ISSN 1755-5361

University of Essex

Department of Economics

Discussion Paper Series

No. 736 April 2013

Why are households that report the lowest incomes so well-off?

Mike Brewer, Ben Etheridge and Cormac O'Dea

Why are households that report the lowest incomes so well-off?*

Mike Brewer†
Ben Etheridge‡
Cormac O'Dea§

12th April 2013

Abstract

Using data from the Living Costs and Food Survey in the UK over 1978-2009 we document that households with extremely low measured income (below 10% of median income) on average spend much more than those with merely moderately low income (those below 50% of median income): in short, the graph of median expenditure against income contains a sharp non-monotonicity (or 'tick'). We show that this tick appears, to a greater or lesser extent, over the whole period 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. Finally, and whatever the reason for the tick, we document that low consumption is better correlated with other measures of living standards than having low income.

JEL Classification: D31, I32

Keywords: Consumption, Measuring living standards, Poverty

*This work builds on several past projects based at the Institute for Fiscal Studies, and we are very grateful to our current and former colleagues. We are also very grateful to Thomas Crossley for advice and encouragement, 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 Institute for Social and Economic Research (Brewer). Material from Department for Work and Pensions Research Report 577 is Crown Copyright and reproduced under the terms of the Open Government Licence v1.0. 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 (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. All errors remain the responsibility of the authors.

†University of Essex and Institute for Fiscal Studies
‡University of Essex
§Institute for Fiscal Studies and University College London

1 Introduction

Economists have consistently argued, on both theoretical and practical grounds, that consumption provides a better measure of poverty than income. On a theoretical level, consumption in any one time period is more closely aligned with life-time utility than is income.[1] In particular, households with high life-time resources can borrow to smooth consumption when income is temporarily low. On a practical level, Meyer and Sullivan (2003) argue that, in US surveys at least, income is more likely than expenditure[2] to be under-reported by those with the lowest resources.

Nevertheless, there are obvious reasons why poverty is generally measured using income. Principally, governments dispense benefits according to income rather than expenditure, which 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 2020-21; see Brewer et al. (2011)), all of which define poverty in terms of a low household income.[3] 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 or the practical reasons listed above or some combination of the two? This question is important because these reasons 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 and that more effort should be put into survey design and implementation. It further implies 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 surveys contains vital information on economic behaviour, and shows how important stocks of assets, credit markets and other smoothing mechanisms are to household welfare.

This paper attacks this question by analysing a puzzling, and, as we show, pervasive, feature of the joint distribution of income and expenditure at the bottom end of the income distribution. The puzzle is documented in Figure 1 using data from the UK Living Costs and Food Survey (LCFS). While median

[1] See Poterba (1989), Cutler and Katz (1992) and Slesnick (1993)

[2] It is important at the outset to distinguish between expenditure and consumption. The former is cash spent on goods and services, the latter will additionally include the flow of consumption obtained from durables but will not include cash spent on such durables. These concepts play important separate roles in our analysis.

[3] Two of the four measures compare the income of households containing children to poverty lines which are fixed in real terms, and which move in line with median income; the third is defined in terms if persistently low relative income, and the fourth is defined in terms of having a low relative income and being materially deprived according to an index. See http://www.legislation.gov.uk/ukpga/2010/9/contents for details.

income is monotonically increasing in household expenditure, median expenditure is higher for those with *very* low income (below £50 per week or 10% of median income) than for those with moderately low income (below £250 per week or 50% of median income). In fact, those at the bottom of the income distribution have median expenditure equal to overall median expenditure. In short, the graph of median expenditure by reported income contains a sharp non-monotonicity; that is it maps out a 'tick' (or for North Americans: a 'check'). This tick has existed to a greater or lesser extent both 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 paper investigates what roles each of these candidates has to play.

Figure 1: Median expenditure by income, and median income by expenditure

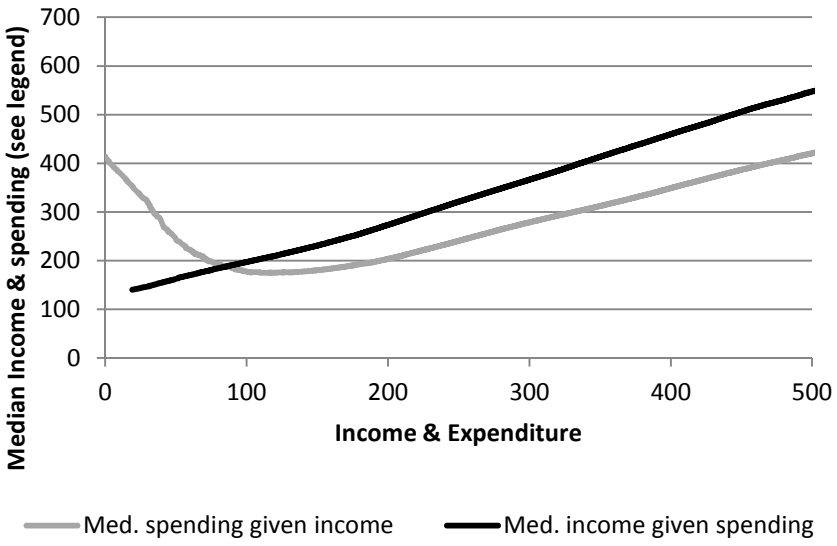

Notes: Data is from 2006 to 2009. Values on the vertical axis are given in equivilized real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale. The lines in the graphs represent smoothed conditional medians and are generated by performing locally-weighted median regressions. Further details are given in Section 3.

We find that under-recording of income likely plays the major role. We provide several pieces of evidence. First we show that a similar tick exists when we look across the income distribution and condition on work status (employment and unemployment). Given that most transitions into poverty are between rather than within work status (see, for example, Altonji et al. (2009)), 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 under-recorded. In particular, we find that, as Meyer and Sullivan (2003) also show in the US, UK data sources 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.

Furthermore, 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 most favours the consumption smoothing hypothesis, then to show that it cannot reproduce the tick. The main result of the simulations is that, even with the most 'generous' model (in particular, without borrowing constraints) only around 60% of the tick can be explained. If we choose less 'generous' specifications, then the fraction comes down. In particular, when borrowing is constrained (a more realistic setting, perhaps, than the model without borrowing constraints), then only a small fraction of the tick can be explained.

To complete the argument, we note two reasons why over-reporting of expenditure is unlikely to be playing a large role here. First, we compare the estimates of total household expenditure implied by the microdata with those reported in the UK National Accounts. We argue that, as in many other developed countries expenditure is, if anything, under- rather than over-reported across the distribution. Meyer and Sullivan (2011a) show this is the case in the US and Barrett et al. (2012) compare the phenomenon across the US, the UK, Canada and Australia. Second, we show that at the bottom of the reported income distribution expenditure is more closely correlated with other measures of living standards than is income. Absenting consumption smoothing, this implies that expenditure is better measured than income at the bottom end of the distribution.

Our final contribution, therefore, is to describe briefly how income and expenditure correlate with other measures of living standards, such as owning a car or the size of dwelling.[4] Leaving aside whether households report extremely low incomes because of misreporting or consumption smoothing we show that consumption is much more closely correlated than income with a variety of measures of living standards.

As mentioned, this paper follows several papers by Meyer and Sullivan, which discuss the relative merits of consumption versus income for measuring poverty, and apply this to important groups, such as single mothers.[5] However, our paper presents a comprehensive assessment across all groups in society (and not just low-education lone parents, or other groups thought to have a low income), and across four decades of micro-data. We also decompose the bottom of the distribution, and use a theoretical model to complement our more descriptive evidence. Our paper also builds on, Saunders et al. (2002), Attanasio et al. (2006), Brewer et al. (2006) and Brewer et al. (2009) who show, using the similar data for the UK, that those with the lowest income do not have the lowest expenditures.[6]

[4]Note here, we use the term consumption rather than expenditure. The distinction that we make between the two is explained in Section 6. However the key difference is our adding to our measure of expenditure an estimate of the net consumption flow coming from home ownership.

[5]Meyer and Sullivan (2003), Meyer and Sullivan (2004), Meyer and Sullivan (2008), Meyer and Sullivan (2011a) and Meyer and Sullivan (2011b).

[6]See also an earlier report by Department of Social Security (1993).

This paper also relates closely to Sabelhaus and Groen (2000). They observe that the bottom decile of the income distribution in the US Consumer Expenditure Survey (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. They conclude that income must be mismeasured in the CEX. We take particular care about modelling income dynamics explicitly and formulating and solving an explicit consumption and saving model. We also take their analysis further by showing that it is for those with the very lowest incomes that the apparent dis-juncture between consumption and income is at its most stark.

The paper proceeds as follows: after briefly describing the data in Section 2 we examine the tick in detail in Section 3. Section 4 examines the direct evidence that income is under-recorded at the bottom end of the distribution, while expenditure is likely to be better recorded. In Section 5 we examine whether consumption smoothing can explain the tick using the dynamic model. Section 6 shows that having low consumption is better correlated with other measures of deprivation than having low income. Section 7 concludes.

2 Data

Our main dataset is the Living Costs and Food Survey (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 a nationally-respresentative survey with an annual sample of between 5,000 and 7,000 households. Its primary purpose is to provide data on household spending patterns to inform the derivation of 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.[7] The LCFS is no longer the best survey for current income,[8] but it is the only UK survey which contains comprehensive measures of both income and expenditure.

Our aim in this paper is to assess whether the differences between income and expenditure over a particular period are indicative of a systematic type of measurement error in either of those series or can plausibly be explained by saving behaviour. It is therefore important that our measures of both income and expenditure are as broad and consistent as possible. Our measure of income (which we call

[7]The number of months varies between items on the questionnaire. 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 3 months.

