

Local indicators of child poverty – developing a new technique for estimation

Note by Donald Hirsch and Laura Valadez, Centre for Research in Social Policy, Loughborough University, July 2014

Background

The Centre for Research in Social Policy makes annual estimates for the End Child Poverty Coalition of the number of children in poverty in each ward, local authority and parliamentary constituency in the UK. These estimates are not accurate counts of how many children are in poverty in each area. Rather, they use the best local data available to give an indication of where child poverty is particularly high, and therefore where there need to be the strongest efforts to tackle it.

Up to the publication of the latest figures in early 2013, which apply to 2012, the data has been assembled on the basis of:

- The HMRC figures estimating local child poverty based on administrative data. This combines a count of children on out of work benefits and children in families on tax credits whose reported family incomes are below 60 per cent of the median. Child Benefit data are used to count the total number of children in each area.
- An estimated update to the above data, which is about 2½ years out of date at the time of publication. The update uses more recent regional Labour Market Survey data on trends in the number of children in out of work households to adjust figures for out-of-work child poverty.

Two drawbacks of the current method

These local data were originally designed as a snapshot estimate allowing comparisons of areas with higher and lower overall child poverty levels. However, in updating them over time, it is important that the method shows change that in some way reflects reality. The present method is highly problematic for that purpose, for two main reasons.

- (i) The local data produced by HMRC do not change in the same way as the official poverty rate**

First, and most crucially, the HMRC estimate is unduly skewed towards out-of-work poverty, which it overcounts (by assuming that everyone on out of work benefits is in poverty), compared to in-work poverty, which it undercounts. As a snapshot, these inaccuracies have roughly balanced out overall, originally producing an overall national figure not very different from the official HBAI estimate of the number of children living below 60% median income before housing costs. However, over time,


the measure has become increasingly unsatisfactory, to the point where it could soon be suggesting that trends are going in the opposite direction to reality.

This is because the HMRC calculation barely picks up trends in in-work poverty, which it counts as only 21% of all child poverty, when in reality it is now 63% in the HBAI figures. The relative share of in-work poverty has grown from under a half to nearly two thirds in the past decade. Moreover, in an economic recovery typified by falling rates of worklessness but also falling levels (in real terms) of in-work benefits, a measure driven by the numbers on out of work benefits could well show reductions in child poverty when the rate shown by HBAI is rising.

Figure 1 looks at how the HMRC figure has performed across time in estimating change in overall poverty rates. It has not proven particularly useful regard in this regard. Between 2007 and 2011 the official child poverty rate fell six percentage points but the HMRC data picked up only a quarter of this change. As a consequence, a measure that in national terms was quite close to the official before housing cost measure up to around 2008 is now significantly above it. This was influenced in particular by:

- The constant assumption that all families on benefits were in poverty, whereas the official figures showed a significant decline in risk due to improvements in benefits over this period
- The lack of impact on the HMRC figures of any trend in in-work poverty rates since they play a very small part in this estimate

**Fig 1 - Risk of poverty, different measures, UK
2006-11**




It is certainly a risk, and even likely, that these figures may start to move in different directions, as the relationship between in and out of work poverty continues to change.

As well as wanting a method that measures change over time in some meaningful way, it is naturally also important that it provides a reasonable estimate of differences across geographical areas. While survey evidence does not allow us to check the figures against the HBAI measure at a local level, because the numbers in the survey in any one area are too small, this is at least possible at the regional level. Figure 2 compares the HMRC and HBAI results across regions, using a three-year period, over which regional HBAI figures are reported in order to enhance sample size.

Figure 2 shows that while in general terms regions with higher actual poverty rates have higher reported poverty rates in the HMRC figures, the relationship is not consistent. The West Midlands, for example, has considerably higher poverty in HBAI than the North East, but lower on the HMRC measure. The gaps between the HMRC and HBAI figures both before and after housing costs vary considerably. For example, the London rate BHC is 11 percentage points below the HMRC figure, whereas for Yorkshire and Humber it is 2 points higher.

