How should we, and how do we, measure living standards? An Appendix to *What Do We Know and What Should We Do About Inequality?*

In *What Do We Know and What Should We Do About Inequality?*, I showed trends in inequality in income, earnings and wealth in the UK. When looking at income, I used two different measures. The main measure was the 'net income of a household, adjusted for household size' (disposable income), but when looking at the income shares of the very rich I used 'gross (pre-tax) individual taxable income'. The main government publication that looks at economic inequality in the UK says that disposable income should be thought of as 'a proxy for the material living standards of individuals or, more precisely, for the level of consumption of goods and services that people could attain' (DWP (2019)). In other words, we look at net income because it gives us a good guide to what standard of living a household could obtain.

This Appendix discusses the issues involved in defining and measuring income (for more on the issues covered, see Atkinson (2015)).

Is income the best measure of living standards?

An alternative to income is to measure a household's 'consumption', or consumption expenditure, as a proxy for its standard of living. 'Consumption' here is calculated as the value of all the things that a household spends money on in a certain period (with some exceptions), plus an estimate of the benefits that a household gets from the goods that it already owns, the

main ones being housing and vehicles. At any point in time, a household's consumption and income could be different. There are two reasons. First, households can borrow or save: if a household borrows, or runs down its savings, then it can spend more than its income; if it is saving, or paying off debts, then consumption will be lower than income. Second, households that already have a lot of possessions may well have a level of consumption that is a lot higher than their income.

Economists have long argued that a household's current level of consumption should in theory be a better guide to its long-run, or average, standard of living than its current income. Others have argued that income can be hard to measure accurately among households with low resources, and that this is not true for expenditure (see Brewer et al. (2017) or Meyer and Sullivan (2017, 2018) and references therein). *What Do We Know and What Should We Do About Inequality?* does not use consumption to look at economic inequality, largely for two reasons. First, the household survey data in the UK on expenditure is not as rich as that on income, and it is impossible to use it to look at the very rich (as we can with top incomes if we use the tax data). Second, there is a concern that, in recent years, the household survey data in the UK on expenditure is not very accurate, and that it significantly underestimates spending among well-off households – not just those in the top 1% or 2% but for most of the top decile group (for the UK, this can be seen in trends of consumption spending shown in Brewer and O'Dea, 2012).

There are also ways of measuring living standards that do not involve financial concepts, including multi-dimensional measures or subjective measures.

Where does the data on income inequality in the UK come from?

The main facts on income inequality in the UK are derived from data from a large-scale household survey that is run by government statisticians. Every year, they ask tens of thousands of adults to tell them about their income and other issues. The data is now published every year by the government and, thanks to heroic efforts by researchers at the Institute for Fiscal Studies, it is also available in a consistent form back to 1961. The data series is called 'Households Below Average Income', which is a

misleading name as it covers the whole of the income distribution (there is another UK government publication on inequality, the latest one being ONS (2019a), which uses a different and smaller household survey). At the time of writing, the latest edition was DWP (2019).

But is it any good? The underlying survey and questionnaire used to collect the information on people's incomes is thought of as being high quality, and often known as the gold standard of income surveys. But no one actually knows how accurate it is. The survey spends a lot of time asking detailed questions about all sorts of aspects of different income sources, to try to avoid respondents unintentionally omitting any or double counting, and then asks people to look for evidence where possible, such as in payslips or bank statements. We know a little bit about how accurate people are at remembering what they get paid (they are pretty good), and how much money they get in social security benefits (they are less good), because we can compare what they say with what their payslips say, or the Department for Work and Pensions says it pays out. But we do not have anything to compare what people tell us with what we think really is a household's disposable income: no single institution knows for sure about all the different sources of income coming into a household.