[8]The Family Resources Survey has, since it was started in 1994, been more often used for poverty measurement as it has a larger sample size.

'cash 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 (for example the value of free TV licenses that the elderly 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.

Our measure of expenditure records all cash outlays by a household in a given period. This information is based on the two-week diary for items that tend to be purchased frequently (e.g. food, drink, household consumables, petrol) and the questionnaire for household expenses that are incurred less frequently (e.g. household furnishing and appliances, household bills, mortgage and rental payments). Our expenditure series can be interpreted therefore as average (weekly) expenditure in recent months.

Both income and spending are therefore measured over shorter periods (though periods that are similar to each other), than in the US Consumer Expenditure Survey which measures income over the previous year and spending over the previous quarter. However 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.[9]

Our measures of both income and spending are at the level of the household and are adjusted for household composition using the modified OECD scale.[10] We express income and expenditure quantities in terms of real (December 2009) UK pounds per week. We use the Retail Prices Index (RPI) to convert nominal quantities to real quantities.

[9]See Brzozowski and Crossley (2011). The Canadian surveys also make use of a balance edit, where participants are probed if they report annual spending and income figures that are too dissimilar.

[10]This is usually expressed as giving a weight of 1 to the household head, of 0.5 to each additional adult member and of 0.3 to each child. We follow usual UK practice and re-base so that a two-adult 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.

The LCFS contains some imputed data where satisfactory responses were not obtained from respondents. The analysis in this paper retains those households whose responses contain imputed data. However we have redone the analysis in this paper dropping households which we know contain imputed data. None of the results presented are sensitive to the exclusion of these households.

3 Examining the 'Tick'

Figure 1 in the introduction plots, for all households in our sample in the four most recent years of our data (2006/07-2009), the median household spending[11] given household income[12], and the median household income given household spending.[13] The clear pattern is that households with very low reported cash income (below £50 a week, say) have reported cash outlays associated with households with a weekly income of around £400 (above median income) – that is, the relationship between reported income and median reported expenditure has a 'tick' (or for North Americans a 'check') 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 (2011b) 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.

We first explore what might be causing this tick by examining expenditure patterns for various subgroups. Each figure shows the profile of median reported spending conditional on income reported for a number of household types. These are generated using locally-weighted median regressions.

Figure 2 shows that the tick pattern appears for all working-age family types. The only household type for which it does not appear is pensioners.[14] Figure 3 shows that the tick is also apparent across education categories. Together these figures show that the tick is a robust feature of the data for working-age households and cannot be explained by demographic factors.

To complement the smoothed profiles that these figures show, we have checked, and confirmed, that the 'tick', where shown, is statistically significant. That is, in results not shown, we run a median regression of reported spending on a set of indicator variables for £50 income bands. The reported expenditure

[11] This is 'cash outlays', as defined in Section 2.

[12] This is 'cash income', as defined in Section 2

[13] Similar analysis exists, for countries other than the UK or covering a subset of the years used in this study in Meyer and Sullivan (2003); Brewer et al. (2006); Brewer et al. (2009); Brzozowski and Crossley (2011).

[14] See Brewer and O'Dea (2012) for further details on patterns for pensioners.

of those in the lowest income group (£0-£50) is statistically significantly lower than each of the next seven income groups (that is groups representing households with equivalised reported income of less than £400). The 'tick' is significant in this sense for each work-status group (employed, self-employed, workless), for each education group, and for each non-pensioner household type.

Figure 2: Median expenditure by income by family type (working-age households)

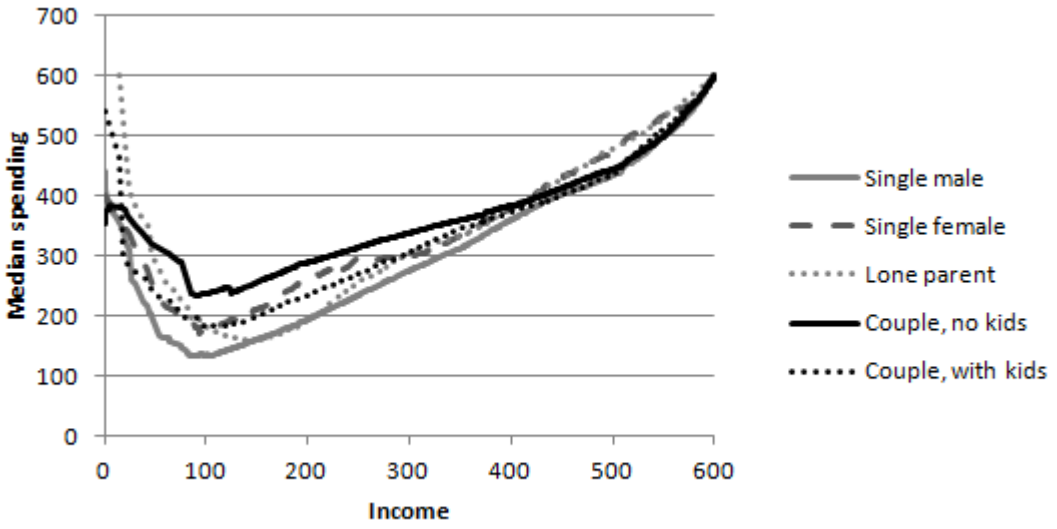

Notes: Data is from 2006 to 2009. Values on the vertical axis are given in equivalised real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale.

Figure 3: Median expenditure by education group

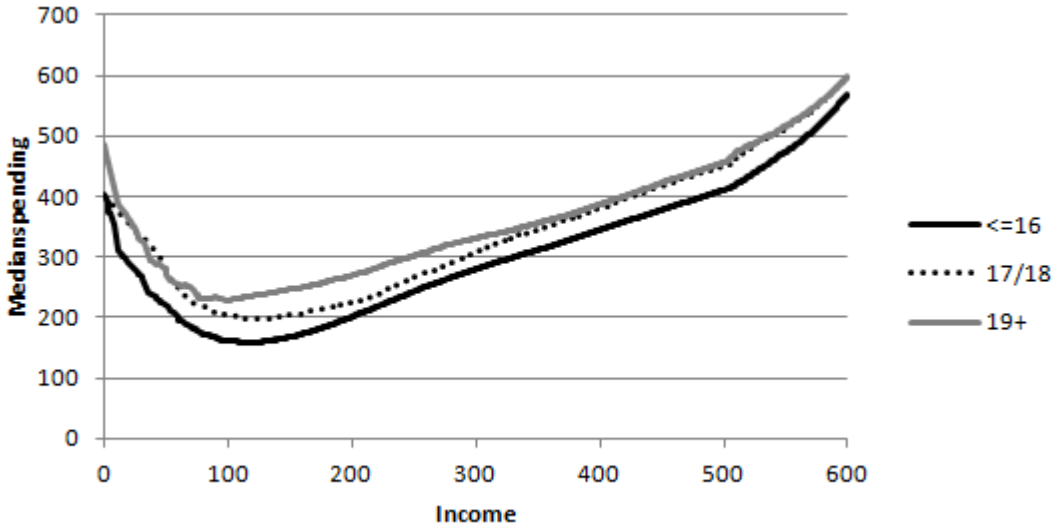

Notes: Data is from 2006 to 2009. Values on the vertical axis are given in equivalised real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale.

Figure 4 provides our first piece of evidence that the tick is not explained solely by consumption smoothing alone. It shows median expenditure across the income distribution when we split by work status. It shows that the tick is evident for all households, whether the head is working, is out of work or is self-employed. Given that most large temporary fluctuations are between employment states[15], we would expect that if consumption smoothing explains all 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 a role for measurement error.

Figure 4: Median expenditure by employment status

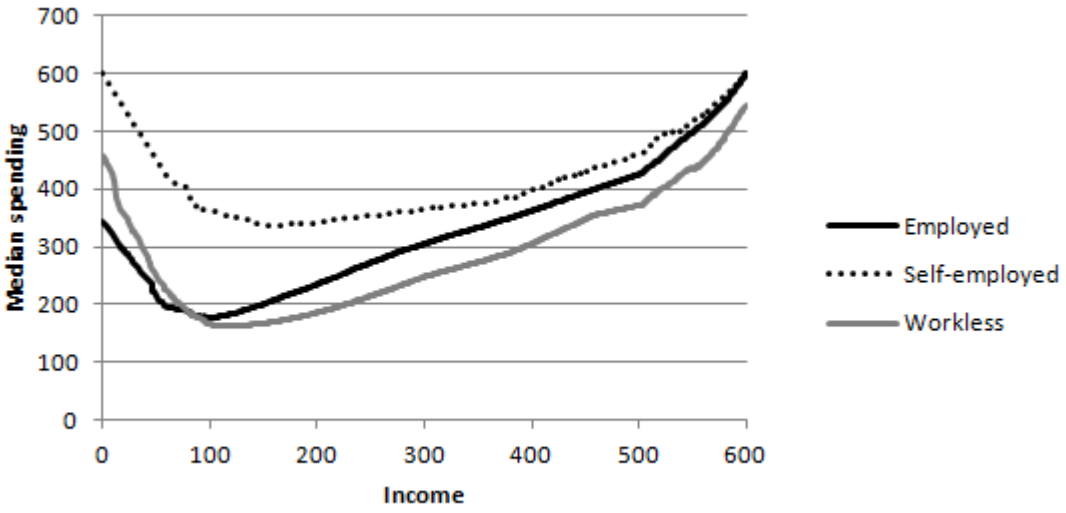

Notes: Data is from 2006 to 2009. Values on the vertical axis are given in equivalised real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale.