Fig 2 Risk of poverty, different measures, 2009-11 average by region



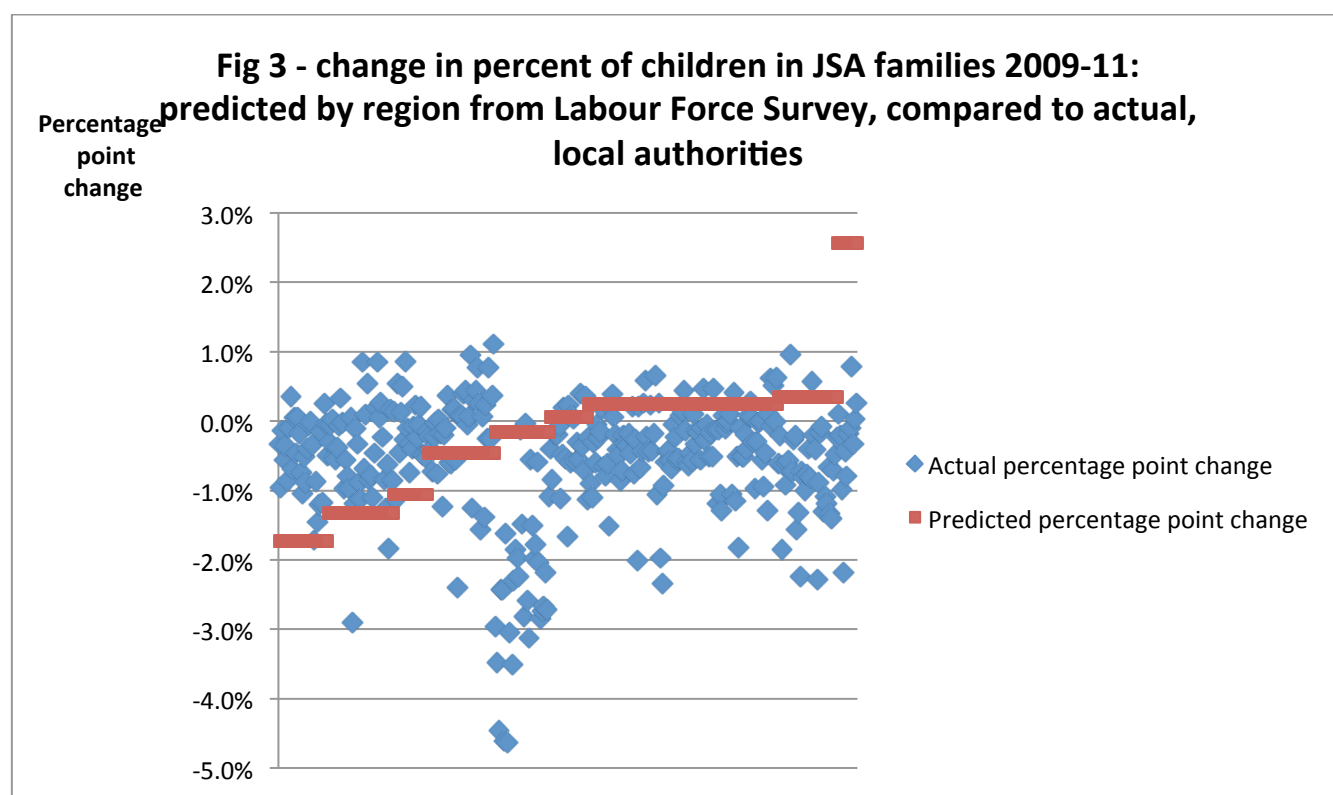
(ii) The updating method used up to now is starting to become misleading

The second, and related problem concerns the method for updating. ECP puts a strong emphasis on getting estimates of data that are as timely as possible, and Labour Force Survey data produce the most promising method of doing this. However, since all that these data can tell us is how many families in each region are in work, rather than anything about the changing risk of families with a given work status being in poverty, they are very crude. In the past two reported years, they have shown sharp drops in children in out-of-work households in London, without being able to say whether declining real wages or rising housing costs are in reality causing more families in this region to struggle, offsetting favourable employment trends.

A further risk is that in attributing very different trends to different localities according to what employment surveys show about what is happening at a regional level, the adjustment has a high level of inaccuracy about what is going on at a subregional level. For example, the estimates infer that in a region with falling worklessness rates, child poverty is falling in every local authority and ward, whereas in a region with rising worklessness rates the reverse is true. There are bound to be many exceptions to this within regions.