Also, and technically speaking, I should call the information about income inequality that I show in *What Do We Know and What Should We Do About Inequality?* estimates of the true level of income inequality. No one knows for sure the income of every household in the UK – not even the government – so what I show in *What Do We Know and What Should We Do About Inequality?* is a guess (or an estimate) based upon these tens of thousands of interviews. Because so many households are interviewed, the estimates (or guesses) are pretty accurate, but inevitably they are better guesses in some years than in others. If, in a particular year, the survey happened by accident to go to (for example) fewer very rich or fewer very poor households, then we might underestimate what inequality really is. For this reason, one should not make too much out of the wobbles in the level of inequality shown in (for example) Figure 3.1 of *What Do We Know and What Should We Do About Inequality?*: they might just represent this so-called 'sampling error'.

There are also reasons to think that household surveys are not very good at working out how high are the very highest and how low are the very lowest incomes. At the bottom of the distribution, we do know that the people who report zero or very low incomes over the past month

often appear to be well off, if we look at the sort of houses in which they live, the stuff they own, the amount they say they spend each month, or whether the state of their finances makes them stressed (see Brewer et al. (2009) or Brewer et al. (2017)). Some people do have incomes that vary dramatically from month to month, of course, but it looks very likely that part of the discrepancy between the very low income that is reported and their apparently high standard of living is due to some individuals not reporting big chunks of their income (and I showed in Chapter 1 that, in the most used data on income inequality in the UK, about 0.6 million people seem to be surviving on no income at all). Because of this, one should not pay too much attention to the level of income reported by the bottom 1% or 2%.

It is also hard to record accurately the incomes of the rich in a household survey. First, very rich people tend to be very busy and not especially keen on talking to government-backed survey interviewers. Second, very rich people have complicated financial set-ups, and even if an interviewer could get through the front door, it might be hard to learn enough in a typical interview of 60–90 minutes. But we do know a lot about the incomes of the very rich from the information that they report to the tax authorities. Although every individual's own tax return is confidential, it is possible for researchers in the UK to get hold of a sample of tax returns, suitably anonymised, in which it turns out that the amount of income being declared to the taxman by the very rich is a good deal more than the amount of income that appears to be being received by the very rich according to the government surveys (for the UK, this discrepancy is shown in detail by Burkhauser et al., 2018a, b). Because of this, government statisticians have long made an adjustment to the household survey data so that it better matches the information from tax data (see Burkhauser et al., 2018b; ONS, 2019b). Recently, it has become clear that it is not doing this in the most effective way, so the main statistics on overall inequality in the UK are probably still underestimating the true level (it is not possible to use the tax data to estimate the Gini coefficient, because this data does not cover the income received by the whole population). It also means that we should not pay too much attention to what the survey data thinks is the level of income reported by the top 1% or 2%.

For these reasons, researchers who want to zoom in on the circumstances of the very rich prefer to use tax data rather than survey data. The idea of using information from tax authorities to learn about top incomes

was pioneered by a team of economists including Thomas Piketty (of *Capital* fame), Emmanuel Saez and the late Anthony Atkinson. They and their team of collaborators have now built an enormous database of measures of income and wealth inequality over hundreds of countries (most countries have a form of income tax), and, in some cases, over the past hundred years (this work was first collected together in Atkinson and Piketty (2007), and is now available at the World Inequality Database https://wid.world). They argue that the advantages of using data from tax authorities to learn about the income of the very rich are that there is lot of data (because everyone has to pay taxes), the information is usually confirmed against what employers and financial institutions think to be the case, and there are stiff penalties for getting it wrong. On the other hand, tax authorities care only about the sort of income that is taxable, and they know only about income it if is declared to them: we can see in the UK data that changes to tax rates and regulations since 2010 do affect whether, how and in what form the very rich report their income to the taxman. (You may well wonder whether the very rich would declare absolutely all of their income to the taxman, but tax evasion is illegal, so I am sure that they would declare everything) Chapter 3 of *What Do We Know and What Should We Do About Inequality?* features some new analysis of the tax data for the UK to show levels of top incomes, changes in top income shares, and what sort of people can be found at the very top of the income distribution.