Figure 5 shows the 25th and 75th percentiles of expenditure alongside median expenditure conditional on level of reported income. The tick that we previously described is also present at each of these percentiles, although it is less obvious at the lower quantile. As the graph implies, the variance of spending increases at low levels of income. This figure suggests that the bottom of the measured income distribution contains more 'types' than income levels even slightly higher up. This further suggests that the bottom tail includes both those with persistently low income as well as those who are better off (but either under-report income or are smoothing consumption).

[15]See, for example Altonji et al. (2009) who decompose earnings shocks into employment transitions and wage shocks.

Figure 5: Expenditure at selected percentiles by income

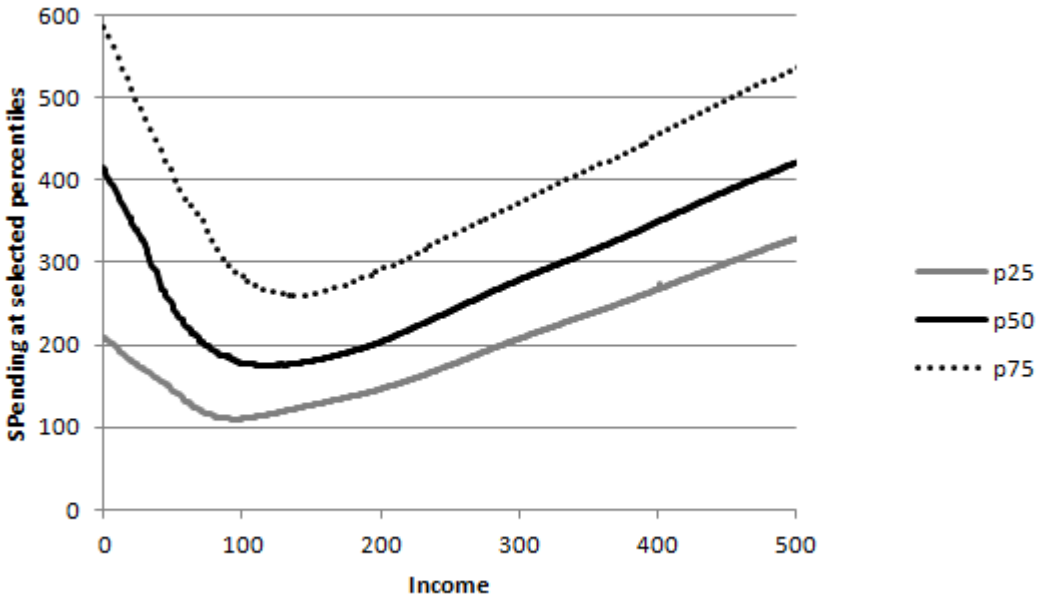

Notes: Data is from 2006 to 2009. Values on the vertical axis are given in equivilised real pounds per week, expressed in December 2009 prices. Equivalisation is carried out using the modified OECD equivalence scale.

Figure 6 shows the median (real) expenditure conditional on (real) reported income for five-year periods starting in each of 1978, 1983, 1987, 1993, 1998 and 2003. In each of the periods the 'tick' is evident, though its magnitude has been growing (i.e. median expenditure at the bottom of the reported income distribution has been growing). This graph, at least is consistent with both measurement error and consumption smoothing. On the one hand measurement error may have got worse, because survey participation has declined over time. On the other hand, access to credit markets has become more widespread (at least until the financial crisis of 2008) and households are better able to smooth consumption.

Figure 6: Median expenditure by income (equivalised real £/wk), 5 year averages

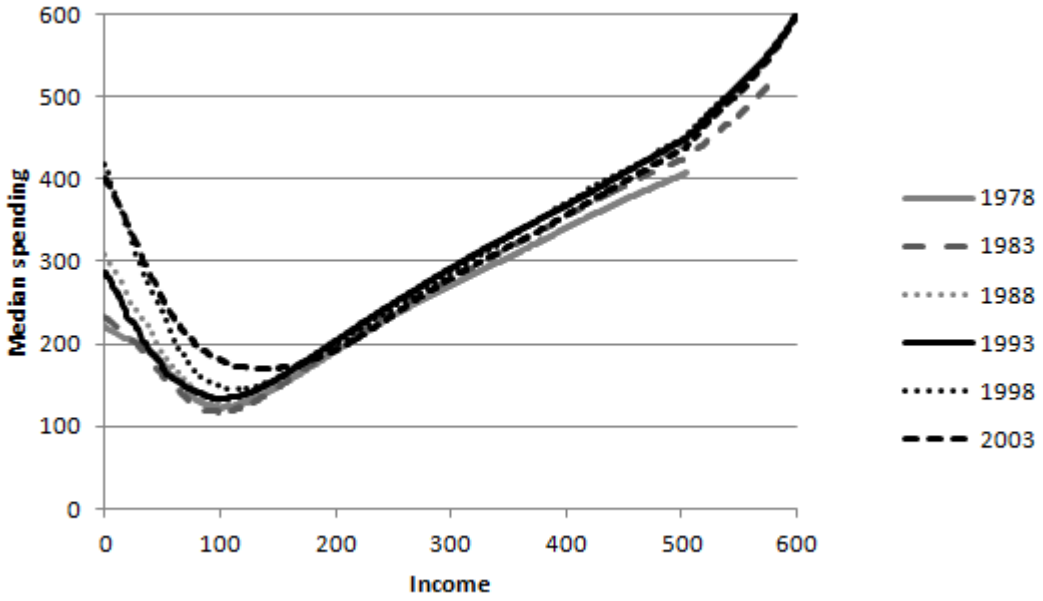

Figure 7 looks a little more closely at this pattern – it shows median expenditure conditional on income as a proportion of overall median expenditure, and does this for those in centiles 1 through 8 of the income distribution. It shows that the tick, at the median, is most clearly evident among the bottom 1%. The tick therefore represents a small fraction of the population. However, as discussed above, the phenomenon we are investigating - expenditure far in excess of reported income - is also evident further up the distribution. We concentrate on the bottom 1% because it is a particularly salient feature of the data. Nevertheless, we emphasize that many of our arguments, particularly in Sections 4 and 6, apply also to households further up the distribution.

Figure 7: Median normalised expenditure by centiles of equivalised real income

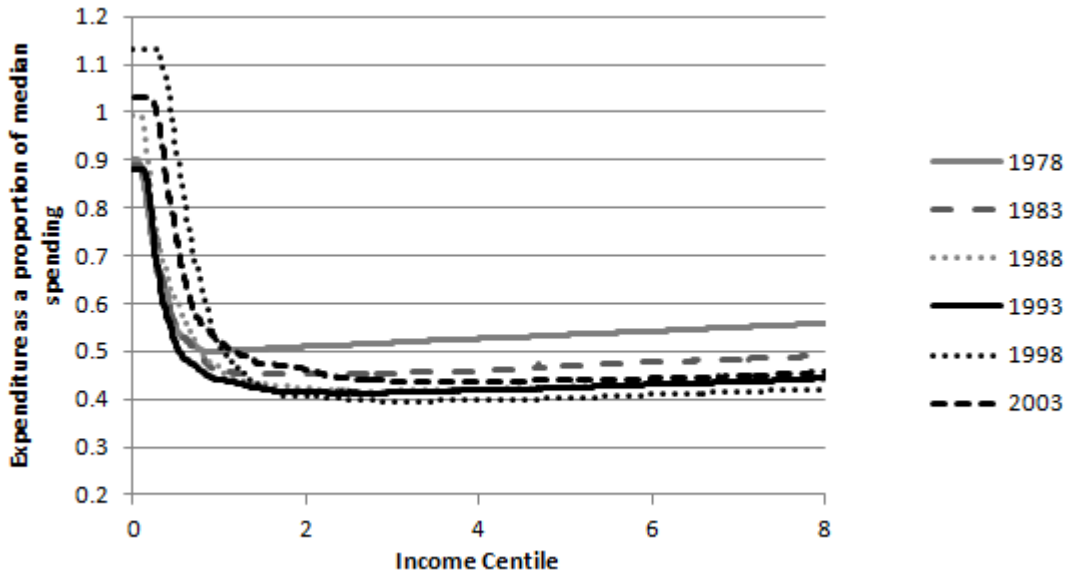

Notes: Data in each series is for the five years starting in the year shown in the legend.

4 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-recorded income, over-recorded spending, or consumption-smoothing (in this case, running down assets or increasing debt, or "dissaving" for short). 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.

4.1 What evidence is there that income is mis-measured?

Meyer and Sullivan (2003) show that income from welfare benefits is under-reported in the CEX and that reported values of earnings and hours in that survey are inconsistent with minimum wage rules. They use these observations to argue that under-reporting of income from cash transfers and earnings are likely explanations for the observed mismatch between reported cash income and reported cash outlays. Table 1 shows the fraction of the amount that the government says that it pays out in various state benefits that is recorded in the LCFS (having used population grossing weights to give aggregate estimates). Coverage rates are high for the two benefits which are universal or close to it: child benefit

and the basic state pension.[16]. 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 (58%) 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.

Table 1: Coverage in LCFS of spend on cash benefit programmes

	Coverage	Spend (£m/yr)
Retirement pension	0.95	66,480
"Other" non contributory benefits (see notes)	0.52	27,970
Working and child tax credits	0.50	21,270
Rent rebates and allowances	0.83	18,930
Income support & pension credit	0.68	16,580
Child benefit	0.96	11,880
Incapacity benefit	0.74	6,670
Maternity/Statutory maternity pay	1.19	1,900
Jobseekers allowance	0.80	1,200
War pensions	0.33	1,020
Student support	2.36	970

Source: Authors' calculation using Tables 13 and 14 of Barnard (2011)."Other" non-contributory benefits mostly comprises benefits paid to disabled people or those requiring care in their own home.

