed current income in this model, something that would rule out consumption smoothing as an explanation for the tick, *a priori*. We therefore add a specific Markov process for this bottom 1%; this is similar to the approach taken by Jenkins (2011) and Cappellari and Jenkins (2004). Formally:

$$\begin{gathered} y_t = \mu + (1-\theta_t)\max(\tilde{y}_t, u) + \theta_t \underline{y} \\ \theta_t \in \{0,1\} \\ \Pr(\theta_t = j | \theta_{t-1} = k) = \pi_{kj}(\tilde{y}_{t-1}), \end{gathered} \tag{3}$$

where y_t is observed log income, composed of mean income μ, and either (truncated) latent income or a low income draw $\underline{y} < 0$ that would put the household in the bottom 1%. We set $\underline{y}$ such that $\mu + \underline{y}$ is the average of the bottom 1% of equivalised incomes observed in the data. The random draw is governed by the stationary Markov process θ_t. We term $\theta_t = 1$ a poverty shock (as a shorthand for ultra-low realisation of income).[10] When $\theta_t = 0$, the household receives $\mu + \max(\tilde{y}_t, u)$, where $\max(\tilde{y}_t, u)$ is a truncated version of latent income with $u > \underline{y}$. (This truncation has almost no effect on the analysis but is necessary to ensure that non-poverty income is always larger than poverty income, something which is not otherwise guaranteed as the underlying shocks have an unbounded lower tail.) By making ultra-low realisations of income a purely

[10] This should not be understood as poverty in the sense used in UK policy, where poverty is most often thought of as having income below 60% of median income.

temporary phenomenon, we make smoothing consumption over these realisations easy for agents, and so this feature of our model favours the finding of a tick in consumption.

The transition probability from state k to state j (where k and j are binaries indicating poverty status) is given by $\pi_{kj}(\tilde{y}_{t-1})$. We allow the transition probabilities to depend on latent income because it is noticeable from Table 3 that the entry and exit rates are not uniform across the top 99%; note that our specification means that each household's probability of entering and exiting poverty is evolving constantly as their latent income changes.

To complete the set-up, we specify a rule for consumption and saving. In the spirit of Carroll (2009) and Gourinchas and Parker (2002), we let consumption evolve as the solution to the following optimisation problem:

$$V(A_t, v_t, \epsilon_t, \theta_t) = \max_{C_t} C_t^{1-\gamma}/(1-\gamma) + \beta \mathbb{E}_t\{V[(A_t - C_t)R + Y_{t+1}, v_{t+1}, \epsilon_{t+1}, \theta_{t+1}]\}, \quad (4)$$

where households choose consumption C_t to maximise life-time (indirect) utility $V()$ given assets A_t, underlying persistent income v_t, transitory shock ϵ_t and poverty state θ_t. On the right hand side, $Y_t = \exp(y_t)$ is household income, β is the discount factor, and R is one plus the interest rate.[11] When taking the model to data, we use equivalised income as the input, so the model is best thought of as capturing the consumption decisions of an equivalised household. Per-period utility is given by the standard isoelastic function. This utility specification implies that households accumulate precautionary saving to protect against negative income shocks, including the poverty shock. It also implies that households have a finite target wealth level as long as $\beta R < 1$: whenever assets get too high, households will want to dissave. Finally, $\mathbb{E}_t$ denotes the expectations operator at time t: households know the income model and form rational expectations over the stochastic processes for the evolution of v_t, ϵ_t and θ_t given by (1), (2) and (3).

A critical factor in the ability of households to smooth consumption is the tightness of borrowing constraints. We therefore investigate the model with two extremes of constraints. First, to model a lax lending environment, we allow 'natural' borrowing constraints: households are allowed to borrow up to the maximum they would be able to pay back if they receive the minimum possible income draw in every future period. At the specified interest rate, the long-term debt limit is 11.9 units of income (where 1 is the average income across society). Second, at the other extreme, we impose strict no-borrowing constraints. We recognise that each regime is unrealistic but together they capture the bounds of plausible credit limits, with the no-borrowing case likely to be closer to the truth.