Issues to consider when measuring income

To design the questions for the household survey that underpins the statistics on inequality, and to turn the survey answers from a household into a single number of 'income', statisticians have to decide matters such as over what period to measure income, whose income to measure, and what counts and does not count as income. International best practice on how to define and measure income, drawn up under the auspices of the UN, can be found at UNECE (2011).

The basic idea is that income is all money coming into the household, over some period, that is available to spend. As done in *What do We Know and What Should We Do About Inequality?*, total income can then be broken down into income from employment or self-employment, income from financial assets, income from social security benefits or tax credits (government transfers) and other sources. This definition means

that what economists call 'home production' does not count as part of income. Home production is those products or services that households produce for themselves – not just whether they produce their own food, but also basic housework, childcare, and so on: that is, any activity which would otherwise have to be done by someone else in return for payment. It seems reasonable for home production not to count as income: I have to be able to spend income on buying stuff, but I cannot use the value of my services providing childcare to my children to buy anything. But not counting home production as income is one way in which there might be a difference between a household's living standard and its income (e.g. compare two couples with children with the same total disposable (cash) income, but where one couple has two full-time workers and the other has one full-time worker: we might think that the latter is better off (in some sense) than the former).

A comprehensive measure of income clearly should count income that comes to households directly from government, such as social security benefits or tax credits. A grey area is whether to count so-called 'in-kind transfers'. This is the name given to government programmes where a household is entitled to a service (usually for free). For example, in England, parents of 3- and 4-year-old children are entitled to use free childcare; children, those over the state pension age and people receiving certain social security benefits are entitled to free prescriptions; children whose families receive certain social security benefits are entitled to free school meals; those aged 75 or over are entitled to a free TV licence. These entitlements do not seem like 'income' – you cannot actually spend them to buy other goods or services, like you can spend payments of child benefit or the state pension. On the other hand, most people would think that families receiving these in-kind transfers were better off than families that do not: being offered free school meals must save families money, otherwise they would have to spend money providing children with lunch at school. For this reason, the definition of income used in the main UK series on inequality ('Households Below Average Income') does count some of these as income (specifically: free school meals and free TV licences). But once you start counting some of these public services as income, it is hard to know where to stop. What about counting the value of free education? Or the fact that the NHS is free at the point of use to all? Some researchers do add the value of public services to household income (see e.g. Paulus et al., 2010), but this is

not the approach that I took in *What Do We Know and What Should We Do About Inequality?*.

What counts as 'income'?

In a household's budget, the difference between 'income' and 'spending' has to be equal to the change in the household's net worth. If a household spends less than its income, then it will have managed either to increase the amount of savings it has, or to pay off some debt. If it spends more than its income, then it has either increased its stock of debts, or run down some of its savings. We can write this as:

> Income – spend = change in net worth (where 'net worth' = savings – debts)

Now let's rearrange that word equation to give:

> Income = spend + change in net worth

We can then use this as a way of defining income; it helps a lot when considering the following special cases.

***Capital gains*.** Consider a couple who own shares worth £10,000. In 2018, imagine that the company pays a dividend payment of £300 on these, and that the value of the shares rises by 10%. If the couple had no other sources of income, then we would say that their disposable income, according to the definition used in *What Do We Know and What Should We Do About Inequality?,* was just the value of that dividend payment, of £300. However, thanks to this £10,000 of shares, by the end of the year the couple are richer to the tune of £1,300. If the couple did not spend any money over the year (not realistic at all, but it makes the maths easier), and their net worth has increased by £1,300, then their income would have to be £1,300 too, to make the word equation above balance.

