

If I'd known you were coming I'd have specified a plausible counterfactual: we need a better explanation of failure to reach the 2010 child poverty target

1

Why was the 2010 target to halve child poverty missed? There are many different factors that could be alleged to form part of the explanation. In fact there are so many that it's obvious the question needs to be tightened up before it can really be addressed.

The best way to do this is to ask what plausible developments might have led to the target being met and explore the reasons these didn't occur. In other words, any serious explanation of why the target wasn't met involves a counterfactual under which it would have been met. Without specifying a counterfactual, the field is open for anyone to throw in whatever in the vast array of explanatory factors most appeals to them. (If people think there is no such counterfactual, as Frank Field does, they should say so explicitly, as Frank does- erroneously in his case.)

The counterfactual has to be realistic to be of any use. Saying that the target would have been met had the government or anyone else taken some action which lies outside the bounds of plausibility tells us nothing. Part of this plausibility condition is the need for factors to be relevant to the time period we are looking at. Saying that the target could have been met if some possible long-term development had been telescoped into the few years we are looking at also tells us next to nothing. To me, these simple criteria diminish the relevance of many of the factors that people frequently appeal to in explaining perceived lack of progress on poverty reduction in the UK.

Social conservatives can justifiably argue that if UK family structure were different then the target would be easier to meet: relatively high rates of lone parenthood are indeed important in understanding why child poverty is so high in the UK. But any changes in family structure which would be likely to have a significant impact on the achievability of the target would of necessity have a quite different timeframe to the relatively short period (1998/9-2010/11) we are talking about.

On the left, it is common for people to say that meeting the target will always be difficult without the levels of taxation, redistribution and public service provision associated with the Nordic social democracies. Again these considerations are relevant if we want to understand the UK's comparative position, but they have nothing to contribute to explaining why *that* target, in *this* country, over *this* time period, was missed. Is there a plausible counterfactual under which the share of public expenditure in UK GDP moved from its long-term average of 40-41% to the

Swedish average of 54% - as a permanent shift, not a response to recession- between 1998/9 and 2010/11? I think not.

So explaining the failure to hit the 2010 target involves limiting the range of relevant explanatory variables, which is basically a matter of distinguishing between long-term background conditions (such as family structure or the spending share of GDP) and the factors which explain change (or its absence) over the time period in question.

The logical form of this sort of counterfactual explanation is worth noting: If Y , p where p stands for 'the target would have been met' and Y is a set of (one or more) plausible counterfactual conditions relevant to the time period under consideration. Less formally, 'If I'd known you were coming...' (Y) '..I'd have baked a cake' (p). This may seem trivial, because that is what it is, but spelling out the logical form is a good way of unpicking exactly what is being asserted.

It's also worth distinguishing between descriptions and explanations- not always a hard and fast distinction, but nonetheless pretty fundamental. To say for example that child poverty was at $x\%$ because the poverty risks for groups a , b and c were whatever they were and their shares of the population were whatever *they* were is generally speaking a description: to say that a change in the reward to work for parents in group a led to an increase in employment with a consequent effect on the poverty risks for children is an explanation, or part of one.

2

Under the 2010 Child Poverty Act the government is obliged to report to parliament about progress on the child poverty target. That report <http://media.education.gov.uk/assets/files/pdf/c/child%20poverty%20in%20the%20uk%20the%20report%20on%20the%202010%20target.pdf> was published on 14 June 2012, the same day as the annual *Households Below Average Income* publication which contained the figures for child poverty for 2010/11. It was prepared by civil servants at DWP and DfE.

Given that this report came out the same day as *HBAI*, it would be unreasonable to expect the analysts to have produced anything more than a broad –brush account. That said, the status of this report means that its account of what has happened is destined to wide exposure, perhaps more than the departments' analysts would prefer.

So how does the report account for failure to meet the 2010 target? The key passage is this: 'Despite some progress, not enough parents moved into work, and progressed in work. Work did not pay as well as it should and the proportion of

poor children who came from working households increased. Not all poor families received the financial support they were entitled to as the system was complicated and unclear.'