The performance of this updating procedure can be considered by looking at 2011 data on the one hand in terms of what would have been predicted from the 2009

HMRC figures and 2011 LFS figures used to update them, and on the other in terms of what the 2011 HMRC figures eventually did show, broken down by local authority. On average across local authorities, the error in the prediction compared to the outturn was 0.79 percentage points – ie it over or under-estimated child poverty by an average of about three quarters of one per cent. While this is not a large error, it is in fact very similar to the error that would have occurred using the out of date figures – that is, the estimate of change did not produce results closer to the truth than assuming no change. Figure 3 illustrates this inaccuracy. It shows both the large variation of local authorities within regions (each represented by a flat brown line) and the fact that the LFS evidence, based on household data, did not succeed in identifying in which parts of the country particular changes in claimant rates (which are at the family level) would occur.



Adjustment 1: getting a more reliable estimate of local child poverty based on HMRC data

In relation to the first problem outlined above, our analysis tested a variety of ways to reduce the inaccuracies in the way that the local poverty estimates would represent changes in child poverty over time and across regions, as illustrated in Figures 1 and 2 above.

The underlying principles for adjusting the figures are

- (i) that they should be calibrated to produce both an AHC and BHC measure that produce close to the “correct” result at a national level in each year, both for in and out of work poverty, and
- (ii) that this calibration should also produce as close as possible an estimate of regional rates as those produced by HBAI.

The simple calibration method is to look at the national level at how much the HMRC figures miscalculate in and out of work poverty rates, and to adjust data at all geographical levels by this amount. For example, in 2011, a child in an out of work family had a 67 per cent chance of being in poverty AHC, but the HMRC figures assumed the figure was 100% of all people on out of work benefits, so a simple adjustment multiplies each local out-of-work rate by 0.67, assuming that 67% times the number on out of work benefits is a better estimate everywhere of out of work poverty than 100% of this group.. Correspondingly, since the AHC in-work poverty rate is 3.9 times the rate shown by the HMRC figures, an adjustment by this ratio can be made, to produce the correct result overall.

Using such a ratio adjustment recalculated each year ensures that the sum of the local child poverty figures will be close to the numbers shown in HBAI each year, separately in calculations for AHC and BHC, and thus that the differences shown in Figure 1 above will disappear. The two HMRC-based rates will closely track those produced by HBAI. However, it is not inevitable that the regional differences shown in Figure 2 will disappear entirely, since for example the relative risk of in-work and out-of-work poverty may vary by region.

To test the extent to which a measure based on ratios produces a more accurate set of relative poverty rates by regions than those illustrated in Figure 2 above, we can compare these rates produced by this adjustment applied to HMRC figures in 2009-11 with the actual regional rates as shown in HBAI. In this case the error rate is much smaller than without the adjustment, with an average regional inaccuracy of 1.6 percentage points in the child poverty rate. This suggests that the presumption of a fairly consistent ratio between the HMRC and actual poverty risks as formulated above is a reasonable one. One reason for thinking this is that it eliminates large differences between error rates for different places, notably London and other parts of the country, where conditions are far from identical.

However, there remains a risk that for some individual locations, especially those with more extreme characteristics, applying a single ratio will produce substantial distortions. This is especially the case for the in-work adjustment, where multiplying the HMRC rate by a factor of nearly four could potentially produce serious inaccuracies, most particularly where in-work poverty is already reported as high, and therefore an exceptionally high rate with a high margin of error could be produced. This is in fact the case, for example, with a ward such as Bethnal Green South in Tower Hamlets, where the 50 per cent child poverty rate estimated by HMRC rises to 73 per cent AHC on adjustment. It is quite possible that the 73 per

cent figure is correct, bearing in mind that a large proportion of the workforce in this area have high housing costs and low household earnings, but the uncertainty about the validity of such an extreme figure makes it risky to produce it as an estimate, to which considerable attention would be drawn. The extent to which this simple ratio produces an accurate prediction in this case is subject to a wide margin of error.

