

Response of IFS researchers to “Measuring Child Poverty: a consultation on better measures of child poverty”

James Browne, Jonathan Cribb, Andrew Hood, Paul Johnson, Robert Joyce, Cormac O’Dea, David Phillips and Luke Sibieta

Institute for Fiscal Studies

Summary

The government is consulting on defining a ‘multidimensional’ measure of child poverty.¹ It has stated that it wants the new measure to “combine a range of indicators into a single number”. It has suggested eight specific dimensions: income and material deprivation; worklessness; unmanageable debt; poor housing; parental skill level; access to quality education; family stability; and parental health.

We recommend that the government defines distinct sets of indicators that will be used to measure children’s current material living standards; children’s wider wellbeing; the causes of child poverty; and the future life chances of children. We recommend against combining multiple indicators into a single index of child poverty. It would be difficult or impossible to do so in a way that is understood and that achieves consensus – two of the stated aims. And the particular dimensions suggested cover distinct concepts, each providing valuable but different kinds of information. Aggregating them would result in something less informative than multiple indicators.

The consultation document sets out a range of poverty-related indicators. The Child Poverty Strategy,² published in 2011, had also done this. But IFS researchers had argued for a more formal statement of how the numerous and wide-ranging set of potential indicators relate to each other, and how each ultimately relates to child poverty.³ The consultation offers an opportunity to establish this, and we therefore welcome it.

Our reading of the consultation document suggests that the government actually has at least four things in mind: a) children’s current material living standards; b) other inputs to their wellbeing; c) the causes of low living standards; and d) the future life chances of children who grow up in poor families. These are all important issues.

However, indicators belonging to each of these categories should be kept separate. They are different concepts. In particular, they are not all measures of child poverty. It therefore does not make sense to combine them into a single number. This would be difficult or impossible to interpret.

¹The consultation document is here:

<https://media.education.gov.uk/assets/files/pdf/m/measuring%20child%20poverty%20consultation%20document%20final.pdf>. Henceforth we refer to this as “Measuring Child Poverty”.

² <https://www.education.gov.uk/publications/eOrderingDownload/CM-8061.pdf>.

³ Cribb, Joyce and Phillips (2012).

Our suggested approach would give policy-makers an incentive to tackle the long-run causes of child poverty, and of poor adult outcomes stemming from childhood deprivation. Such incentives would exist even if the ultimate benefits of policy action – for levels of child poverty, or the adult lives of current children – would not be fully realised for many years, beyond the political cycle. IFS researchers have argued that one risk of fixating too much on current income is that this may skew the political incentives towards policies which have immediate and predictable effects on current incomes, like increasing benefits.⁴ There is no question that income-based child poverty measures are highly – and immediately – sensitive to fiscal redistribution.⁵ A government aiming for quick results with these measures may well therefore have too strong an incentive to spend money on increasing benefits, rather than other things which may ultimately be more powerful and cost-effective at improving lives.

But our approach would avoid confusing the causes of child poverty with child poverty itself, or with the adult prospects of current children. It would also separate material aspects of low living standards (by which we mean a lack of economic resources) – which is typically what people have in mind when they think of someone as ‘poor’ or ‘in poverty’ – from other aspects of ‘quality of life’ or ‘wellbeing’. An approach to poverty measurement that is credible, meaningful and, as the consultation document puts it, “understood and accepted”,⁶ must recognise these conceptual distinctions.

There is then a further question: *within* the conceptually distinct categories that we have outlined, should different indicators be combined into one index? In principle, if one has several indicators of (e.g.) material living standards and wants to form a comprehensive view, combining them in some way does look appealing. But there are serious practical obstacles, and we argue that the resulting need to attach weights to different indicators is incompatible with the stated aims of defining a measure which achieves a broad consensus and is widely understood. It is also not clear how a multidimensional index could cope with a (perfectly sensible) decision to track both absolute and relative measures of poverty or life chances.