The choice of consumption model warrants further discussion. The model is in the spirit of, for example, Carroll (2009), who looks at the consumption responses to income fluctuations in a dynamic programming model. As such, we focus on income

[11] As the model is written, the income concept, Y_t, is more commonly interpreted as labour income rather than total income. Total income would be captured by $Y_t + (A_t - C_t)r$ where $r = R - 1$. However, we interpret Y_t here as total income, because $(A_t - C_t)r$ likely also reflects capital gains, which are not included in the data. The distinction between income concepts is slight and we ignore it for the purposes of exposition. See the online Appendix for more discussion.

dynamics as a possible explanation of the tick and abstract from other possibly relevant factors, such as non-separabilities with leisure. These non-separabilities might be particularly relevant if low income is caused by unemployment. For example, if those at the bottom have high leisure, the tick could be explained by leisure and consumption being complements. However, most studies (Browning and Meghir, 1991; Blundell *et al.*, 1994) find that consumption and leisure are substitutes. Further, Aguiar and Hurst (2005) show that by switching from market to home production, those who acquire more leisure can obtain a given quantity of consumption with less expenditure. Therefore, if the ultra-low realisations of income are mainly due to unemployment shocks, then introducing these unemployment states explicitly into the model would lower the tick and strengthen our argument. Finally, note that we use an infinite horizon and so abstract from life-cycle considerations such as income growth and changes to family size. We do this because an infinite horizon is more likely to generate a tick and because the tick is observed in the data across the whole of working life and for different groups.

As a final word on methodology, we note that we are using panel data on household incomes which are likely to be measured with just as much error as the LCFS possibly is. To emphasise, our approach is to assume as a falsifiable hypothesis that income is measured correctly and to see if the tick can be reproduced by consumption smoothing alone.

4.3. *Estimation and Calibration*

We fit the model to the data in two stages. First, we estimate the parameters of the income process using the transition matrix. Second, we solve the consumption model using numerical techniques and calibrate it using relevant moments. Full details are given in the online Appendix; here we discuss the important and intuitive details.

To fit the income process in the main analysis, we merge the decile groups into courser quantile groups. The courser transition matrix is given on the left hand side of Table 4. We adopt a simple and transparent approach to estimation. First, we take transition probabilities for the poverty shock directly from the matrix. The probabil-

Table 4

Empirical and Fitted Transition Matrices

	Empirical						Fitted				
$t-1$ \ t	0–1	1–10	10–20	20–50	50–100		0–1	1–10	10–20	20–50	50–100
0–1	24.8	35.9	9.2	18.0	12.1	0–1	24.8	17.3	11.8	22.8	23.2
1–10	3.7	48.0	23.9	18.2	6.1	1–10	3.7	64.4	21.9	9.8	0.2
10–20	1.2	21.5	39.0	32.1	6.2	10–20	1.2	22.5	33.9	39.9	2.5
20–50	0.6	5.3	10.7	62.2	21.1	20–50	0.6	3.4	13.4	58.2	24.4
50–100	0.4	1.5	1.5	12.8	83.8	50–100	0.4	0.0	0.5	14.7	84.4

Notes. Column and row headers give quantile groups. See online Appendix A.4.2 for more details on estimation.

ities of entering the bottom 1% are given by the lower four entries in the first column. Second, we impose that the probability of remaining in the bottom 1% is constant across the latent income types and given by the top left-hand entry: 0.248. All that remains is to estimate the parameters of the underlying latent income process: ρ, σ_ζ^2 and σ_ϵ^2. These parameters are not separately identified from the first-order transition matrix. We therefore fix $\rho = 0.95$, in the range of conventional estimates, as discussed in Meghir and Pistaferri (2011). This leaves σ_ζ^2 and σ_ϵ^2, the variances of persistent and transitory shocks, to identify. We identify the relative magnitudes of these shocks as follows: letting $\sigma_v^2 = \sigma_\zeta^2/(1 - \rho^2)$ be the long-run, stationary variance of persistent incomes, note that the total variance of (latent) incomes is given by the sum of σ_v^2 and σ_ϵ^2. We then use the empirical matrix to estimate, by minimum distance, the share of persistent variance in this total variance: we estimate this to be 0.93 (standard error of 0.006). We later calibrate the overall variance of income shocks when we calibrate the consumption model. This overall variance and the shares of persistent and transitory components yield us σ_ζ^2 and σ_ϵ^2.