To deal with this problem, we tested two alternative techniques for producing ratios for estimating in-work poverty that dull the effect of wide variations from the initial figure and therefore reduce the risk of extreme outliers subject to large uncertainty. The first method is to assume that the relationship between HMRC and actual in-work child poverty rates is exponential rather than a simple multiple, and hence to assume that the actual rate is proportional to the log of the HMRC rate rather than to the rate itself. The other is to assume that the relationship is governed by odds ratios rather than simple percentages: that there is a consistent factor by which the HMRC figures under- or over-estimate the *odds* of a child living in poverty compared to not living in poverty. Both of these relationships have the advantage of making it impossible to produce a total poverty risk rate above 100%, which would be possible using the simple ratio method described above. Moreover, it is possible to test which method produces the most accurate results at a regional level, and hence which seems best to describe how the relationship between predicted and actual results work out in practice.

Table 1 shows the results of this exercise. It shows that the average regional error rate after a ratio based on log values is slightly lower than one that uses crude ratios. However, an adjustment based on relative odds ratios reduces the error rate further, by about a third. With this method, regional child poverty rates get to within less than one percentage point of the real ones, and in six of the nine English regions it is below half a per cent. This suggests that odds ratios are a good means of applying a standard adjustment, that is not greatly confounded by area differences. Note also that in the case of the Bethnal Green South example mentioned above, the adjusted AHC poverty rate using this basis drops to 57 per cent. This is seven percentage points higher than reported by HMRC, similar to the national gap observed between the crude HMRC poverty rate and the actual AHC figure. In other cases where the simple ratio approach would produce much higher local data than reported by HMRC, an odds ratio similarly produces an adjustment more in line with national trends.

Table 1: Accuracy of estimated regional child poverty rates using different methods of adjusting HMRC figures, 2009-11 average.

Region	HBAI <u>estimate</u>	Simple ratio		Ratio of log values*		Relative odds ratios*	
		<i>Rate</i>	<i>Difference with HBAI</i>	<i>Rate</i>	<i>Difference with HBAI</i>	<i>Rate</i>	<i>Difference with HBAI</i>
North East	28.0%	32.2%	-4.2%	30.0%	-2.0%	31.4%	-3.4%
North West	31.0%	32.1%	-1.1%	29.7%	1.3%	31.1%	-0.1%
Yorkshire and the Humber	31.0%	32.3%	-1.3%	29.5%	1.5%	31.2%	-0.2%
East Midlands	25.0%	27.1%	-2.1%	26.4%	-1.4%	26.8%	-1.8%
West Midlands	32.0%	33.8%	-1.8%	30.5%	1.5%	32.4%	-0.4%
East of England	24.0%	23.9%	0.1%	24.4%	-0.4%	24.1%	-0.1%
London	35.0%	36.5%	-1.5%	32.8%	2.2%	34.9%	0.1%
South East	21.0%	20.8%	0.2%	22.5%	-1.5%	21.3%	-0.3%
South West	25.0%	22.5%	2.5%	23.6%	1.4%	22.9%	2.1%
Average size of difference			1.6%		1.4%		0.9%

*These different adjustment methods apply only to in-work poverty rates. Simple ratio used in all cases for out-of-work adjustment

On the basis of this evidence, the “relative odds ratios” method will be used to adjust HMRC figures for in-work poverty, while the simple ratio will be used for out-of-work poverty. (Note that for out-of-work poverty neither an exponential nor an odds-ratio adjustment would be possible, since the assumed risk in the HMRC calculation is 100%. There is no exponential factor that transforms 100% to 67% or any other percentage, since 100% to any power remains 100%, ie 1. Similarly 100% cannot be expressed as an odds ratio since with a 100% poverty risk, the odds of poverty compared to not poverty is infinite!)

Adjustment 2: finding a useful way of updating data collected over two years previously

As discussed above and illustrated in Figure 3, the method used so far to get more up to date child poverty estimates than those published by HMRC has outlived its usefulness. This method makes only small adjustments, since poverty rates change only slowly, but the size of these adjustments do not seem to produce more accurate estimates of current levels than the more out-of-date figures.

The Labour Force Survey remains the most promising source of recent information to get timely information relevant to people’s incomes. Data on how many children live in working compared to non-working households, used in the old calculations, continue to appear to have the greatest pertinence to child poverty rates among those appearing in the survey. The risk of poverty is so much higher for out of work families than those in work that changes in worklessness seems bound to affect poverty rates. Other data in the LFS including earnings figures could also potentially have a bearing on poverty rates, but it would require complex modelling to discern

whether there is a measurable relationship of this kind, or whether the effect of earnings changes is mixed with so many other simultaneous changes that an independent effect is impossible to measure.