Figure 8 plots trends in the amount of "missing" income from state benefits as a fraction of total household income, 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.[17]

[16]It should be noted that there are legitimate reasons why the fraction captured should be below 100% some benefits are paid to people outside the UK, and some are paid to people in the UK who do not live in private households and who would therefore be outside the sampling frame of the LCFS

[17]Manipulation of the same data (Tables 13 & 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 per cent, over the same period.

Figure 8: "Missing" income from cash benefits as a percentage of total household income

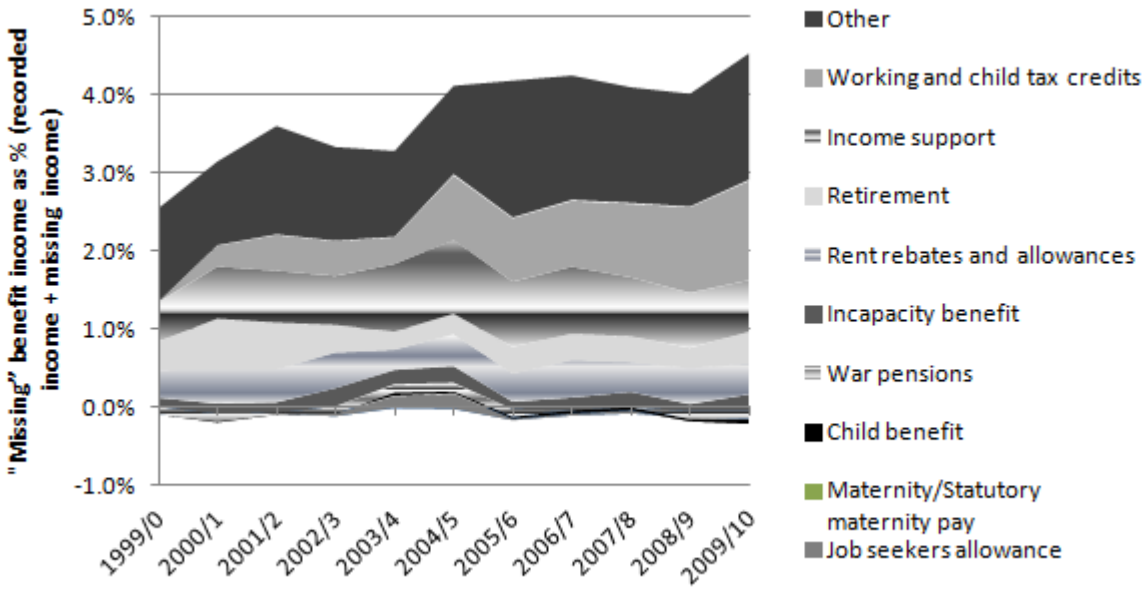

Notes: Authors' calculation using Tables 13 and 14 of Barnard et al. (2011) and previous versions.

It is not clear whether this low (and declining) microdata coverage of benefit income is due to differential patterns of non-response to the LCFS that is going uncorrected when grossing weights are calculated, or whether it is due to incorrect item response amongst households who are genuinely receiving benefits, or whether it is due such households reporting that they receive benefits, but under-reporting the amounts. However, we do know that the phenomenon of under-recording the total spend on cash benefits in the UK is not unique to the LCFS: Brewer et al. (2008) report that, in 2006-7, around one third of government spending on child and working tax credits was not captured in the main UK household survey used for recording income, the Family Resources Survey (FRS), and 43% of spending on the pension credit – the main means-tested programme for pensioners – was missing (see Bound et al. (2001) for a general discussion of this phenomenon and see Lynn et al. (2012) for what little is known from UK survey data that has been linked to administrative data).

Turning to the question of whether individuals are reporting wages that are below legal minimums, there has been a nationally-set, legally-binding, minimum wage in the UK since 1999. There are very few exemptions (although employers can count the value of some employer-provided benefits). 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 5% and 11%) are substantial, although we cannot tell whether earnings have been under-reported or hours worked over-reported.

Table 2: Proportion of employees aged 23 and over reporting implied hourly earnings below the national minimum wage

Year	Proportion
2001	5.9
2002	6.8
2003	6.7
2004	8.0
2005	7.0
2006	8.5
2007	8.6
2008	8.8
2009	10.5

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

4.2 What evidence is there that spending is mis-measured?

The phenomenon we investigate in this paper - the divergence between income and expenditure for those with the lowest reported income could in principle, be driven by over-reporting of expenditure. Previous studies, however, have noted that total expenditure (and indeed total income) captured by the UK's household budget survey is less than expenditure as measured in the UK National Accounts (see especially Deaton (2005); Attanasio et al. (2006)).[18] Figure 9 illustrates this phenomenon by showing the ratio of aggregate income and expenditure implied by the LCFS to a similar measure derived using the UK National Accounts. The figure shows that the LCFS is under-recording spending relative to the National Accounts, doing so more substantially than is the case with income, and doing so at an increasing rate. While these are, of course, economic aggregates and some care should be taken in using them to inform statements about particular groups, this low (and declining coverage) is certainly indicative of the fact that, as in the US (see Meyer and Sullivan (2011a)) expenditure is, if anything, under- rather than over-reported in aggregate.

[18]Similarly Crossley and O'Dea (2010) look at longitudinal data on changes in household wealth between 2000 and 2005 using data from the British Household Panel Survey and note that the very high saving rates implied by the LCFS are inconsistent with (i.e. too high given) the observed wealth accumulation over the same period.

Figure 9: Income and Expenditure Coverage of UK National Accounts Data

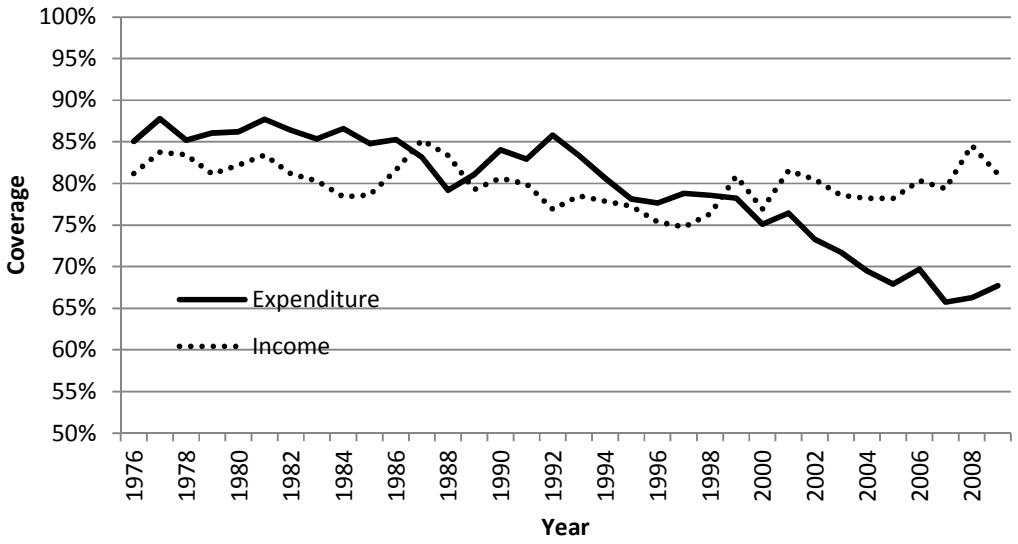

Source: Authors' calculation from Living Costs and Food Survey (and predecessors) and National Accounts 1976-2009.

Additionally, Brzozowski and Crossley (2011) argue that over-reporting of spending is unlikely to be a contender for Canadian data. Brewer et al. (2009) show (for households with children) that the tick-shaped pattern between income and other measures of living standards exists when using many measures of (or proxies for) living standards (and when looking at income measured in four different British household surveys); this strongly suggests that over-recording of spending is not an important cause of the pattern depicted in Figure 1.

4.3 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 dis-saving (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, and so we cannot investigate in detail the extent to which households whose reported spending far exceeds their reported income are dis-saving or borrowing.[19] 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

[19]This 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.

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. 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 10 plots the relationship between net assets and reported household income in 2005 and shows that most of those with the lowest report income have low levels of net assets.

Figure 10: Distribution of financial assets by £25 income bands amongst households with children (BHPS), 2004-05 to 2006-07

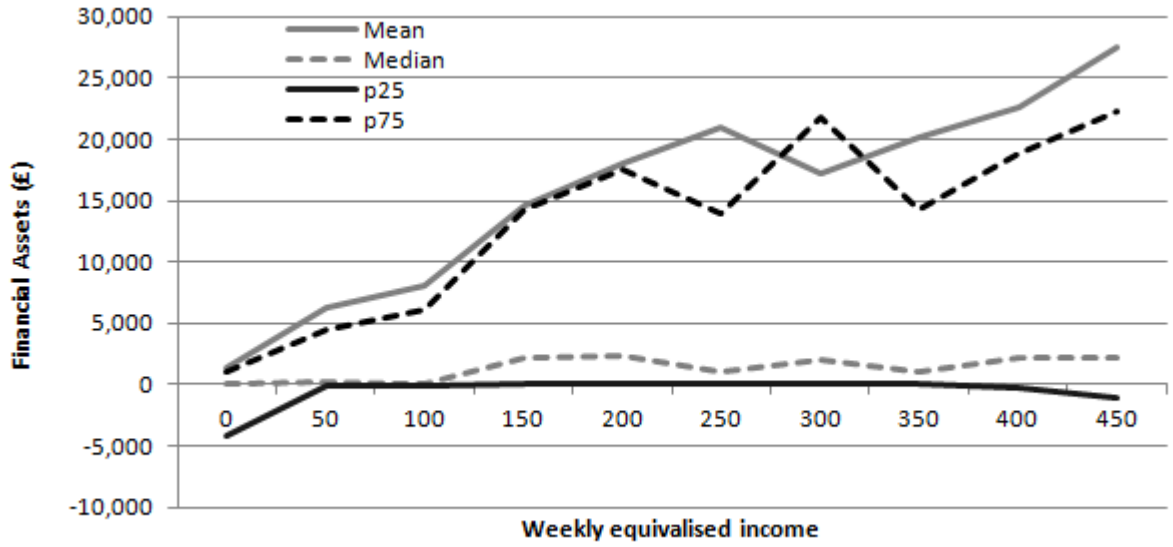

Notes: Authors' calculations using British Household Panel Survey 2005. Financial assets comprise all liquid financial wealth (cash savings, stocks and shares) and are net of non-mortgage debt.