The fitted transition matrix is given on the right-hand side of Table 4. It lines up with the empirical matrix reasonably well, given that the right hand four columns are estimated using a single free parameter. If anything, the transitions from the bottom 1% to the top half of the income distribution are more frequent in the fitted model than in the empirical data (and especially so relative to the transitions of the 1–10% group to the top half of the income distribution). To the extent that, by overestimating the latent income (and therefore permanent income) of those in the bottom 1% of the income distribution, we are overestimating their consumption, this small divergence between the empirical transition matrix and our fitted matrix will magnify the tick produced by our simulations.

Having estimated the income process from the transition matrix, we then fit the consumption model by calibrating the discount factor β and the variance of incomes. We do this by fitting to the average wealth-to-income ratio and by fitting to the average consumption of the 1–10% group. The average wealth-to-income ratio is 3.1, a standard target in the macroeconomics literature (Storesletten *et al.*, 2004). The consumption of the 1–10% group is approximately 45% of median consumption (see Figure 9). Full details of the model solution, parametrisation, calibration and simulation are in the online Appendix. Our results are robust to other sensible calibrations, some of which are discussed in the online Appendix, and to changing the variance of incomes. We discuss the model fit when discussing the results, to which we now turn.

4.4. *Results*

Figure 11 shows median simulated consumption by income centile for the models with and without borrowing constraints (the income distribution is the same in both specifications). There is a tick in consumption under both regimes because the bottom 1% of the income distribution contains a broader mix of skill types than the 1–10% group. However, if we take as a benchmark from the data that those at the bottom have consumption equal to median consumption, then even the economy without credit constraints cannot reproduce all the tick. As a rough calculation using Figure 9 as a

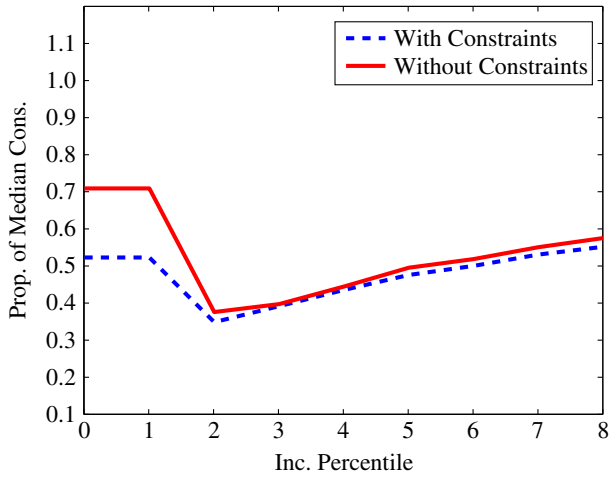

Fig. 11. *Simulated Median Consumption Across the Income Distribution*

basis, to reproduce the tick in full, the bottom 1% would need to consume just under three times as much as the the second percentile (which here consumes just over 35% of median consumption). Actually, in our most generous model, it consumes only around twice as much, meaning that it explains around 50% of the tick. The economy with tight borrowing constraints explains even less of the tick, with the same calculation revealing that such a regime can explain only around 30% of the tick.

Neither model of borrowing constraints matches the reality of credit markets perfectly but the model with borrowing constraints is likely to be closer to reality (that a household could borrow almost 12 times its earnings potential without collateral is implausible). Furthermore, the distribution of wealth in the model with constraints much more closely matches the empirical distribution than does the model without constraints: Figure 12 shows the simulated wealth distribution in both economies, together with mean financial wealth-to-income ratio calculated using the same data that is used from Figure 8.

We conclude that consumption smoothing cannot account for all the tick in either model. This is driven by two factors. First, the bottom 1% does not contain, on average, households of median productivity level; second, households cannot perfectly smooth their consumption through a period of low income, even when we specify that the shocks that induce such low income are temporary.