This note proceeds as follows. We begin in Section 1 with discussion of the current headline income-based measures, as in Chapter 1 of the consultation document. There are various potential limitations to these measures, but it is not always clear which of these the government is most concerned about, or even whether it appreciates the distinctions between them. In Section 2 we discuss the principle of tracking multiple poverty-related indicators. We argue that this is desirable, but in doing so it is crucial to be clear about how each indicator actually relates to child poverty. Finally, in Section 3 we discuss the idea of aggregating indicators together into one index. We argue that there are major disadvantages of doing this.

⁴ For example, see the response of IFS researchers to the consultation on child poverty under the previous government in 2009: http://www.ifs.org.uk/docs/cp_consultation09.pdf.

⁵ Cribb, Joyce and Phillips (2012); Dickens (2011); Brewer, Browne, Joyce and Sibieta (2010).

⁶ “Measuring Child Poverty” consultation, pp. 16.

1. Current measures: areas of contention

The government clearly does not believe that current income alone constitutes a complete measure of child poverty. There are several references in the consultation document to the notion that current income alone is too narrow. We agree with this notion. In the next subsection we set out the reasons for this belief, which also lead naturally to suggested solutions.

However, Chapter 1 of the consultation document conflates distinct issues relating to the current income-based measures of child poverty set out in the Child Poverty Act. One consequence is that the limitations of income-based poverty measures specifically are confused with much more general issues related to poverty measurement. Some of these issues are not in fact specific to income-based measures of poverty: they would apply whether measuring income or something else. Greater clarity over the government's main areas of concern is needed, as different problems suggest different solutions.

In particular, it is crucial to distinguish the narrowness of just measuring income from the following issues:

- i) the choice of absolute or relative poverty lines;
- ii) the choice of 'binary'/'headcount' measures of poverty versus 'poverty gap' or 'shortfall' -type measures which are sensitive to the depth of poverty.

Much of Chapter 1 of the consultation document conflates either i) or ii) (or both) with the question of whether we measure only income or whether we also measure other things. For example, a sub-heading on page 15 says "Income does not tell us enough", but the subsequent critique applies only to *relative* income, and an analogous critique would apply to any relative measure of poverty, whether income-based or not. If this was the only limitation, one could simply use a measure based upon absolute income, as is already defined in the Child Poverty Act. (As it turns out, it is not the only limitation.)

Absolute versus relative measures

Using non-income based indicators would not avoid the need to choose between absolute and relative measures (or, as we firmly recommend, to decide to measure both). For example, if measuring educational outcomes, the government could choose to do so in a relative or absolute sense. Of course we care whether absolute standards of education are improving over time. But in judging the education levels of the poor, the context of time and place may be relevant too. In the 1950s a child leaving school at 15 would have been leaving at the school leaving age. By 2015, a child leaving school at 15 would be dropping out 3 years before the official school leaving age.

We would argue that it is not advisable to try to choose between absolute and relative measures of poverty, regardless of whether it is income that is being measured. People have long argued about whether poverty is better thought of as an absolute or relative concept, and this will not be resolved by a government consultation. Both are limited on their own. Relative measures do not reveal what has happened to absolute living standards and, as the consultation document repeatedly states, this can look particularly perverse during recessions (if considered in isolation from absolute measures). But

absolute measures cannot reveal whether there is a group towards the bottom of the distribution falling further and further behind everyone else over time, which is also a perfectly legitimate concern.⁷

In the long run, our view of what makes someone 'poor' is probably sensitive to movements in general living standards. This may be less likely in the short run: year-to-year volatility in median income may not reflect changes in what we think of as 'poor'. For these reasons, it may be reasonable to focus primarily on absolute measures when looking at year-to-year movements in poverty, whilst tracking relative measures closely when making longer-term comparisons. But more generally the measures simply provide different information, and there is nothing to be gained from ignoring some of this information.