We infer from this is that the majority of households whose low reported income is less than their cash outlays will not be running down savings (as they don't have any), but 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 shows 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 medium term (which is our interpretation of this three year period), the level of access to credit markets would have to be implausible.

5 Can Consumption Smoothing Explain the Tick?

We now examine the tick through the lens of a theoretical model. To show that consumption smoothing cannot explain all the tick we formulate a model of income dynamics and consumption behaviour that both fits panel income data prima facie and is most likely to induce high consumption at the bottom end of the distribution. We then show that such a model cannot reproduce the tick.

There is no clear consensus in the literature on the correct specification of the income process. Our approach therefore is to formulate a model of income dynamics that ascribes all ultra-low income realisations to be transitory. This specification is much more likely than one in which such shocks are persistent to generate a tick similar to that observed in the data without appealing to measurement error as an explanation.

The ability to smooth consumption critically depends on the tightness of borrowing constraints, which are difficult to identify. We therefore specify one model with the loosest possible borrowing constraints and another with the tightest. Even with the loosest constraints, consumption smoothing cannot explain all the tick.

As discussed in Section 3, it is likely that unusually high consumption (due to either consumption smoothing or income mismeasurement) is not confined to the bottom 1% of incomes. Figure 2, for example, shows that implied saving is negative much further up the distribution. Nevertheless the phenomenon is particularly stark for the bottom 1%. Therefore in this section we focus on this bottom 1% as being 'unusual' and use the 1st-10th percentiles as a comparison/control group. Also we use as a basis for discussion that expenditure for the bottom 1% in the data is equal to the median across the population. This is a good approximation even though Figure 7 shows that it is just below median expenditure for many years.

5.1 Empirical Evidence on Transitions into and out of the Bottom 1%

Our first task is to examine the income transitions into the bottom 1% empirically using panel data from the British Household Panel Survey (BHPS). Table 3 shows these transitions from other quantile ranges. The rates here use equivalized current net income.[20] The data are split into yearly quantile bands conditional on cohort,[21] year and the household appearing in the sample over the two periods.

Table 3 shows that households enter and exit the bottom 1% from across the entire income distribution. A noticeable feature of the data, however, is that the entry and exit rates are not uniform across the

[20] This uses the variable 'hhnetde' from the derived income variables, documented by Bardasi et al. (1999))

[21] We use decade of birth.

top 99%. On first sight, the data are consistent with a simple permanent-transitory model. Households from the top 10% of the distribution, for example, may enter the bottom 1% temporarily, but more households enter from the bottom two bands because downwards permanent shocks from these bands need not be so large to push the household to the bottom. However, within this permanent-transitory model, the bottom 1% will be filled mainly with those of low permanent income, whose consumption will be low. We therefore seek an alternative model that fits these data but ascribes ultra-low income realizations to be more temporary.

Most of the existing literature on income dynamics models the entire income distribution. Our interest centres on transitions into and out of poverty, about which there is much less research. Jenkins (2011) stands out as a recent work on this specific topic and emphasizes the persistent nature of poverty. He discusses Cappellari and Jenkins (2004), who use a Markov model to analyse poverty dynamics. We adopt a approach similar here. Cappellari and Jenkins, however, examine movements into and out of the official poverty threshold (defined as 60% of median income), which captures around 10% of households. We are instead focusing on the poorest 1% of households.

5.2 An Income and Consumption Model

In the light of the above discussion, we specify the following dynamic process for income y_{it}:

$$\begin{aligned} y_{it}^P &= \rho y_{it-1}^P + \zeta_{it} && (1) \\ y_{it} &= \theta_t y_{it}^P + (1-\theta_t)\,\underline{y} \\ \theta_t &\in \{0,1\} \\ Pr\left(\theta_t = j | \theta_t = i\right) &= \pi_{ij,y^P} && (2) \\ \zeta_{it} &\sim \text{ln-}N\left(0, \sigma_\zeta^2\right) \end{aligned}$$

where y_i^P is 'permanent' income for household i. It evolves as a persistent AR(1) process. The parameter of persistence is $\rho \in [0,1)$. Although the focus in this paper is income fluctuations into and out of poverty, we include income risk across the distribution because this causes a higher (and more empirically valid) accumulation of wealth. Wealth holdings have a large effect on consumption choices when the household enters poverty. Aggregate wealth holdings in this model are precautionary in nature, because there is no life cycle: the horizon is infinite.[22]

Complementing the risk to 'permanent' income, θ_t is a markov chain governing transitions into and out of poverty. We assume that the process is stationary, i.e. the probabilities are constant over the sample

[22] To ensure a long-run stationary distribution ρ must be strictly less than 1.

period. The pure transitory process would be captured by $\pi_{0,0} = \pi_{1,0}$ and $\pi_{0,1} = \pi_{1,1}$ (i.e. the rows of the associated transition matrix would be identical). However, as mentioned above, there are two clear ways in which the pure transitory process contradicts the empirical transition probabilities in Table 3. First, there is substantial persistence in the bottom 1%. Second the probability of entering the bottom 1% is not uniform across the upper 99 centiles.

Table 3: Transitions Into and Out of Bottom 1%

quantile	% transiting into bottom 1%		% transiting out of 1%*	
90-100	0.48	(0.08)	3.7	(0.2)
80-90	0.24	(0.06)	2.7	(0.2)
70-80	0.32	(0.07)	3.9	(0.2)
60-70	0.41	(0.08)	4.8	(0.3)
50-60	0.39	(0.07)	5.3	(0.3)
40-50	0.45	(0.08)	5.9	(0.3)
30-40	0.79	(0.11)	6.3	(0.3)
20-30	0.69	(0.10)	6.6	(0.3)
10-20	1.13	(0.13)	9.7	(0.4)
1-10	3.56	(0.23)	27.5	(0.6)
0-1	23.80	(1.56)	23.7	(1.6)

Notes: Source: BHPS.

Asymptotic standard errors in brackets. Standard errors given by $\sqrt{\frac{p(1-p)}{n}}$.
Data created using hhnetde.
* These numbers sum to 100%.

To match the transitions in Table 3 we therefore first impose different probabilities of entering the bottom 1% for 4 quantile groups. These probabilities are given in Table 4 (first column). To complete the description of poverty dynamics we must then specify the probability that a given type remains in the bottom 1% after a poverty shock. For simplicity, we assume that, across the productivity types, this probability is constant. In order to match the flows out of the bottom 1% with the flows into the bottom 1% this probability must be 24%. This probability matches the bottom right hand entry in Table 3. We can check how good is the assumption of a uniform probability of exit by comparing the exit rates from the model and in the data. These exit rates are summarized in the 2nd and 3rd columns of Table 4. Notice that the two columns line up well, implying a good fit of the model. As a final point about the poverty dynamics, note that each household's probability of entering and exiting poverty is changing constantly as the level of permanent income changes.

Table 4: Income Transitions By Quantile

Quantile range	$\Pr(\theta_t{=}0 \mid \theta_{t-1}{=}1)$	Simulated exits from bottom 1%	Exits from bottom 1% from data*
50-100	0.0037	23%	27%
20-50	0.0064	24%	25%
10-20	0.0113	14%	13%
1-10	0.0356	39%	36%

Notes: Source: Simulations and BHPS (final column).
* This is the row sum from the final column of Table 3, divided by the total fraction leaving the bottom 1%

We specify that consumption evolves in a manner implicit in the following functional equation:

$$V\left(A_t, y_t^P, \theta_t\right) = \max_{c_t} \frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t \left(V\left(\left(A_t - c_t\right) R + y_{t+1}, y_{t+1}^P, \theta_{t+1}\right)\right)$$

where y_t^P and θ_t evolve as specified in equations 1 and 2, and where we have dropped i subscripts for notational convenience. Here we use the standard isoelastic CRRA utility function. This functional equation has a solution (i.e. there exists a fixed point of the value function) as long as $\beta < 1$. The utility specification implies that households accumulate precautionary saving to protect against negative income shocks, such as a poverty shock. It also implies that households have a finite target wealth level as long as $\beta R < 1$.[23] .

We investigate the model with two extremes of borrowing constraints. First we allow natural borrowing constraints. This implies that households can borrow up to the point at which debt would escalate to infinity if the household only receives minimum possible income in perpetuity. At the specified interest rate the long-term debt limit is 14.9 units of income (where 1 is the average income across society). Second, at the other extreme, we impose strict non-zero borrowing constraints. We recognize that each regime is unrealistic but captures the bounds of likely credit limits. An alternative approach would be to match results in, for example, Browning and Crossley (2001) on the average consumption drop on entering unemployment. We could then calibrate the borrowing constraint to match those estimates. However, that literature focuses on the effect of unemployment benefits which we do not model closely here, and therefore that approach seems out of spirit with the current analysis. We therefore prefer to bound the effects using the two extreme models.

Further details of the model solution, parametrization, calibration and simulation are in Appendix A.1.