The other £1,000 is known as a 'capital gain'. Capital gains can be 'realised' or 'unrealised'. If I buy £1,000 of shares now and sell them a year later for £1,100, then I have realised a capital gain of £100, because I have £100 more cash in my hand that I can use to buy other things. If I buy £1,000 of shares now and their value increased to £1,100, then I have an

unrealised capital gain. Realised capital gains are usually subject to tax, but unrealised capital gains are not. However, both should count in a comprehensive measure of income because capital gains make people better off. For people who own property, or who own financial assets, unrealised capital gains are an important part of their overall income, and one that is missed by the usual data on living standards and inequality.

Contributions to pensions. For reasons that are not clear, the measure of income used by government statisticians in the UK takes your headline pre-tax earnings and then subtracts from those any payments that are made from your salary into a private pension (it is common to save in this way to take advantage of tax relief on pension saving). But this is not consistent with how we treat people who save in other ways, and it is not consistent with the definition of broader income that we introduced above.

To see this, consider two people on identical salaries of £1,500 after tax. Person A saves £100 a month into a private pension in the way that we describe above, and Person B saves into a regular bank account; both spend the remaining £1,400. Person A will appear to have a lower disposable income, at £1,400, than Person B, but this makes little sense: both earn the same (and both have the same amount of money to spend on stuff that is not savings). Alternatively, we could use the word equation above. Both individuals have increased their net worth by £100 over this past month. If we assume that both have spend £1,400 a month on things that are not savings, then for both their measure of income should be £1,500 a month.

Hopefully, this will convince you that we ought to count the money that people save as part of their income, rather than pretending that money saved is lost in the same way that income tax payments are lost. However, money that people save into pensions is missed from the conventional statistics on living standards and inequality in the UK. This gets even more complicated for people who have 'final salary' pensions, such as teachers, nurses, university lecturers, MPs and other workers in the public sector. For these people, the amount of pension that they will be paid depends not on what they have paid in, but on complicated formulae that relate how many years they have been working and their salary over their working lifetime. What this means is that the value of their pension goes up each year that they work – because, when they retire, they will be entitled to a larger pension having worked that extra year. So their net worth has definitely gone up, because they are going to receive a higher income over some

portion of their life. In principle, this should be accounted for in this wider definition of income; in practice, it is hard to calculate by exactly by how much the value of the pension has gone up.

***Housing*.** The measure of disposable income also misses out on something that is a very large factor in determining someone's living standards: whether they own the house in which they live. This is a different, and more subtle, issue than the fact that disposable income is a different concept from the stock of wealth or assets.

Imagine three houses in the same street, all occupied by people earning the same monthly salary of £3,000. But they differ in their living arrangements. Person A pays rent of £800 a month. Person B owns her house, which is otherwise identical to Person A's, outright. Person C pays rent of just £600 a month, because the house she lives in is smaller than A's or B's. Also, Person C owns the house that Person A lives in, so she has an extra income from the rent of £800 a month.

Which of the three is the worst off? If we rank them by their disposable income, then Person A and B are tied for last place, on £3,000 a month. But most people would agree that Person B, who owns her own house, must be better off: Person B can afford to spend £3,000 on things other than rent (without going into debt), whereas person A can afford to spend just £2,200. Another way of thinking about this is that Person B must be better off than Person A because she is living in an identical house and can afford to buy all the same things as Person A, but she would still have £800 a month left over. So Person B is better off than A, despite having the same disposable income using the standard definitions.

But what about Persons B and C? Person C has the greater disposable income, for sure, at £3,800 vs £3,000. But what about the comparison of how much they can spend on stuff that is not rent? Well, Person B can spend £3,000 and C can spend £3,200. So perhaps Person C is better off? Maybe. (This concept of how much they can spend on stuff that is not rent is used in official statistics, where it is known as 'income after housing costs', meaning income after we have subtracted housing costs.) A comparison of income after housing costs suggests that Person C is better off, on the grounds that she has more money to spend on things that are not rent. But this misses out on the fact that Person C lives in a cheaper and, presumably, less nice house. Does that not count for something?