5. Conclusion

This article analyses a puzzling feature of the joint distribution of income and expenditure. Although median income is monotonically increasing in household expenditure, median expenditure is higher for those with very low income than for those with moderately low income. In fact, those at the very bottom of the income distribution have expenditure equal to the population-level median expenditure. In short, the graph of median expenditure by reported income maps out a 'tick'. This

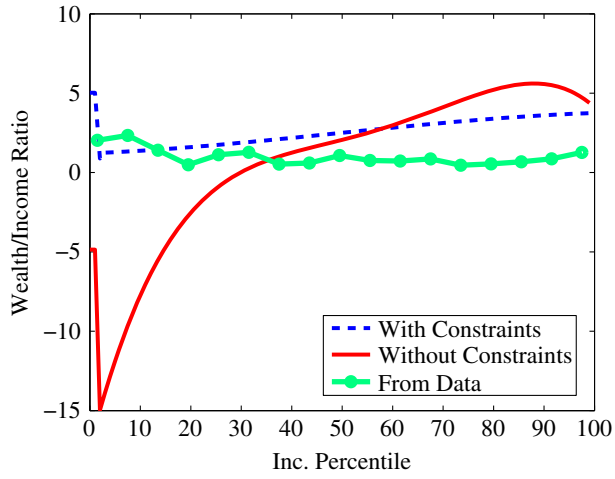

Fig. 12. *Simulated and Empirical Mean Wealth/Income Ratio Across the Income Distribution*
Notes. Data from BHPS taken from Figure 8. Wealth data are financial assets, and do not include housing and pension wealth.

pattern could be explained by any combination of under-reporting of income, over-reporting of expenditure, or that households smooth consumption over time, and this article investigates what roles each of these candidates has to play. Of the three (non-mutually exclusive) hypotheses which could explain the tick, we argue that under-reporting of income plays the most important role.

Our results beg the question of whether expenditure levels, rather than income, should be the criterion which governments use to determine eligibility for benefits or to measure poverty. There are two reasons why we would not recommend the use of expenditure to determine who should receive state benefits or be enrolled in government programmes. First, the collection of accurate expenditure data is substantially more difficult and more expensive than income data. Second, it is possible that our finding – that income is more vulnerable to under-reporting than expenditure – would not survive if the latter were made the instrument for delivering poor relief. Indeed, one candidate explanation for our results is that, even though the statistical authorities do not share households' responses with other government authorities, households may feel that their responses to the survey may lead them to have higher tax bills or reduced entitlement to benefits.

But there is a stronger case for the use of expenditure data in the measurement of national poverty trends, because our results strongly suggest that a fraction of those with reported low income are unlikely to have genuinely low resources. However, we do not suggest, in the UK at least, that poverty should be measured with expenditure data alone. One obstacle (that can be overcome, but at a financial cost) is the relatively small sample size of the UK's LCFS, which at present makes it hard to judge whether poverty trends are statistically significant or to measure poverty reliably amongst certain sub-groups. A less readily remediable problem is the concerning (and increasing) divergence seen in the UK and elsewhere between aggregate consumption

implied by household budget surveys and that reported as part of National Accounts. Until it is understood to what extent (if at all) this phenomenon is related to the way that households in poverty report their spending patterns, we are not prepared to recommend the exclusive use of expenditure data for poverty measurement. Indeed, the contribution that expenditure data can make to the measurement of poverty provides an important reason (over and above the many others that have been discussed) to understand and resolve the tension between micro and aggregate data on household spending better. In spite of this, the monitoring of expenditure (or rather, consumption) poverty as well as income poverty, or their use together, perhaps with other measures of deprivation in a multidimensional index, would represent an improvement on the use of income alone, and especially when measuring trends in extreme poverty.

Finally, our results imply the need for careful and thorough consideration of what aspect of the survey design causes this mismeasurement of income, and how survey instruments can better capture income at the bottom of the distribution.

University of Essex Institute for Fiscal Studies
University of Essex
Institute for Fiscal Studies

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

Appendix A. Supplementary Details.
Data S1.