[23] I.e. assets will not tend to infinity: whenever assets get too high, households will want to de-cumulate.

5.3 Results

Figure 11 shows simulated median consumption by income centile. It does so for both the models with and without borrowing constraints. The income distribution is the same under both specifications. Under both regimes, there is a tick in consumption, because the bottom 1% of the income distribution contains a broader mix of skill types than the bottom 1-10%. However, if we take as a benchmark from the data that the bottom 1% have consumption equal to median consumption, then even the economy without credit constraints cannot reproduce all the tick. As a rough calculation, to reproduce the tick in full, the bottom 1% would need to consume over twice as much as the the 2nd percentile (which here consumes just under 45% of median consumption). Actually, it consumes only around 75% more. Therefore it seems, this model explains at most 60% of the tick. Recall that this is the most 'generous' model we can devise.

The economy with tight borrowing constraints explains even less of the tick. If we repeat the same calculation, then in order for the bottom 1% to have median consumption they would need to consume 3 times as much as the bottom 1-10%. Actually they only consume around 30% more.[24]

Figure 11: Simulated Consumption Across the Income Distribution

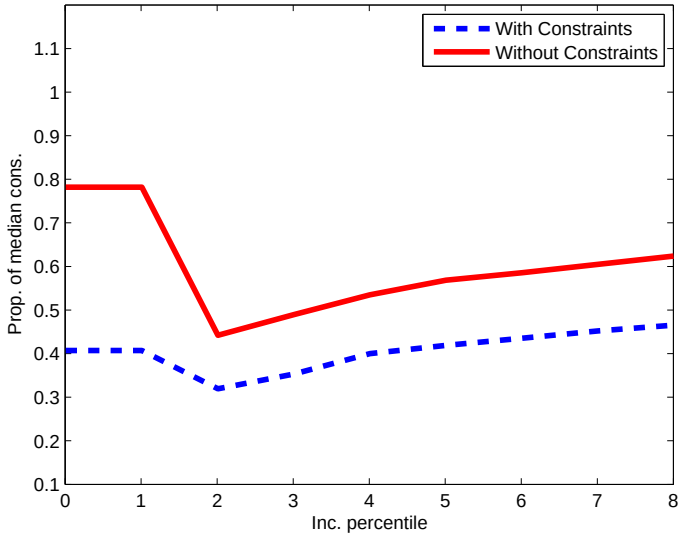

Consumption smoothing cannot account for all the tick in either model because of two factors. First, the bottom 1% does not contain, on average, households of median productivity level; second, households cannot perfectly smooth their consumption through a period of low income. We now investigate, in each model, to what extent these factors apply.

[24]This result is robust to changing the variance of permanent incomes.

Table 5 shows how well households can smooth consumption: it shows the ratio of consumption to permanent income across the 4 quantile groups. The columns give consumption across each state of the transitory shock.[25]

Table 5: Simulated Consumption By Income Quantile: With Credit Constraints

	No Constraints		Constraints	
Quantile range	$\theta = 1$	$\theta = 0$	$\theta = 1$	$\theta = 0$
50-100	93.9%	89.7%	99.9%	83.6%
25-50	120.8%	116.2%	115.7%	83.4%
10-25	134.3%	129.1%	117.4%	79.8%
1-10	147.1%	142.9%	115.5%	86.2%

Notes: Table reports consumption as a % of skill level y_i^p

We first look at the effect of a transitory shock across the quantile groups in the economy without constraints. The consumption drop upon a negative transitory shock is only around 4% across the population and slightly less for the low skilled. Therefore, in this version of the model the size of the consumption drop has little effect on the tick. And therefore, in this case, consumption is less than median for the bottom 1% almost entirely because the bottom 1% is predominantly made up of low skilled rather than high skilled.

At the other extreme, the right hand two columns of Table 5 show average consumption in the model with constraints. As discussed above, notice that the consumption drop upon entering poverty is far greater than without the constraints for all productivity levels. Interestingly, this is not only directly because of binding constraints. Binding constraints affect only around half of the bottom centile.[26] However, the tight credit constraints affect everyone. Even those with positive assets (for example, the high skilled) must protect their wealth buffer if they enter poverty. Therefore they smooth consumption very badly. To conclude, for this model, the tick is muted both because the bottom 1% disproportionately includes the low skilled and because consumption is very imperfectly smoothed.

[25] Consumption is above 'permanent' income for most of these groups for two reasons. Some of these groups are actually dissaving because they expect income to revert upwards in future. But more importantly, these tables don't include asset income, which is positive on average across the income groups. In fact those at the top of the productivity distribution have large positive assets so have high asset income, while those at the bottom of the distribution have negative assets so are repaying debt. Therefore the ratio of consumption out of total income is even more skewed than out of 'labour' income.

[26] For the 1-10% of the income distribution around 70% of households are at the constraint. They would like to consume more than income because they expect income growth in the future, but they are prevented by the constraint.

Figure 12: Simulated Wealth/Income Ratio Across the Income Distribution

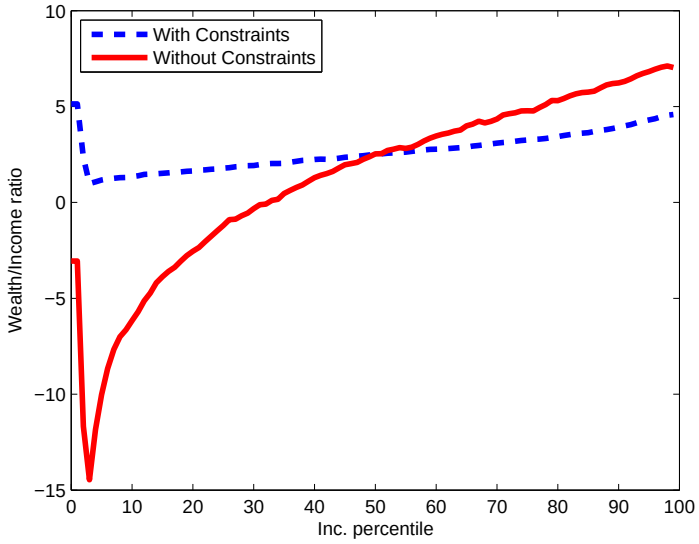

Of course, neither model of borrowing constraints matches reality perfectly. However, it is illuminating to show how the distribution of wealth in these simulated economies matches that in the data. Figure 12 shows the simulated wealth distribution in both economies. Note that wealth dispersion is far greater for the unconstrained economy than the constrained economy and the distribution in the latter economy is much closer to that seen in the data in Figure 10.

6 Comparing consumption and income as correlates of low living standards

Section 3 showed the mismatch between expenditure and income that occurs amongst households with a low recorded income, and Sections 4 and 5 give evidence for our contention that under-reporting of income is at the heart of this problem. This adds to the theoretical and empirical arguments, put forward by many authors, particularly Meyer & Sullivan, for using data on expenditure either instead of, or in addition to, data on income in the measurement of poverty. In spite this, income-based measures of poverty retain a near-monopoly position in political debate and in discussions of social policy in the UK.

In this section we ask whether measures of consumption (as estimated using expenditure data recorded by a household survey) are better correlated with other measures of low living standards than income (as recorded by a household survey). The evidence we present should be interpreted as arguing for some greater use of expenditure data in the measurement of living standards - *regardless of the provenance of the divergence between income and expenditure at the bottom distribution of reported income distribution.*

The results we present in this section use a different measure of income and consumption to those used in the rest of the paper (though the results we present are almost identical when we use the measures of income and expenditure described in Section 2). To the expenditure (cash outlays) measure that we introduced in Section 2, we add the imputed rental value of owning a property or vehicle (the manner in which these imputations are carried out is described in detail in Section 3 of Brewer and O'Dea (2012)). We subtract cash outlays on housing or vehicles and make a number of other small adjustments (such as adding the value of TV licences and school meals for those households who do not have pay for them). For reasons described below, we also subtract all spending on durables. To distinguish this quantity from the measure of expenditure we used previously, and to reflect the fact that we are including the value of consumption flows from housing and vehicles, we describe it as 'consumption' in the discussion that follows.

In deriving a measure of income to compare to that derived measure of consumption, we are very conscious of the point made by Bavier (2008) that one should not compare consumption only to the measure of income used in the official analysis of poverty but to the "best" measure of income that can be derived. This is particularly important here - as the derivation of our consumption measure starts with expenditure data and makes adjustments in keeping with theoretical and empirical evidence about how best that data can be used to predict deprivation, then the odds would be stacked against income predicting living standards better than consumption unless a similar process is carried out to the income data. We therefore make exactly the same adjustments to the income measure that we make when moving from expenditure to consumption - most importantly, subtracting the income that is spent on housing and adding the imputed rental value of any owned property. We refer to this quantity as 'broad income' to distinguish it from that measure of income used in the rest of the paper.

We proceed using an approach suggested by Meyer and Sullivan (2003). We define four groups $BrInc_{low}$. $BrInc_{notlow}$, Con_{low} and Con_{notlow}, where the subscript 'low' refers to those households lying in the bottom 10 per cent of the consumption or income distribution, and 'notlow' to those households lying in the upper 90 per cent of the distribution in question, and BrInc and Con refer to broad income and consumption respectively. Defining X as some outcome that (at least arguably) correlates positively with living standards (for example having health insurance or owning one's own home), and defining $X(y)$ as the mean outcome for group y, we then calculate a difference-in-difference type measure:

$$[X(Con_{low}) - X(Con_{notlow})]–[X(BrInc_{low}) - X(BrInc_{notlow})] \tag{3}$$

This will be negative if being in the bottom decile group of reported consumption is a better indicator of poor outcomes than being in the bottom decile group of reported income.