Another (better!) way to compare Persons B and C is to realise that they both have the same earnings, and both are property owners, owning identical houses, so they must be as well off as each other. What makes the comparison seem complicated is that, although both Person B and C could choose to get some income from their property by renting it out, Person B has decided to live in her own house and, it appears, to go without £800 a month in rental income. But a better way to think about Person B is to think of her as if she is renting out the house she owns to herself, and additionally receiving a rent of £800 a month, the same as Person C. Person B would then have an income of £3,800 a month, and be spending £800 a month on rent. In effect, they both have the same potential income of £3,800 a month and are equally well off. If you do not like that comparison, then how about this? Person C could, if she chose, rent a house like B for £800, leaving her with £3,000 to spend on stuff that is not rent. Also, Person B could move out, get a tenant, have an income of £3,800 and rent a house identical to C for £600, also leaving her with £3,200 to spend on stuff that is not rent. If, without borrowing any money, they can both afford to live the same lifestyles as each other, then they must be equally well off.

The £800 of potential income that we gave to Person B is called the 'implicit rental income'. Because she owns a house, Person B could get this rental income; the only reason we do not see it as an actual income is that *she rents it to herself*. This implicit rental income is missed by the concept of disposable income usually used in the UK. But it is commonly counted as income by researchers in other countries, and there is no technical reason why it could not be done in the UK. Further, because houses tend to be owned by those who are richer or older, looking at a measure of income that includes implicit rent is going to give us a different impression of the nature of income inequality in the UK. This lies behind one of the recommendations I make in Chapter 4 of *What Do We Know and What Should We Do About Inequality?*

Whose income do we count?

In many cases, people living together in a household share their resources in some way. Because of this, the main statistics on oncome inequality in the UK add up the income of all the people who live in the same household and assume that this is shared out equally. This sometimes feels

like a perfectly reasonable thing to do. For example, let's suppose that a rich CEO has a stay-at-home husband, with no income of his own, who looks after the household while she is out earning mega-bucks. Even if his wife also owns all the assets, the house-husband's standard of living will surely be a lot higher than that of an unemployed man living alone. Similarly, we do not think that all children under 16 are poor just because they themselves have no or very little income. In these cases, the idea that everyone in the household has a similar standard of living does not seem too unreasonable.

At other times, however, this might be a less good approximation of the truth. Imagine a single household that contains two or three generations of the same family living together. In some cases, even though they might share some household expenses or have meals together, the house-sharers might keep their income and their outgoings separate from their family members in other respects. In this case, adding up the income of everyone in the household and assuming that all family members have equal access to this may not be so accurate.

On balance, though, if we are trying to assess how much inequality there is in the standard of living that people enjoy, then it is probably less wrong most of the time to work out the total income of a household than it is to look at each individual's own income: this is what is done in the main data on inequality that is used in *What Do We Know and What Should We Do About Inequality?* for both income and wealth. (The exception is that the data from tax authorities used for the analysis of the very rich or for top-income shares is all measured at the individual level, because the underlying data does not tell us anything about who lives with whom.) One implication of this is that we cannot look at inequalities between (say) men and women living in the same household (it is possible to look at inequalities in individuals' income, and this is done for example in section 2.4 of NEP, 2010).

How do we adjust for household size?

There is one further wrinkle to this. Having added up the total income of the household, we then usually adjust it to reflect the number of adults and children living there. This adjustment is meant to recognise that there are savings to be made when living together compared with living apart, so a couple living together (say) generally need less than twice

the income of a single person to reach the same standard of living. This adjustment is known as equivalisation, and what we calculate is something like *income per equivalent adult.* Here, I explain in detail why such an adjustment might be sensible.