References

Aguiar, M.A. and Bils, M. (2011). 'Has consumption inequality mirrored income inequality?', Working Paper No. 16807, NBER.
Aguiar, M. and Hurst, E. (2005). 'Consumption versus expenditure', *Journal of Political Economy*, vol. 113(5), pp. 919–48.
Altonji, J., Smith, A. and Vidangos, I. (2013). 'Modeling earnings dynamics', *Econometrica*, vol. 81(4), pp. 1395–454.
Attanasio, O.P., Battistin, E. and Leicester, A. (2006). 'From micro to macro, from poor to rich: consumption and income in the UK and the US', mimeo, National Poverty Centre.
Balarajan, M. and Collins, D. (2013). 'A review of questions asked about receipt of state benefits on the family resources survey', Working Paper Series No. 115, Department for Work and Pensions.
Barnard, A., Howell, S. and Smith, R. (2011). 'The effects of taxes and benefits on household income, 2009/10. Further analysis and methodology', *Newport: Office for National Statistics*. http://www.ons.gov.uk/ons/rel/household-income/the-effects-of-taxes-and-benefits-on-household-income/2009-2010/the-effects-of-taxes-and-benefits-on-household-income-2009-10.pdf (last accessed: 17 November 2015).
Barrett, G., Levell, P. and Milligan, K. (2013). 'A comparison of micro and macro expenditure measures across countries using differing survey methods', Working Paper No. 19544, NBER.
Bee, C., Meyer, B. and Sullivan, J. (2012). 'The validity of consumption data: are the consumer expenditure interview and diary surveys informative?', Working Paper No. 18308, NBER.
Blundell, R., Browning, M. and Meghir, C. (1994). 'Consumer demand and the life-cycle allocation of household expenditures', *Review of Economic Studies*, vol. 61(1), pp. 57–80.
Bound, J., Brown, C. and Mathiowetz, N. (2001). 'Measurement error in survey data', in (J. Heckman and E. Leamer, eds.), *Handbook of Econometrics*, chap. 59, pp. 3705–843, Oxford: Elsevier.
Brewer, M., Browne, J. and Joyce, R. (2011). 'Child and working-age poverty from 2010 to 2020', Commentary No. 121, Institute for Fiscal Studies.