Sensitivity to the depth of poverty

Any poverty indicator needs to define a poverty threshold (otherwise it would just be a measure of general living standards or of inequality). Having done this, one is faced with a choice of whether only to measure whether someone is above or below that threshold ('headcount' measures), or whether also to account for how far people fall below the threshold.⁸ In principle, accounting for poverty depth as well as the number below the poverty line is better because it provides more information.

This is a real issue – but again, it is not an issue specific to income-based measures. There may however be specific problems with income-based measures which *do* attempt to account for poverty depth, because they will be sensitive to inaccuracies in the recording of those with the very lowest measured incomes. There is evidence that these inaccuracies exist. This is discussed in more detail in the next section.

The limitations of income-based measures

Current income determines the flow of goods and services that a family can buy without depleting their wealth. We therefore care about it because it is a measure of material living standards and, for many families much of the time, it is quite a comprehensive measure. However, it proxies material living standards imperfectly – and this is true regardless of whether the income-based measure in question is absolute or relative, and regardless of whether it is sensitive to the distance from the poverty line.

⁷ And it is separable from inequality more generally, which depends upon inequalities throughout the distribution.

⁸ Simple 'poverty gap' indices capture the average distance from the poverty line of those in poverty. In combination with headcount measures they can in principle be useful, although in isolation they have the undesirable property that if someone just below the poverty line moves above it then the average poverty gap for those left in poverty increases. As with headcount measures, poverty gap measures have also been criticised for their insensitivity to the degree of inequality between those in poverty (for example, see Sen, 1976). More sophisticated indices have long been proposed in the academic literature – but are used much more rarely elsewhere, perhaps because they are more difficult to interpret – which are sensitive to the number below the poverty line, poverty depth for those in poverty, and inequality among the poor, as well as having various other desirable properties. For example, see Clark, Hemming and Ulph (1981).

Income is also a measure only of *material* living standards. Ideally we would also want to account for other factors when thinking about children's wellbeing (although it is less common, and perhaps unnecessarily confusing, for the term 'poverty' to be used in a non-material context). The government does appear to have non-material factors in mind – for example, in the consultation document section on 'worklessness'.

Section 2 sets out the key limitations of income in more detail, and offers suggested solutions.

2. Why and how should we look beyond income?

This Section makes two basic points. First, there are various ways in which a fuller picture of living standards can be gained by supplementing income data with other information. Second, there are several good reasons why one should want to look at things other than income in order to assess whether the ultimate causes of child poverty are being addressed (e.g. parental employment patterns) and to assess the life chances of current children (e.g. educational opportunity). But we argue that the clear conceptual distinctions between these different kinds of measures should not be obfuscated.

A comprehensive picture of living standards

Current income is, for good reason, a standard measure of material living standards. It has been used across the world for a long time by both governments and academics, and should continue to be. But it has always been acknowledged as an imperfect measure, both due to the concept of income in principle and its measurement in practice. Here we outline some of its key deficiencies, and suggest how these can be alleviated by supplementing it with additional information.

Some families have relatively high or low incomes only temporarily. Current income may be a poor proxy for the material living standards of these people. Given that incomes can effectively be transferred from one period to another via borrowing and saving (or paying off previously accumulated debt), what we would really like to observe is lifetime income.

That this can really matter for poverty measures can be seen by considering an increase in income volatility (e.g. due to an increase in short-term unemployment) versus an increase in permanent income inequality between different groups (e.g. due to an increase in long-term unemployment). Both will increase the cross-sectional dispersion of incomes, and hence will tend to increase the number of people with incomes below a poverty line. But we would expect the effects on living standards and lifetime resources to be smaller in the former case, as (at least some of) those with temporarily low incomes will be able to finance consumption either by drawing down accumulated

savings or by borrowing against the expectation of higher income in future. **A way to address this is to use poverty measures based on consumption.**^{9,10}