Let's say we have four households living in identical houses. Household A contains a single person with an annual disposable income of £25,000. In Household B, there is a married couple, and their combined disposable income is also £25,000. Household C contains a couple, but here the combined disposable income is £50,000. Household D contains a married couple with children aged 6 and 10, and has a combined income of £50,000. How should we rank these households from the highest to the lowest standard of living? Most people would agree that Household A is better off than Household B, because B has the same income but more mouths to feed, as it were. But what about Households A and C? Both have the same income per person, so perhaps they are equally well off? In fact, it will generally be the case that Household C is better off than A. Certainly Household C will need to have an income of more than £25,000 to be as well off as Household A: there are more mouths to feed, after all. But there are savings to be made when living together compared with living apart, such as on things like heating or lighting and food, so Household C should be able to enjoy the same lifestyle as A with less than twice the income.

So Household C is better off than A and both are better off than B. What about Household D? Well, we can use similar logic to argue that it is probably worse off than Household C (which has the same income, but fewer mouths to feed), and probably better off than Household B (which has half the income and half the mouths to feed). It is less obvious what we should conclude when comparing Households A and D, though. In fact, when doing an analysis of inequalities in household income, researchers tend to make assumptions that imply that a single person with an income of £25,000 a year is as well off as a couple whose joint income is £37,500 a year, and as well off as a couple with two children with a total income of £52,500 a year. They do this by dividing each household's total income by a number which can be thought of as the number of equivalent single adults living in it. For a couple, the magic number is 1.5 (meaning that we think that a couple needs just one and a half times as much income as a single person to be as well off), and

for a couple with two children, it is 2.1 (meaning that a couple with two children need just over twice as much income as a single person to be as well off). If we do this for our four mythical households, then we would say that Household A has £25,000 in equivalent (or equivalised) income, Household B has £16,667 equivalent income, Household C has £33,333 and Household D has £23,810, so the ranking from richest to poorest is C, A, D, B. It means that we think that Household B, with its two adults on a combined £25,000, *is as well off as if it were* a single adult on £16,667 a year.

So what we have calculated here is something like *income per equivalent adult*, the intuition behind it being that it is cheaper to live apart than separate. Unfortunately, it is common in the UK to express household incomes as 'income per equivalent couple with no children'. This reflects entirely the same ideas, but the sums are different. This time we leave the incomes of the households that contain couples with no children – Households B and C – alone and adjust only the others. For Household D, the magic numbers imply that a couple with two children need 60% more than a couple with no children to be as well off, so we divide their income by 1.6, giving us £31,250. In other words, we think that a couple with two children on £50,000 a year are as well off as the same couple without children on £30,120. For our single adult living alone in Household A, we divide her income by 0.66 to get £37,878 in *income per equivalent couple without children*, meaning that she, living alone on her £25,000, is as well off as a couple without children living on £37,878. The ranking of the four households does not change whether we use income per equivalent adult or income per equivalent couple without children, nor does the size of the gaps between them, but obviously the actual numbers in equivalised pounds do. This means it is important to check the small print when reading off graphs with income on one of the axes, and, to help, I provide some conversions in Table A.1. The table gives the value in pounds of different centiles on the income distribution in 2016–17, and shows what this would be for different family types. It shows that, for example, a couple without children that has an income of above £962 is in the top decile group (the top 10%). A single person living alone is in the top decile group with a disposable income of just £645, and a couple with two children needs to have a disposable income of just £1,347 to be in the top decile group.

Table A.1 Different centiles of the UK income distribution in 2016–17 for different family types (£/week)

	Couple with no children	Single adult	Lone parent with one child	Couple with two children
10th centile	248	166	216	348
60% of median income	296	198	258	415
50th centile (median income)	494	331	430	691
Mean (average) income	594	398	516	831
90th centile	962	645	837	1,347

Over what time period should we measure income?