Brewer, M., Goodman, A. and Leicester, A. (2006). 'Household spending in Britain: what can it teach us about poverty?', mimeo, Joseph Rowntree Foundation.
Brewer, M., O'Dea, C., Paull, G. and Sibieta, L. (2009). 'The living standards of families with children reporting low incomes', Research Report No. 577, Department for Work and Pensions.
Browning, M. and Meghir, C. (1991). 'The effects of male and female labor supply on commodity demands', *Econometrica*, vol. 59(4), pp. 925–51.
Brzozowski, M. and Crossley, T.F. (2011). 'Viewpoint: Measuring the well-being of the poor with income or consumption: a Canadian perspective', *Canadian Journal of Economics*, vol. 44(1), pp. 88–106.
Cappellari, L. and Jenkins, S. (2004). 'Modelling low income transitions', *Journal of Applied Econometrics*, vol. 19(5), pp. 593–610.
Carroll, C. (2009). 'Precautionary saving and the marginal propensity to consume out of permanent income', *Journal of Monetary Economics*, vol. 56(6), pp. 780–90.
Chetverikov, D. (2012). 'Testing regression monotonicity in econometric models', Working Paper No. CWP35/12, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
Crossley, T. and O'Dea, C. (2010). *The Wealth and Saving of UK Families on the Eve of the Crisis*, London: Institute for Fiscal Studies.
Cutler, D.M. and Katz, L.F. (1992). 'Rising inequality? Changes in the distribution of income and consumption in the 1980's', *American Economic Review*, vol. 82(2), pp. 546–51.
Daffin, C. (2009). *Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2006/2008*, Newport: Office for National Statistics.
Deaton, A.S. (2005). 'Measuring poverty in a growing world (or measuring growth in a poor world)', *Review of Economics and Statistics*, vol. 87(1), pp. 1–19.
Department for Work and Pensions (2013*a*). 'Benefit expenditure and caseload tables 2013: outturn and forecast autumn statement 2013', mimeo, Department for Work and Pensions.
Department for Work and Pensions (2013*b*). *Family Resources Survey 2011/12*, London: TSO.
Department for Work and Pensions (2013*c*). 'Households below average income – an analysis of the income distribution 1994/95 to 2011/12', Technical Report, London: TSO.
Department for Work and Pensions, National Centre for Social Research and Office for National Statistics (2012). 'Family Resources Survey, 2009–2010 [computer file]', UK Data Archive [distributor].
Gourinchas, P.O. and Parker, J.A. (2002). 'Consumption over the life cycle', *Econometrica*, vol. 70(1), pp. 47–89.
Hurst, E., Li, G. and Pugsley, B. (2014). 'Are household surveys like tax forms? Evidence from income underreporting of the self-employed', *Review of Economics and Statistics*, vol. 96(1), pp. 19–33.
Jenkins, S.P. (2011). *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*, Oxford: Oxford University Press.
Leicester, A. and Oldfield, Z. (2009). 'Using scanner technology to collect expenditure data', *Fiscal Studies*, vol. 30(Special Issue on Measuring Consumption and Saving), pp. 309–37.
Lynn, P., Jäckle, A., Jenkins, S.P. and Sala, E. (2012). 'The impact of questioning method on measurement error in panel survey measures of benefit receipt: evidence from a validation study', *Journal of the Royal Statistical Society Series A*, vol. 175(1), pp. 289–308.
Marquis, K. and Moore, J. (1990). 'Measurement errors in SIPP program reports', Survey of Income and Program Participation Working Paper Series No. 9008, Bureau of the Census.
Meghir, C. and Pistaferri, L. (2011). 'Earnings, consumption and life cycle choices', in D. Card and O. Ashenfelter (eds.), *Handbook of Labor Economics*, vol. 4 Part B, pp. 7730–854, Oxford: Elsevier.
Meyer, B.D. and Sullivan, J.X. (2003). 'Measuring the well-being of the poor using income and consumption', *Journal of Human Resources*, vol. 38, pp. 1180–220.
Meyer, B.D. and Sullivan, J.X. (2008). 'Changes in the consumption, income, and well-being of single mother headed families', *American Economic Review*, vol. 98(5), pp. 2221–41.
Meyer, B.D. and Sullivan, J.X. (2011*a*). 'Consumption and income poverty over the business cycle', Working Paper No. 16751, NBER.
Meyer, B.D. and Sullivan, J.X. (2011*b*). 'Viewpoint: Further results on measuring the well-being of the poor using income and consumption', *Canadian Journal of Economics*, vol. 44(1), pp. 52–87.
Meyer, B.D., Mok, W.K.C. and Sullivan, J.X. (2009). 'The under-reporting of transfers in household surveys: its nature and consequences', Working Paper No. 15181, NBER.
Office for National Statistics (2014). 'Living costs and food survey technical report for survey year: January to December 2012', Technical Report, London: TSO.
Pissarides, C.A. and Weber, G. (1989). 'An expenditure-based estimate of Britain's black economy', *Journal of Public Economics*, vol. 39(1), pp. 17–32.
Poterba, J.M. (1989). 'Lifetime incidence and the distributional burden of excise taxes', *American Economic Review*, vol. 79(2), pp. 325–30.
Sabelhaus, J. and Groen, J. (2000). 'Can permanent-income theory explain cross-sectional consumption patterns?', *Review of Economics and Statistics*, vol. 82(3), pp. 431–8.

Saunders, P., Bradshaw, J. and Hirst, M. (2002). 'Using household expenditure to develop an income poverty line', *Social Policy and Administration*, vol. 36(3), pp. 217–34.
Slesnick, D.T. (1993). 'Gaining ground: Poverty in the postwar United States', *Journal of Political Economy*, vol. 101(1), pp. 1–38.
Storesletten, K., Telmer, C. and Yaron, A. (2004). 'Consumption and risk sharing over the life cycle', *Journal of Monetary Economics*, vol. 51(3), pp. 609–33.