So this boils down to saying IF Y, p where Y is the set

- (a) if more parents had moved into or progressed at work
- (b) if work had paid more
- (c) if take-up of benefits had been higher

and p is 'the target would have been met'.

In what follows I'm going to concentrate on (a).

3

The approach taken by the analysts is admirably simple: they point out that reducing child poverty involves reducing poverty for both in-work and workless families, and posit that 'If the size of both of these groups had been halved by 2010/11, the total number of children in poverty would have halved and the target would have been met. While this is not the only combination of progress that would have been successful, it provides a useful benchmark for the discussion below.' (para 23, p.12)

Changes in the numbers of children in poverty each group can be split into (a) changes in the number of children (poor or not) in each group and (b) changes in the risk of poverty within each group. When we add (c) changes in the relative population shares of children in in-work and workless families the analysis is complete. (In fact, this is the composition of changes in the *rate* of child poverty rather than the number of children in poverty but that doesn't raise any major issues over this period that I can see.)

From this analysis the departments draw the following striking conclusion:

'Analysis of the change between 1998/99 and 2010/11 shows that **only 13 per cent of the overall reduction in poverty was due to a reduction in the proportion of families who were workless.**'


Bear in mind that this is all purely descriptive: but there is the danger of logical leaps to explanation among readers predisposed to particular types of explanation. The inference which has been drawn from this by government and some commentators is that there is a simple *explanation* for failure to meet the target; i.e. that the previous government was successful in raising the incomes of workless families and to some extent in-work families but failed to make proportional inroads on the number or proportion of children in workless households.

That is the conclusion drawn by Iain Duncan Smith in his introduction to the report: ‘the last government’s approach to poverty was fundamentally flawed - today’s figures show that they did not do enough to make work pay, focusing instead on treating the symptoms of poverty through simple income transfers.’

I don’t think the departments’ report justifies this statement. Nonetheless the aspects which are highlighted in the descriptive analysis are ones which can easily be read as supporting the Secretary of State’s explanation. So it is important to note that as an general account of trends in child poverty under Labour, this is a non-starter, as it completely ignores the impact of recession.

In order to get a sense of how the three factors which the departments have specified played out under Labour, it is necessary to look at the pre- and post-recessionary contexts. First, I’ve attempted to reproduce the departments’ analysis for the whole period from the published *HBAI* data. The chart corresponds to the one on page 14 of the report and shows how the fall in child poverty between 1998/9 and 2010/11 is made up of changes in within-group risks and changes in the share of children in workless families. (So positive figures show a *negative* contribution to child poverty.) Figures are for poverty on a Before Housing Costs basis.

Decomposition of change in child poverty rate (BHC) 1998/9-2010/11




My values are slightly different to those in the report . The fall in poverty risk for children in workless families is 54.5% of the total, compared to 52% in the report, while the compositional effect of changes in shares of children in workless and in-work families is 15.7% compared to 13%. The values are close enough to indicate (to me at least) that we're doing the same kind of analysis of roughly the same data however.

So over the entire period the change in within-group risks accounts for some 84% of the fall (87% on the departments' figures) and shifts between the groups have a minor impact. Employment is playing little role in child poverty reduction, which seems consistent with the message from the Secretary of State.

Then I analysed the composition of change between 1998/9 and 2008/9 , which can reasonably be seen as the last pre-recessionary year for this purpose (as the bulk of the incomes data was collected before the major impacts of recession were registered).