Income has to be measured over a period of time, such as a year, or month or week (we say that income is a 'flow', whereas wealth is a stock and has to be measured at a specific point in time). The usual practice in the UK is to measure income over a fairly short period – the last few weeks, more or less – and express it in a weekly amount. Most other countries measure income over a year, and express it as an annual or a monthly amount. We would expect inequality to be lower the longer the period over which income is measured, because the longer the period, the more likely that extremely low or high values of income will get partially offset by less extreme values (this is 'regression to the mean'). At the limit, one could even measure income over a lifetime. Inequality in lifetime incomes is more important than inequality in incomes over the previous month, because we would normally be less concerned about income inequality if individuals were moving about the income distribution from month to month, so that those who are well off now are likely to be poor later on. However, we can only measure lifetime income once a person reaches the end of their lifetime; it is an excellent concept in principle, but hard to use in practice. And Jenkins (2011) shows that although incomes do change and people do move around the income distribution, the vast majority of movements are over a short distance.

References

Atkinson, A. (2015) *Inequality: What Can Be Done?* Cambridge, MA: Harvard University Press.

Atkinson, A. and Piketty, T. (eds) (2007) *Top Incomes Over the Twentieth Century.* Oxford: Oxford University Press.

Brewer, M., Etheridge, B. and O'Dea, C. (2017) Why are Households that Report the Lowest Incomes So Well-off?. *Economic Journal*, 127, F24–49. doi:10.1111/ecoj.12334

Brewer, M., O'Dea, C., Paull, G. and Sibieta, L. (2009) *The Living Standards of Families Reporting Very Low Incomes*, Department for Work and Pensions Research Report 577.

Brewer, M. and O'Dea, C. (2012) Measuring living standards with income and consumption: evidence from the UK, ISER WP 2012–05, www.iser.essex.ac.uk/research/publications/working-papers/iser/2012-05

Burkhauser, R., Hérault, N., Jenkins, S. and Wilkins, R. (2018a) Top Incomes and Inequality in the UK: Reconciling Estimates from Household Survey and Tax Return Data. *Oxford Economic Papers*, 70, 301–26. doi:10.1093/oep/gpx041

Burkhauser, R., Hérault, N., Jenkins, S. and Wilkins, R. (2018b) Survey Under-Coverage of Top Incomes and Estimation of Inequality: What Is the Role of the UK's SPI Adjustment? *Fiscal Studies*, 39(2), 213–40.

DWP (Department for Work and Pensions) (2019) Households Below Average Income: An analysis of the UK income distribution: 1994/95–2017/18, www.gov.uk/government/statistics/households-below-average-income-199495-to-201718

Jenkins, S. (2011) *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*. Oxford: Oxford University Press.

Meyer, B. and Sullivan. J. (2017) Consumption and Income Inequality in the US since the 1960s. NBER WP 23655.

Meyer, B. and Sullivan. J. (2018) Consumption and Income Inequality in the US since the 1960s. [blog], 15 January, https://voxeu.org/article/consumption-and-income-inequality-us-1960s

NEP (The National Equality Panel) (2010) *An Anatomy of Economic Inequality in the UK*, London School of Economics CASE report 60.

ONS (Office for National Statistics) (2019a) The Effects of Taxes and Benefits on Household Income, Disposable Income Estimate: 2018, www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/datasets/householddisposableincomeandinequality

ONS (Office for National Statistics) (2019b) Using Tax Data to Better Capture Top Earners in Household Income Inequality Statistics, www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/articles/usingtaxdatatobettercapturetopearnersinhouseholdincomeinequalitystatistics/2019-02-26

Paulus, A., Sutherland, H. and Tsakloglou, P. (2010) The Distributional Impact of In-kind Public Benefits in European Countries. *Journal of Policy Analysis and Management*, 29, 243–66. doi:10.1002/pam.20490

UNECE (United Nations Economic Commission for Europe) (2011) *Canberra Group: Handbook on Household Income Statistics*, Second Edition, ECE/CES/11, www.unece.org/fileadmin/DAM/stats/groups/cgh/Canbera_Handbook_2011_WEB.pdf