We calculate this measure for ownership of various consumer durables (dishwasher, washing machine,

central heating, computer, DVD player, access to the internet at home, a TV, subscription TV), having health insurance, owning one or more cars, owning their own house, and the number of rooms in the house. We noted above that the measure of consumption used here does not include any spending on durables. This is to avoid the generation of mechanical (and not particularly informative) relationships between the measure of consumption and ownership of durables.

The results are shown in Table 6. All but one of the statistics have a negative sign and are statistically significant; the exception is owning a TV (perhaps unsurprisingly as the vast majority of all households now own a TV). Tables 7, 8 and 9 show the results of this analysis carried out on sub-groups defined (in turn) by family type, work status, and education group. A handful of estimates are positive for some measures for some family types, but the vast majority of the estimates are either negative and significant, or insignificantly different from zero. Although the LCFS provides a limited number of alternative measures of living standards, overall we conclude emphatically that having a low recorded consumption is a better guide to who has a low living standard than having a low reported income.

Table 6: The relationship between low consumption, low income and other outcomes, all households

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	N
	$X(Inc_{low})$	$X(Inc_{notlow})$	(1)-(2)	$X(Con_{low})$	$X(Con_{notlow})$	(4)-(5)	(6)-(3)	
Wsh. Mch.	0.92	0.96	-0.04	0.84	0.96	-0.12	-0.083***	52,796
Cent. Heat.	0.92	0.94	-0.03	0.89	0.95	-0.06	-0.030***	52,796
Dishwash.	0.20	0.36	-0.16	0.06	0.37	-0.31	-0.151***	52,796
DVD	0.59	0.61	-0.03	0.42	0.63	-0.22	-0.189***	52,796
TV	0.98	0.99	-0.01	0.98	0.99	-0.01	0.002	52,796
Pay TV	0.30	0.39	-0.09	0.18	0.40	-0.22	-0.134***	52,796
PC	0.54	0.64	-0.10	0.24	0.68	-0.44	-0.335***	52,796
Internet	0.41	0.56	-0.15	0.15	0.59	-0.44	-0.286***	52,796
Car	0.53	0.78	-0.25	0.24	0.82	-0.57	-0.321***	52,796
Two cars	0.14	0.33	-0.19	0.02	0.34	-0.32	-0.130***	52,796
Own hse.	0.37	0.74	-0.38	0.32	0.75	-0.43	-0.055***	52,796
No. rooms	4.97	5.38	-0.41	4.5	5.44	-0.94	-0.528***	52,796
Health ins.	0.05	0.13	-0.08	0.02	0.13	-0.12	-0.036***	52,796

Source: Authors' calculation using Expenditure and Food Survey/Living Costs and Food Survey 2001/02- 2009.
Notes: Column (7) gives the quantity expressed in equation (3). Negative numbers indicate that consumption is better correlated with the outcome in the left-hand column (e.g. having a washing machine, owning one's own home) than is income. Other columns give the individual components of the quantity in equation (3). *** indicates significant at the 1% level, ** indicates significant at the 5% level, * indicates significant at the 10% level. Confidence intervals are calculated by bootstrapping with 999 replications.

Table 7: The relationship between low consumption, low income and other outcomes, non-pensioner households

	Single male	Single female	Lone Parent	Couple, no kids	Couple, kids
Wsh. Mch.	-0.114***	-0.041**	-0.021***	-0.027*	-0.010
Cent. Heat.	-0.039**	-0.01	0.000	-0.025*	-0.036***
Dishwash.	-0.106***	-0.072***	-0.110***	-0.172***	-0.135***
DVD	-0.123***	-0.077**	-0.064***	-0.098***	-0.020
TV	-0.022*	0.019*	-0.010	0.010	0.000
Pay TV	-0.100***	-0.048**	-0.055***	-0.116***	-0.062**
PC	-0.262***	-0.268***	-0.174***	-0.307***	-0.156***
Internet	-0.253***	-0.223***	-0.171***	-0.271***	-0.159***
Car	-0.325***	-0.232***	-0.194***	-0.328***	-0.295***
Two cars	-0.060***	-0.025***	-0.021***	-0.254***	-0.190***
Own hse.	-0.226***	-0.141***	-0.134***	-0.224***	-0.200***
No. rooms	-0.601***	-0.144**	-0.355***	-0.467***	-0.385***
Health ins.	-0.044***	-0.027**	-0.011**	-0.081***	-0.055***
N	4,188	2,957	3,436	9,675	10,478

Source: Authors' calculation using Expenditure and Food Survey/Living Costs and Food Survey 2001/02- 2009.

Notes: This table gives the quantity expressed in equation (3) by household type. Negative numbers indicate that consumption is better correlated with the outcome in the left-hand column (e.g. having a washing machine, owning one's own home) than is income.*** indicates significant at the 1% level, ** indicates significant at the 5% level, * indicates significant at the 10% level. Confidence intervals are calculated by bootstrapping with 999 replications.

Table 8: The relationship between low consumption, low income and other outcomes, pensioner households

	Pensioner Single male	Pensioner, Single female	Pensioner couple
Wsh. Mch.	-0.020	-0.049**	-0.025**
Cent. Heat.	-0.030	-0.006	-0.029*
Dishwash.	-0.142***	-0.132***	-0.164***
DVD	-0.093**	-0.073***	-0.033
TV	-0.010	0.001	-0.001
Pay TV	-0.048*	-0.038**	-0.048*
PC	-0.195***	-0.122***	-0.160***
Internet	-0.147***	-0.096***	-0.175***
Car	-0.264***	-0.257***	-0.313***
Two cars	-0.042**	-0.006***	-0.154***
Own hse.	0.020	-0.066**	0.001
No. rooms	-0.283**	-0.342***	-0.342***
Health ins.	-0.056***	-0.042***	-0.029**
N	2,799	5,859	6,449

Source: Authors' calculation using Expenditure and Food Survey/Living Costs and Food Survey 2001/02- 2009.

Notes: This table gives the quantity expressed in equation (3) by household type. Negative numbers indicate that consumption is better correlated with the outcome in the left-hand column (e.g. having a washing machine, owning one's own home) than is income.*** indicates significant at the 1% level, ** indicates significant at the 5% level, * indicates significant at the 10% level. Confidence intervals are calculated by bootstrapping with 999 replications.

Table 9: The relationship between low consumption, low income and other outcomes by education and workstatus

	Age left full-time education			Employment status		
	<=16	17/18	19+	Employed	Self-employed	Workless
Wsh. Mch.	-0.087***	-0.075***	-0.055***	-0.057***	0.000	-0.084***
Cent. Heat.	-0.033***	-0.027*	0.000	-0.037***	-0.076**	-0.025***
Dishwash.	-0.144***	-0.166***	-0.145***	-0.105***	-0.238***	-0.133***
DVD	-0.199***	-0.121***	-0.135***	-0.093***	-0.089*	-0.173***
TV	-0.007***	0.000	0.010	-0.010	0.010	0.000
Pay TV	-0.145***	-0.129***	-0.077***	-0.088***	-0.175***	-0.103***
PC	-0.293***	-0.311***	-0.275***	-0.214***	-0.244***	-0.278***
Internet	-0.244***	-0.294***	-0.292***	-0.192***	-0.311***	-0.226***
Car	-0.326***	-0.346***	-0.277***	-0.258***	-0.224***	-0.276***
Two cars	-0.124***	-0.129***	-0.167***	-0.118***	-0.226***	-0.063***
Own hse.	-0.077***	-0.143***	-0.100***	-0.062***	-0.06	-0.029**
No. rooms	-0.509***	-0.694***	-0.508***	-0.368***	-0.285**	-0.468***
Health ins.	-0.034***	-0.030**	-0.065***	-0.041***	-0.092***	-0.023***
N	31,833	8,532	10,087	29,204	6,031	17,561

Source: Authors' calculation using Expenditure and Food Survey/Living Costs and Food Survey 2001/02- 2009.

Notes: This table gives the quantity expressed in equation (**??**) by household type. Negative numbers indicate that consumption is better correlated with the outcome in the left-hand column (e.g. having a washing machine, owning one's own home) than is income. *** indicates significant at the 1% level, ** indicates significant at the 5% level, * indicates significant at the 10% level. Confidence intervals are calculated by bootstrapping with 999 replications.

7 Conclusions

This paper analyses a puzzling feature of the joint distribution of income and expenditure at the bottom end of the income distribution. While median income is monotonically increasing in household expenditure, median expenditure is higher for those with *very* low income (below 10% of median income) than for those with moderately low income (below 50% of median income). In fact, those at the bottom tail of the income distribution have expenditure equal to overall median expenditure. In short, the graph of median expenditure by reported income maps out a 'tick'. 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 paper 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.

This result has very important implications for the measurement of poverty and deprivation in countries, such as the UK and US, where these are typically measured with income. Even if poverty is 'eradicated' (a policy aim for households with children in the UK), this achievement might never be evident in the data. This calls for careful and thorough consideration of how survey instruments can better capture income at the bottom of the reported distribution. It also calls for measures of deprivation other than income, including, but not limited to consumption, to be considered in official analyses of poverty.

The prospects for the greater use of consumption (derived using expenditure data) in official analyses of poverty provides a further important reason to better understand and resolve the concerning (and increasing) divergence seen in many developed countries between aggregate consumption implied by household budget surveys and that reported as part of National Income and Product Accounts.

The analysis here uses the LCFS, the only survey in the UK which includes good measures of both income and expenditure. A key question for future research is to what extent the tick appears in surveys in other countries. Another key question is what aspect of the survey design causes this mismeasurement of income.

References

Aguiar, M. A. and M. Bils (2011): "Has Consumption Inequality Mirrored Income Inequality?" NBER Working Paper 16807, National Bureau of Economic Research.