Therefore, the analysis here has looked at what alternative measures using the LFS data on household working status might be used. This is now published at both regional and local authority level, raising the possibility that an updated set of estimates could be more sensitive to local difference than the regional figures used up to now. However, given that the evidence is from a survey, the risk is that by using smaller samples, more local estimates may actually be *less* accurate.

Table 2 looks at how many children would have been estimated as being in families on out of work benefits using various forms of updating the HMRC's 2009 data to 2011, and compares these results to the actual HMRC figures for 2011.

Table 2

Accuracy of LFS-based forecasts that could have been made for 2011

Estimates for percentage of children on out-of-work benefits predicted by 2009 HMRC and 2011 LFS data, compared to actual change according to HMRC data for 2011 (published in 2013)

Assumption	Estimated overall change (%age points)	Difference with actual change (-0.76% age points)
No change	0%	0.76
Change in each local authority at the same rate as national LFS change (in %age points)	-0.11%	0.65%
Change in each local authority at the same rate as national LFS change (proportionately)	-0.17%	0.60%
Change in each local authority at the same rate as regional LFS change (%age points)	0.03%	0.79%
Change in each local authority at same rate as regional LFS change (proportionately)	-0.04%	0.72%
Change in each local authority at same rate as local authority LFS change (%age points)	3.98%	4.74%
Change in each local authority at same rate as local authority LFS change (proportionately)	4.82%	5.59%

Table 2 shows what may at first seem a rather counter-intuitive result. This is that, when predicting the actual change in out-of-work benefit claimants rates in local

authorities over the least accurate measures is to use Labour Force Survey data on workless households at local authority level, which is less accurate even than assuming no change at all, while the most accurate measure is to apply the rate of change in workless families at *national* level to all local areas. This can be explained by the distorting effects of sample error when using Labour Force Survey data at the sub-national level.

On this basis, in estimating 2013 data from the 2011 HMRC figures, we assume that the percentage of children in poverty in out of work families has everywhere changed in direct proportion to the national percentage of children in non-working households shown in the Labour Force Survey. As Table 2 shows, this produces somewhat more accurate figures at local authority level than sticking with the earlier data, even though the results should just be presented as estimates.

An equivalent adjustment can also be made in the in-work poverty assumptions. This is that the 2011 estimates of in-work poverty are multiplied by the proportionate change in the percentage of children in working households between 2011 and 2013. This only makes a very small difference, but it is appropriate to reflect the fact that, for example, the proportion of children in working households rose by about 3% from 2011 to 2013, so if risks remained the same, working poverty can be expected to have risen (as well as non-working poverty falling).

Conclusion

This paper shows that it is possible to improve local estimates of child poverty (i) to make them more compatible than up to now with the way in which poverty is measured through HBAI and (ii) to ensure that an approximation of change since the period covered by the latest available administrative data actually improve the estimate rather than showing figures no more accurate than the more out of date ones. On this basis, in summary, the following steps will be taken in 2014 to estimate 2013 data. In future years, similar estimates will be possible, using data available in March of the year after the one that they cover.

Step 1: Compile HMRC data for in and out of work local child poverty rates, published for two years preceding the year under review

Step 2: Adjust both of these rates to create local estimates of actual child poverty for the same year, both AHC and BHC. For out of work rates, make these estimates by multiplying the numbers on IS/JSA by the overall risk of a child in a non-working family being in poverty. For in-work rates, make the estimates by adjusting odds of a child in a working family being in poverty compared to not in poverty, by a factor that produces the “correct” (HBAI) national poverty figure AHC, and repeat for BHC.

Step 3: Adjust each of the out of work poverty rates for most recent year, by multiplying the out-of-work rates by the proportionate change in the national risk of a

child being in a non-working household, as shown by the Labour Force Survey over the relevant two year period. Repeat for working families on the same basis.

Step 4: Combine the resulting in and out of work poverty rates for all wards, local authorities and parliamentary constituencies to produce the final estimates of child poverty.

These figures should always be presented as estimates. The procedures used to create these estimates are far from perfect. However, this paper has argued that their merit is that they are likely to be closer to the true level of child poverty (defined as below 60% of median income) than the HMRC data unadjusted, and are thus the best estimates available.