Decomposition of change in child poverty rate (BHC) 1998/9-2008/9

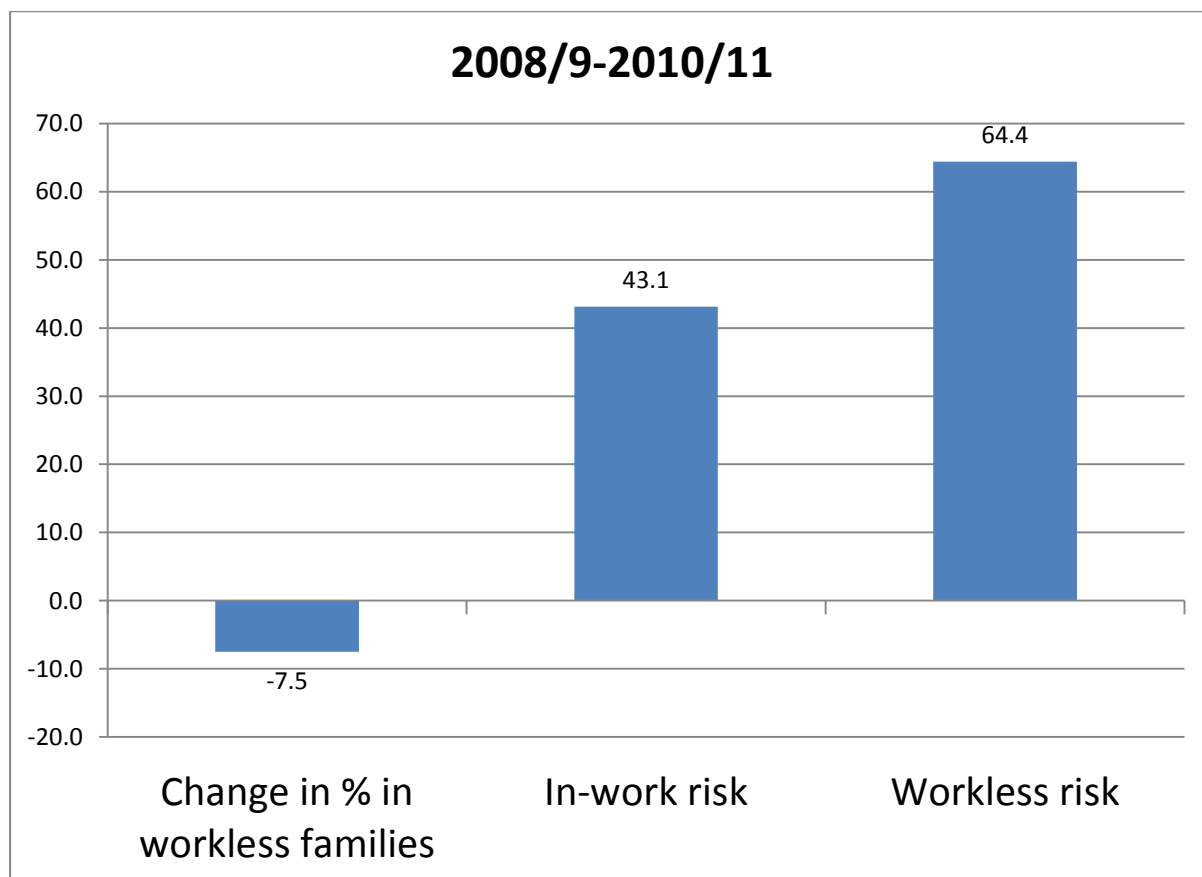


This gives a very different picture. The child poverty rate fell by four percentage points over this period. 43% of the fall is due to the composition effect, 38.2% to reduced risk of poverty for children in workless households and 18.8% to reduced risk for children in in-work households. Changes in risks still account for a majority of the fall, but it is a slender majority indeed (57%) compared to the previous chart (84%).

In other words over this period reductions in the number of children in workless households had a major impact on child poverty, larger than the impact of either of the within-group changes in risk. So Iain Duncan Smith's characterisation of poverty reduction under the last government as 'treating the symptoms of poverty through simple income transfers' is simply wrong. Bear in mind, this chart shows what happened over eleven of the twelve years in which Labour was responsible for child poverty.

Finally I analysed the change between 2008/9 and 2010 when there was a further fall in child poverty of 4 percentage points, on which as has been widely noted falls in median income had a significant impact due to the consequent lowering of the poverty threshold.

Decomposition of change in child poverty rate (BHC) 2008/9-2010/11



The main point here is that given these (one hopes) quite exceptional economic circumstances, it is very important to analyse the post-recession years separately to the preceding period if we want to understand what has happened over the entire period since the child poverty target was introduced. All of the reduction over this period is due to reduced within-group risk with the composition effect being modestly negative, reflecting a relatively small rise in the proportion of children in workless households in the FRS data. None of this, with the exception of the last point, should be surprising. We would hardly have expected the composition effect to be positive during a recession- that it was not more strongly negative is rather striking. But what is really surprising is the scale of the poverty reduction effect, which comes largely from reductions in median income, particularly between 2009/10 and 2010/11, unlike in the earlier period when redistributive measure offset the poverty-increasing effect of rises in median income.