Consumption is generally considered a better proxy than income for lifetime resources and for current living standards.¹¹ And it can be used to assess whether changes in the distribution of current income are large and persistent enough to have caused substantial changes in material living standards. For example, relative poverty in the UK in the 1980s grew less quickly when measured using consumption than when measured using income.¹² This likely reflects the fact that some (though not all) of the increase in income inequality over the period reflected an increase in income volatility – which can, at least partly, be smoothed via saving and borrowing - rather than higher lifetime inequality.¹³

Standard income-based measures also do not account for housing quality. This is true whether they measure incomes before or after deducting housing costs (BHC or AHC).¹⁴ For example, BHC income measures do not account for the fact that owner-occupiers with the same income and housing quality as those in other tenures would be able to achieve a larger consumption flow with their income (since they do not have to spend any of it on housing costs, and can instead spend it on other things); and AHC income measures do not account for the fact that two families with equal BHC incomes and housing costs may have different qualities of housing. This latter point is particularly pertinent for the poor, who are relatively likely to be in social rented accommodation where rents do not necessarily reflect market values. **Accounting for housing quality, as the consultation document suggests, is therefore a good idea.**

A consumption-based measure of poverty could account for housing quality by imputing a consumption flow from housing, based on observed market rents for those in the private rented sector and based on information on housing quality for social renters and

⁹ A measure of consumption needs to go beyond current expenditure, because goods bought previously may still be being consumed. This is true most obviously for housing, but also for other durable goods. The aim should be to measure expenditure on non-durables plus consumption flows from durables. Ideally, a similar imputation procedure would be carried out to estimate consumption flows from other durable goods as for housing (described below). This is dependent upon having information on other durables owned, and ideally their quality. Otherwise, expenditure on new durables could be used as an approximation for consumption flows from those durables (the resulting imprecision would be prohibitive in the case of housing, but housing information already exists in the Living Costs and Food Survey). See Brewer and O’Dea (2012) for further discussion of these issues.

¹⁰ We are not recommending a wholesale replacement of income-based measures with consumption-based ones. The currently available nationally-representative expenditure data in the UK - the Living Costs and Food Survey - has a considerably smaller sample size than the Family Resources Survey used to measure income. And there are some concerns regarding measurement error in consumption as well as income (see Brewer & O’Dea (2012) for discussion). However, for the reasons just described and elaborated on further below, consumption data can be a very valuable complement to income data.

¹¹ For example, see Blundell and Preston (1996); Cutler and Katz (1992); Poterba (1989).

¹² Brewer and O’Dea (2012).

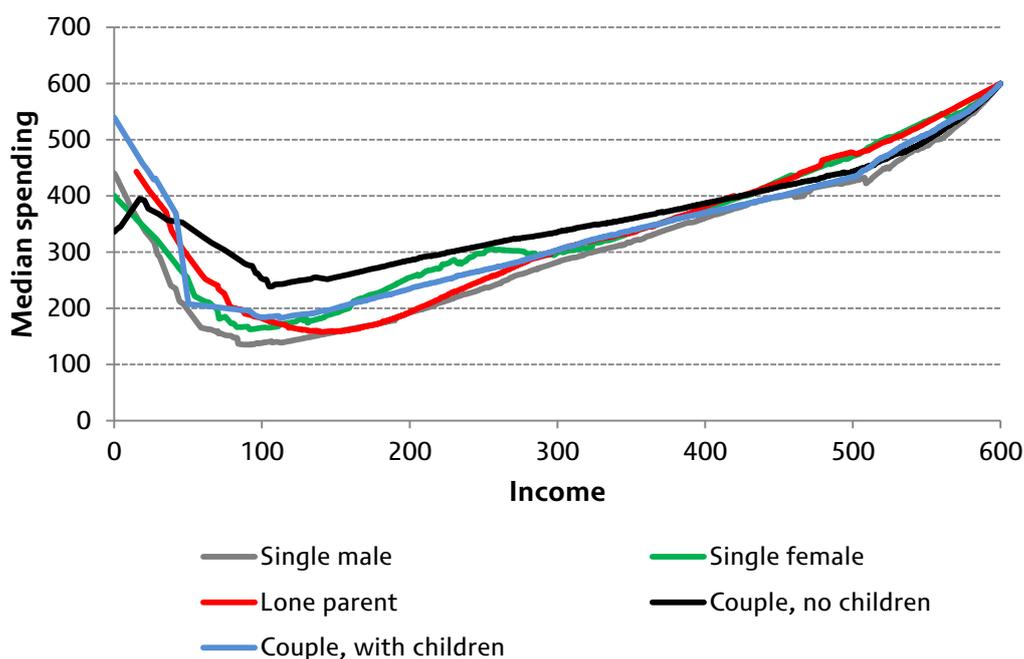
¹³ Blundell and Preston (1998).