Altonji, J. G., A. Smith, and I. Vidangos (2009): "Modeling Earnings Dynamics," NBER Working Paper 14743, National Bureau of Economic Research.

Attanasio, O. P., E. Battistin, and A. Leicester (2006): "From Micro to Macro, from Poor to Rich: Consumption and Income in the UK and the US," mimeo.

Bardasi, E., S. Jenkins, and J. Rigg (1999): "Documentation for derived current and annual net household income variables, BHPS Waves 1-7," Institute for Social and Economic Research Working paper 99/25, Institute for Social and Economic Research.

Barnard, A., S. Howell, and R. Smith (2011): "The effects of taxes and benefits on household income, 2009/10. Further analysis and methodology," *Newport: Office for National Statistics*.

Barrett, G., P. Levell, and K. Milligan (2012): *A Comparison of Micro and Macro Expenditure Measures across Countries using Differing Survey Methods*, University of Chicago Press.

Bavier, R. (2008): "Reconciliation of income and consumption data in poverty measurement," *Journal of Policy Analysis and Management*, 27, 40–62.

Bound, J., C. Brown, and N. Mathiowetz (2001): "Measurement error in survey data," in *Handbook of Econometrics*, ed. by J. Heckman and E. Leamer, Elsevier, vol. 5 of *Handbook of Econometrics*, chap. 59, 3705–3843.

Brewer, M., J. Browne, and R. Joyce (2011): "Child and working-age poverty from 2010 to 2020," Institute for Fiscal Studies Commentary 121, Institute for Fiscal Studies.

Brewer, M., A. Goodman, and A. Leicester (2006): "Household spending in Britain: what can it teach us about poverty?" mimeo.

Brewer, M., A. Muriel, D. Phillips, and L. Sibieta (2008): "Poverty and Inequality in the UK: 2008," Institute for Fiscal Studies Commentary 105, Institute for Fiscal Studies.

Brewer, M. and C. O'Dea (2012): "Measuring living standards with income and consumption: Evidence from the UK," Institute for Fiscal Studies Working paper 12/12, Institute for Fiscal Studies.

Brewer, M., C. O'Dea, G. Paull, and L. Sibieta (2009): "The living standards of families with children reporting low incomes," Department for Work and Pensions Research Report 577, Department for Work and Pensions.

Browning, M. and T. F. Crossley (2001): "Unemployment insurance benefit levels and consumption changes," *Journal of Public Economics*, 80, 1–23.

Brzozowski, M. and T. F. Crossley (2011): "Viewpoint: Measuring the well-being of the poor with income or consumption: a Canadian perspective," *Canadian Journal of Economics*, 44, 88–106.

Cappellari, L. and S. Jenkins (2004): "Modelling low income transitions," *Journal of Applied Econometrics*, 19, 593–610.

Crossley, T. and C. O'Dea (2010): "The wealth and saving of UK families on the eve of the crisis," Institute for Fiscal Studies Report 71, Institute for Fiscal Studies.

Cutler, D. M. and L. F. Katz (1992): "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's," *The American Economic Review*, 82, pp. 546–551.

Daffin, C. (2009): "Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2006/2008," Tech. rep.

Deaton, A. S. (2005): "Measuring Poverty in a Growing World (or Measuring Growth in a Poor World)," *The Review of Economics and Statistics*, 87, 1–19.

Department of Social Security (1993): "HBAI 1979-1990/91," Tech. rep., London: TSO.

Jenkins, S. P. (2011): *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*, no. 9780199226436 in OUP Catalogue, Oxford University Press.

Lynn, P., A. Jäckle, S. P. Jenkins, and E. Sala (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*, 175, 289–308.

MaCurdy, T. E. (1982): "The use of time series processes to model the error structure of earnings in a longitudinal data analysis," *Journal of Econometrics*, 18, 83–114.

Meghir, C. and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72, 1–32.

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, pp. 1180–1220.

——— (2004): "The effects of welfare and tax reform: the material well-being of single mothers in the 1980s and 1990s," *Journal of Public Economics*, 88, 1387–1420.

——— (2008): "Changes in the Consumption, Income, and Well-Being of Single Mother Headed Families," *American Economic Review*, 98, 2221–41.

——— (2011a): "Consumption and Income Poverty over the Business Cycle," .

——— (2011b): "Viewpoint: Further results on measuring the well-being of the poor using income and consumption," *Canadian Journal of Economics*, 44, 52–87.

Poterba, J. M. (1989): "Lifetime Incidence and the Distributional Burden of Excise Taxes," *American Economic Review*, 79, 325–30.

Sabelhaus, J. and J. Groen (2000): "Can permanent-income theory explain cross-sectional consumption patterns?" *Review of Economics and Statistics*, 82, 431–438.

Saunders, P., J. Bradshaw, and M. Hirst (2002): "Using household expenditure to develop an income poverty line," *Social Policy and Administration*, 39, 217–234.

Slesnick, D. T. (1993): "Gaining Ground: Poverty in the Postwar United States," *Journal of Political Economy*, 101, 1–38.

Storesletten, K., C. Telmer, and A. Yaron (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.

A Appendix to Section 5

A.1 Parametrization and Solution

The model presented in Section 5 cannot be solved analytically so is solved numerically by iterating on the policy function.

Table 10 shows the parameter values we impose in the model solution. We take γ from micro studies. r is observable as the long-run average risk-free return. We calibrate β to match a steady-state asset-income ratio of just over 3.[27] The parameter of persistence, ρ is hard to identify. The value of 0.95 is within sensible bounds from micro studies.[28] We impose a long-run variance of (log) incomes (σ_y^2) of 0.22 to match the consumption of the 1-10% productivity types. Together, ρ and σ_y^2 imply a variance of idiosyncratic shocks of 0.022, which lines up well with micro studies (for example in Meghir and Pistaferri (2004)). The choice of σ_y^2 is more subtle than first appears. In this simulation study we are abstracting from measurement error and all transitory shocks other than those into and out of poverty. Therefore we cannot take this moment straight from cross-sectional income data. We document robustness checks of varying ρ and σ_y^2 in appendix Section A.2 below.

Table 10: Key Parameters

Parameter	Description	Additional Information	Value
ρ	Income Persistance	Also try $\rho = 0.85$	0.95
β	Discount Factor	Fit to match wealth/income ratio	0.9724
r	Risk-free Interest Rate	From BoE	0.018
γ	Coefficient of Risk Aversion	Matches Attanasio and Weber (1995)	1.43
σ_y^2	Steady state variance of incomes	Matches cons. of 1-10 pctiles	0.25

To simulate the model we draw permanent incomes from a log-normal distribution with variance (of the underlying gaussian) of σ_y^2, as above. The upper 99% of this distribution forms the distribution of base persistent productivities. We denote mean income of the bottom 1% by $\underline{y}$. This is the income households receive if they enter poverty.[29] We simulate a population of 1000000 households (1000 for each income type) and compute average consumption for each of these groups and for the bottom centile at the steady state level of assets. We then compute each group's consumption as an average of median consumption. We do so for both the natural- and the non-zero borrowing limits. We calibrate β to be 0.9778 without constraints and 0.9679 with constraints. Households have a greater reason to save with constraints so they must be more impatient to fit the same wealth moments.

[27] This is common in the macro literature, see for example Storesletten et al. (2004)

[28] See e.g. MaCurdy (1982) and Storesletten et al. (2004)

[29] In the data obviously, the bottom 1% of income is distributed continuously. However, in this model we cannot distinguish the bottom 1% by type.

A.2 Robustness checks

We now display robustness of the results by displaying some alternative choices for ρ and σ_y. As described in Table 10 we now let $\rho = 0.85$. We adjust the variance of productivities to match the same dispersion in consumption. The variance of productivities is now 0.5. We then adjust the variance of shocks in order to match the same dispersion in incomes. The resulting variance of yearly shocks is 0.14, which is counterfactually high. We fit β to be 0.9697 without constraints and 0.9444 with constraints.

Figure 13 shows the simulated plots of median consumption as a proportion of overall median consumption with and without constraints with the alternative parametrization. Figure 14 shows the distribution of wealth under the 2 scenarios. We see that the results are similar to before. The average consumption drop on the low shock is 8% with no constraints and 23% with constraints.

Figure 13: Simulated Consumption Across the Income Distribution: $\rho = 0.85$

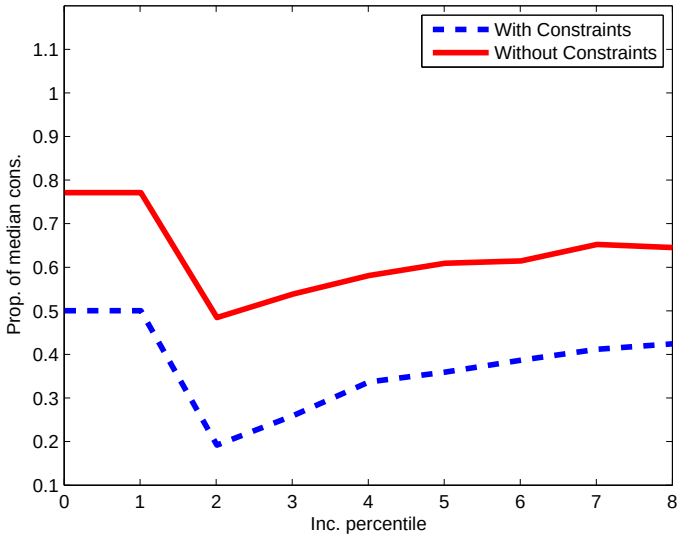

Figure 14: Simulated Wealth/Income Ratio Across the Income Distribution: $\rho = 0.85$