The charts above are based on the published aggregate data for in-work and workless families in *HBAI*. The IFS has published a more detailed analysis of change over the period up to 2010/11 <http://www.ifs.org.uk/comms/comm124.pdf> looking at specific combinations of family type and employment status, which reaches similar conclusions to the departments about the effects of changes in composition and risk. (The departments' analysis may in fact be based on this more detailed data, but this is not what they've published.) Reproducing the IFS analysis for 1998/9-2008/9 appears to show some puzzling discrepancies in the published *HBAI* data (between aggregate figures for in-work and workless families on the one hand and detailed categories on the other). Using the more detailed categories, the share of compositional change in child poverty reduction up to 2008/9 falls from 43% to 27%, which is still a lot more than the 13% reported for the entire period to 2010/11. Those are my calculations: in fact an earlier IFS publication <http://www.ifs.org.uk/wps/wp1023.pdf> which analysed change from 1998/9-2008/9 implies a slightly higher share of 30% for compositional change.

So it looks as if the current downplaying of employment as a factor in child poverty reduction under the previous government is over-influenced by the impacts of recession.

These are descriptive statistics: any explanation of trends in child poverty needs to draw on a good description of what has happened, but an explanation has to add

some sort of theory as well. The government is leaning heavily on an explanation which turns on insufficient falls in worklessness among families with children over the period 1998/0- 2010/11.

Clearly employment played a major role in reducing child poverty between 1998/9 and 2008/9, contrary to what has been claimed by Iain Duncan Smith. But the fall of four percentage points pre-recession is less than we would have hoped for, so it seems reasonable to argue that had family-level worklessness fallen more, we would expect child poverty to be lower.

But the question is not whether poverty would be lower if worklessness was lower: it is whether a further fall in family-level worklessness sufficient to make a significant impact on the child poverty measure over this period meets our plausibility criterion.

I am not convinced it does. The main source of reductions in family worklessness since 1998/9 (and indeed a bit earlier) has been lone parent families. Lone parent employment was only 44% in the mid-1990's, having fallen from about 60% at the end of the 1970's. It is now 57%. That is a rise of 13 percentage points over a decade and a half. I don't know of any country which has seen a more rapid rise in lone parent employment: not even the United States in the wake of the Clinton-era welfare reforms. The UK, Canada, France and the US, under very different approaches to reform, have seen pretty similar rises. If there is a country which has done better than this, I'd be delighted to revise my expectations, but until then, I think we should avoid assuming that there was much more to be achieved through this route than actually was achieved- which was a lot.

So if read as an explanation rather than a description this part of the departments' account for why the 2010 target was missed seems to me to be unconvincing. The descriptive statistics suggest a potentially misleading diagnosis while the counterfactual of a much bigger fall in family worklessness than actually took place seems implausible. (Interestingly, the departments stress the limitations of plausible employment changes in the report, albeit on a rather different basis- see para 17.)

Are there other, plausible counterfactuals that might give a better account of how the 2010 target might have been achieved? The departments mention take-up problems due to the complexity of the system, although they don't quantify the effect of this over the entire period. The potential for increased employment for partners in working couple families is important: we could ask whether a different approach to tax credit design, or a more universalist approach to childcare, might have made a difference. Similarly we could ask whether a different approach to incentives for both employees and employers might have led to less reliance on short working hours. And of course we could ask whether there might have been better (not

necessarily larger) redistributory measures. No doubt there are many other possibilities.

But the government's favoured explanation – too much redistribution and not enough work- seems to me to be no explanation at all: it looks more like a pre-conceived partisan accusation against the previous government, opportunistically exploiting descriptive statistics which are far too strongly influenced by extraordinary economic circumstances to be a useful guide to developments over the longer term.