¹⁴ The measures defined in the Child Poverty Act use before-housing-cost income only.

owner-occupiers. Income-based measures could also account for housing quality in an analogous way, by adding to cash income the imputed consumption/income flow from housing (and subtracting current spending on housing, e.g. rents and mortgage interest payments). IFS research has shown that using either consumption or this broader measure of income makes a substantial difference to impressions of relative living standards between families of different sizes, and between different cohorts, in the UK.¹⁵

A practical problem with income is that household survey datasets on which income statistics are based tend not to record income perfectly. This can be a particular problem close to the bottom of the income distribution. There is substantial under-recording of some benefit and tax credit incomes in the Family Resources Survey.¹⁶ Empirical work suggests that under-reporting of income is in fact the primary reason why income is an imperfect measure of low living standards in the UK: as shown by Figure 1, those families with children with the lowest measured current incomes do not tend to have the lowest expenditures. There is evidence that this is primarily due to under-reporting of incomes rather than over-reporting of expenditure, and one finds the same pattern when mapping income against measures of deprivation.¹⁷ Therefore, using consumption-based measures (which are based in large part on expenditure data) may also have practical advantages when trying to measure poverty accurately.

Figure 1. Median expenditure by income (equivalised, pounds per week)



Source: Brewer and O’Dea (2012), based on Living Costs and Food Survey 2006/07-2009.

¹⁵ Brewer and O’Dea (2012).

¹⁶ Cribb, Joyce and Phillips (2012).

¹⁷ Brewer, Paull, O’Dea and Sibieta (2009).

Note that greater confidence about the reliability of the data towards the bottom of the expenditure distribution would also be likely to mean that the depth of poverty could be accounted for more robustly when using consumption-based measures. Moving beyond simple ‘headcount’ poverty measures – which, in principle, is appealing (see Section 1) – may therefore be more workable when using consumption-based poverty measures than income-based ones.

Finally, income-based measures (and consumption-based measures) do not capture the fact that families use publicly provided services. If the government spends an additional £1 billion on tax credits, this would increase people’s current incomes (and probably their consumption too). If it spent the same money on improving public services for people in poverty, this could also improve their living standards but would have no direct impact on their income.¹⁸ If income-based measures of poverty alone were driving policy, then the incentives being acted on by policy-makers would be skewed and sub-optimal. **A way to address this is to track measures of the quality and usage of public services and, as far as possible, how quality and usage differs across the population.**¹⁹

The discussion so far has focused on how one can obtain a more comprehensive picture of material living standards than is obtained using standard income measures. One might also want to track more than just factors indicative of *material* living standards (by which we mean economic resources). It seems possible that the government has non-material factors in mind. For example, Chapter 2 of the consultation refers to possible effects of worklessness on “parents’ confidence, relationships, health and wellbeing” which “can also have negative implications for their children” (although it is possible that the government just has in mind knock-on effects of this for current material living standards, or for children’s life chances²⁰ – it is never made explicit). There is a case for tracking non-material indicators. But this does move further away from the typical use of the term ‘poverty’, and such indicators should be distinguished from measures of material living standards.

Distinguishing between low living standards, its causes, and life chances

The eight specific suggested dimensions in the consultation document cover more ground than the discussion so far, which has focused on measuring current living standards. The government actually seems to have at least four types of things in mind: a) children’s current material living standards; b) other inputs to their wellbeing; c) the causes of low living standards; and d) the future life chances of children who grow up in

¹⁸ It could even reduce measured consumption if families stop buying services privately in response, and increase saving.

¹⁹ The Office for National Statistics already attempts to account for public services in its measure of ‘final income’. For example, see http://www.ons.gov.uk/ons/dcp171778_267839.pdf.

²⁰ There is evidence that poor mental health during childhood can affect economic success during adulthood. See Goodman, Joyce and Smith (2011).

poor families. It is entirely appropriate to be interested in each of these things. But not enough is done to distinguish between them.

For example, the summary of parental skill in the consultation document suggests that it is included at least primarily because it is a *cause* of poverty and/or low life chances: “...low skills increase the likelihood of parents being unemployed, low paid and so in poverty, and also negatively impact on children’s development. Low skills are a key determinant of the length of time for which a family will be in poverty.”²¹ The factual accuracy of these statements is not in doubt. But they suggest that low skills should be targeted as part of a strategy for reducing poverty or improving life chances (if there are cost-effective ways of doing so); they do not suggest that low skills should feature in a measure of poverty or life chances.²²

It is understandable that the government is keen to avoid a situation where its political incentives to devote resources to things like improving parental skills as a way of reducing poverty are too weak, because spending the money in other ways (such as increasing benefits) may have more immediate and detectable effects on the most high-profile poverty indicators. But this can be achieved without confusing very different concepts. It does not make sense to say that something constitutes poverty merely because it causes it.

We recommend that the government defines distinct sets of indicators that will be used to measure children’s current material living standards; to measure children’s wellbeing more widely; to assess whether the causes of child poverty are being addressed; and to assess the likely future life chances of children. The relative prominence given to each of these sets of indicators, and the resources devoted to tackling each, would be separate decisions (and would be partly political choices).

Some factors may belong to more than one category. For example, something might be relevant both for the current material living standards of children and for their life chances. That is fine, and indeed desirable – if a factor genuinely matters for more than one reason, then it should be counted more than once. Our approach would force the government to think carefully and systematically about what matters and why. We believe this would lead to better outcomes.

²¹ “Measuring Child Poverty” consultation, pp. 33.

²² Perhaps the government believes that parental skill is not only a causal determinant of material living standards, but is also inextricably linked to children’s quality of life for non-material reasons, i.e. independently of effects on economic resources. That would put parental skill in category b), as well as category c), of our typology. If so, that argument should be made clearly, and should be distinguished from the very different claim that parental skill tends to have consequences for material living standards (which is what ‘poverty’ is typically taken to refer to).

3. A ‘multidimensional’ measure

Here, we use the term ‘multidimensional’ in the same way as the consultation document – to mean not merely the principle of measuring several things, but the idea of aggregating them into a single index.

We have already argued that there are at least four different concepts that the government appears to be trying to capture. They are very different concepts, and it therefore cannot make sense to combine them. The resulting measure would be difficult or impossible to interpret. Therefore, if there are to be any multidimensional measures, there should be at least four – and only one should be called a measure of child poverty.

But the question remains: *within* the conceptually distinct categories that we have outlined, should different indicators be weighted and combined into one index? In principle, if one has measures of different dimensions of (e.g.) material living standards and wants to form a comprehensive view, combining indicators in some way does look appealing. But we think that in practice there are major obstacles to doing this sensibly. It also seems inconsistent with the government’s stated aims for the new measure, which include obtaining a broad consensus and being widely understood.

First, the appropriate weights given to different indicators would have to depend on subjective judgments about their relative importance. For example, where the aim is to measure the non-material aspects of a child’s quality of life, one would want weights that reflect the relative importance of different factors in determining this. But this is not known. The government could state its own view by suggesting a set of weights. Indeed, it may be beneficial for the government to be explicit about its preferences in this way. But this is not something around which a consensus should be expected, and forming a measure which is “widely accepted by the public and experts”²³ is one of the stated aims of the consultation.²⁴

Second, there is a problem of transparency and ease of interpretation. Again, a stated aim is for the measure to be “understood and accepted by the public”.²⁵ But, since movements in indices which are weighted combinations of different factors may be very sensitive to the weights chosen, understanding the measure implies knowing what all the weights are. It seems implausible that such knowledge would be prevalent among the public, or indeed the majority of journalists who would report on this. That is an undesirable property for a high-profile poverty measure to have. Even if familiar with

²³ “Measuring Child Poverty” consultation, pp. 46.

²⁴ There is no way around this. For example, it might look tempting to define an index which is an unweighted sum of different measures in an attempt to be ‘neutral’ and avoid debate about the appropriate weights. But there is no reason to expect this to yield a better measure of anything – it simply amounts to an implicit assumption that all components of the index are identically important. It would be easier to interpret in the limited sense that it would be easier to understand the arithmetic of how the measure is computed. But it would not make it any easier to interpret what is actually happening to children’s living standards, which is the ultimate aim. That depends upon the relative importance of different factors, which is unknown and contested.

²⁵ “Measuring Child Poverty” consultation, pp. 46.

the weights, movements in the measure may be very difficult to interpret. A useful thought experiment is: “what does a one percentage point change (or similar) in the measure mean?” In the case of multidimensional measures, the answer is likely to be convoluted. Any serious analysis of it would almost certainly begin by breaking it down into its constituent parts.

Third, we have already discussed how absolute and relative measures of poverty are conceptually distinct. There are good reasons to track both types of measure: something would be lost if either were disregarded. But it does not make sense to aggregate them within a single index, because the measures provide completely different kinds of information. Indeed, aggregation might mask important changes in the living standards of poor children. Suppose, for instance, that the incomes of poor children are falling, but that median income is falling even faster. This will act to reduce relative income poverty and increase absolute income poverty. The two measures considered together would tell us that the incomes of poor children have fallen, but by less than median income. A composite index of the two would obscure both of these pieces of information, and indeed it might show no change (the falls in relative poverty being offset by increases in absolute poverty). This would clearly be far less informative.

In summary, the temptation to summarise everything in a single number should be resisted wherever there is more than one concept that one is ultimately interested in.

Data

The construction of a multidimensional measure would also raise practical issues regarding the data to be used, as multiple indicators would need to be included in the same dataset. And even if one did not attempt to combine various indicators into a single number, observing multiple indicators in the same dataset can still be advantageous. It allows the overlaps between indicators to be analysed (without necessarily relying on judgements about how they should be aggregated, as a multidimensional measure would need to do), which can shed light on important questions. Here we discuss some important issues relating to how one might go about obtaining data containing multiple indicators.

One would require a nationally representative micro dataset with information on multiple dimensions. There is a tradeoff between getting very high quality information on each dimension and having a questionnaire concise enough that response rates are reasonable. The Family Resources Survey is generally considered to provide the best measure of household incomes in the UK, and this requires a survey that is almost entirely oriented around providing information on income.

If looking for a large sample with a large range of information, ‘Understanding Society’ may look like the obvious candidate dataset. This is a panel of about 40,000 households who were selected in 2009. Surveys are sought from household members every year afterwards.²⁶ There is no question that detailed longitudinal data of this kind can be

²⁶ See <https://www.understandingsociety.ac.uk>.

extremely valuable in answering important research questions about the nature of poverty, its causes and consequences, and how it is transmitted across generations.

However, longitudinal data is not ideal if the aim is simply to obtain a robust and consistent measure of poverty over time. Sample members will drop out of the survey in a non-random fashion, and new individuals will join the sample as a result of changes in household formation. One is therefore entirely dependent upon reweighting techniques to ensure that the sample remains nationally representative in each wave. These techniques can only attempt to balance the sample on the basis of characteristics observed in the data, and there is never any guarantee that this is sufficient to end up with a sample that is truly representative with respect to what is being measured (in this case, child poverty). And different techniques for reweighting will inevitably result in different measured poverty rates, both in terms of poverty levels in a given year and in terms of changes over time as the sample continues to change. This all seems unnecessarily confusing, and risky for the robustness of the measure. Repeated cross-sectional data would avoid these problems.

In addition, anything based on Understanding Society could only be constructed from 2009 onwards, meaning that there would be no historical context (not even pre-recession context) in which to place the new measure.

For these reasons it is extremely important that any measure based upon 'Understanding Society' is viewed *alongside* existing measures based on repeated cross-sections of data. The Family Resources Survey, in combination with the Family Expenditure Survey (now the Living Costs and Food Survey), provide consistent measures of household incomes all the way back to 1961. The latter surveys also provide a long historical series of information on expenditure and consumption, which we have argued would be valuable for poverty measurement. These surveys also already provide some of the other information that the government has said it is interested in – for example, information on worklessness, housing, and education levels. It would also be quite straightforward to map in area-based measures of things like public service use and school quality without further augmenting the questionnaire. Given the concerns highlighted above about using longitudinal data, these surveys may therefore provide the best basis for measuring multiple indicators in a single dataset. As discussed, whether such indicators are then aggregated into a single index is a separate question, and we recommend against it.

References

- Blundell, R. and Preston, I. (1996), 'Income, expenditure and the living standards of UK households', *Fiscal Studies*, 16(3), pp. 40–54.
- Blundell, R. and Preston, I. (1998), 'Consumption inequality and income uncertainty', *Quarterly Journal of Economics*, vol. 113, pp. 603–640.
- Brewer, M., Browne, J., Joyce, R. and Sibieta, L. (2010), 'Child poverty in the UK since 1998–99: lessons from the past decade', Institute for Fiscal Studies (IFS), Working Paper no. 10/23 (<http://www.ifs.org.uk/publications/5303>).
- Brewer, M. and O'Dea, C. (2012) 'Measuring living standards with income and consumption: evidence from the UK', IFS Working Paper 12/12 (<http://www.ifs.org.uk/wps/wp1212.pdf>).
- 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. (<http://eprints.ucl.ac.uk/18312/1/18312.pdf>)
- Clark, S., R. Hemming and D. Ulph, (1981) 'On Indices for the Measurement of Poverty,' *Economic Journal*, 91, pp. 515-526.
- Cribb, J., Joyce, R. and Phillips, D. (2012), '*Living standards, poverty and inequality in the UK: 2012*', IFS Report C124, London: Institute for Fiscal Studies (<http://www.ifs.org.uk/publications/6196>).
- Cutler, D and Katz, L. (1992), 'Rising Inequality? Changes in the distribution of income and consumption in the 1980's', *American Economic Review*, 82, pp546-551.
- Dickens, R. (2011), 'Child poverty in Britain: past lessons and future prospects', *National Institute Economic Review*, no. 218, pp. R7–19.
- Goodman, A., Joyce, R. and Smith, J. (2011), 'The long shadow cast by childhood physical and mental problems on adult life', *Proceedings of the National Academy of Sciences of the USA*, vol. 108, no. 15 (<http://www.pnas.org/content/early/2011/03/14/1016970108.full.pdf+html>).
- Poterba, J. (1989), 'Lifetime Incidence and the distributional burden of excise taxes', *American Economic Review*, 79, pp325-30.
- Sen, A. K. (1976), 'Poverty: An Ordinal Approach to Measurement,' *Econometrica*, 44, pp. 219-231.