8. Measuring the distributional impact of public service cuts

Cormac O'Dea and Ian Preston (IFS)

Summary

- The fiscal tightening currently under way will rely on cuts to spending on public services to a greater extent than on cuts to social security spending or increases in taxation.
- Distributional analyses of changes to spending on public services are not common. This is because, unlike with changes to taxation and cash benefits (which directly affect the income of taxpayers and recipients), there is no readily calculable *quantitative* measure for valuing the benefit the public get from services that are provided in kind rather than in cash (such as hospitals, schools, the army and government administration).
- In spite of the unavoidable difficulties associated with carrying out this type of analysis, the imminent deep cuts in public service spending have provoked a good deal of interest in evaluating the distributional impact that they will have.
- The studies that have aimed to evaluate these distributional impacts (including that published by the Treasury alongside the Spending Review) typically assume that the value of a public service is equal to the cost of providing it. But the problems implicit in this approach (problems that are typically noted by those carrying out the studies) mean that the results should be interpreted extremely cautiously.
- It is certainly to be welcomed that the Treasury has considered the distributional impact of changes to spending on public services. However, we make a number of recommendations for any future analyses. In particular, given the absence of an established methodology for carrying out this type of analysis, it is crucial that details are published of how distributional impacts are estimated. Without such details, the robustness of the analysis cannot be assessed. It is also important that the Treasury take a more consistent approach to determining which measures are included in a distributional analysis.

8.1 Introduction

The fiscal tightening that is currently under way will, by 2014–15, be composed as follows: 27.2% net increases in taxation, 15.3% cuts to spending on cash benefits and 57.4% cuts to spending on public services. These measures, as they reduce government borrowing, will inevitably involve a reduction in the resources available to households. Given their scale,[1] an understanding of how the impact of each is distributed across the population is important.

[1] For a discussion, see Chapter 6.

Distributional analyses of changes to taxation and the payment of cash benefits are frequently produced in the aftermath of Budgets[2] and follow an established methodology.[3] For reasons that we discuss in this chapter, distributional analyses of changes to spending on public services (health, education, defence, government administration etc.) are much less common. However, the prospect of large cuts to departmental budgets has sparked a good deal of interest in this sort of analysis recently – with studies claiming to show the distributional impact of these cuts being produced both inside and outside government. The issue is not, however, new: the distributional impact of changes to public service spending was just as much an omission in the evaluation of fiscal plans during earlier times of expanded public expenditure.

This chapter proceeds as follows. Section 8.2 outlines the considerable difficulties in evaluating precisely the distributional impact of public services that are provided in kind rather than in cash. Section 8.3 uses evidence from studies that examine how usage of public services varies over the income distribution to discuss the distributional impact of those services. Informed by these studies and the analysis published by the Treasury alongside the Spending Review, we discuss what the distributional impact of those cuts in spending that have been announced over the past six months might be. Section 8.4 concludes.[4]

8.2 Can we measure the distributional impact of spending on public services?

What is the monetary value to a household of spending on a local school? Or a local hospital? Or the armed forces? Or the Treasury? These are the type of questions that must be answered *in quantitative terms* if one wants to evaluate *quantitatively* and compare the distributional impact of spending on schools, hospitals, the armed forces or the Treasury.

The impact of public spending decisions is substantially harder to document precisely than the impact of decisions that relate to taxation or the payment of cash benefits. To see why, consider the following two questions:

1. Who loses most, and how much do they lose, from a cut to Child Benefit?
2. Who loses most, and how much do they lose, from a cut to spending on the NHS?

The part of the question that asks 'how much do they lose?' calls for the answer to be expressed in quantitative terms. The first of these questions is much easier to answer precisely. In assessing the effect of a cut in Child Benefit on households, the obvious answer is a quantitative one – the amount (in pounds per week, for example) that the household has lost. In assessing the effect of a cut in spending on the NHS, on the other

[2] For a recent example, see J. Browne and P. Levell, *The Distributional Effect of Tax and Benefit Reforms to be Introduced between June 2010 and April 2014: A Revised Assessment*, IFS Briefing Note 108, 2010 (http://www.ifs.org.uk/publications/5246). See also Chapter 12 in this Green Budget, which gives the distributional impact of changes that will take effect in April 2011.

[3] For an overview of this methodology, see R. Hancock and H. Sutherland (eds), *Microsimulation Models for Public Policy Analysis: New Frontiers*, STICERD Occasional Paper 17, London School of Economics, 1992, and G. Redmond, H. Sutherland and M. Wilson, *The Arithmetic of Tax and Social Security Reform: A User's Guide to Microsimulation Methods and Analysis*, University of Cambridge, Department of Applied Economics, Occasional Paper 64, Cambridge University Press, 1998.

[4] A more detailed treatment of the same issues is considered in C. O'Dea and I. Preston, *The Distributional Impact of Public Spending in the UK*, 2020 Public Services Trust, London, 2010 (http://www.ifs.org.uk/publications/5234).

hand, it may be possible to identify who has been affected but there is no uncontroversial or readily calculable quantitative measure of the loss each household experiences. Cuts to spending on public services typically have no direct cash impact on households, but the welfare of households will, of course, be affected. The size of the impact will depend on, among other things, the amount of the service a household uses and how much they value the service. If a quantitative assessment of the distributional impact of that spending cut is required in a way that allows comparison with other budgetary changes, then that reduction in welfare must be expressed in cash terms.

This raises the question of why one would want to express quantitatively something (such as the value of public services) that has no simple quantitative interpretation.[5] There are a number of reasons. First, it is necessary if one wants to compare the distributional impact of cuts to spending on public services with cuts to benefit payments or increases in taxation. Second, it allows a comparison of the distributional impact of changes to various types of public spending (e.g. closing libraries compared with closing hospitals). Putting these together, it would allow a distributional analysis of the *entire* fiscal consolidation, i.e. an assessment of the progressivity or regressivity of all measures introduced (including both those with a cash effect and those with an in-kind effect).

In the rest of this section, we first discuss a relatively intellectually satisfying (but perhaps impractical) approach to the valuation of public services. This equates a user's value of a public service with what their willingness to pay for it would be were it not provided publicly. We then discuss a less-than-ideal (but more practical and therefore more common) approach to the valuation of public services. This equates a user's value with the cost of provision.

Value as 'willingness to pay'

A natural way to think about valuation, we suggest, is to think about the willingness to pay. To value the services provided by a library, for example, we could evaluate how much households would have been willing to pay for its services if it were not provided by the government. Or if we want to assess the impact of a cut in the library's funding, we could evaluate how much households would have been willing to pay to avoid that funding cut.

There are a variety of methods that could be informative about the willingness to pay for some public services. For example:

- The premium paid for a house with easy access to a public amenity (e.g. a public park or transport hub) relative to a similar house with no such access can be used to assess willingness to pay for that amenity.
- For some publicly provided goods, relatively close substitutes exist whose market price can be used to assess valuations. A leading example of this is health insurance, the price of which can be used to inform an exercise that aims to value public healthcare.

[5] The danger of thinking that 'any number is better than no number' is an obvious one. For a perspective that argues that, in the context of goods that are consumed collectively, no number might often be better than some number, see P. Diamond and J. Hausman, 'Contingent valuation: is some number better than no number?', *Journal of Economic Perspectives*, 1994, 8(4), 45–64.

- Some surveys contain questions that try directly to elicit valuation by asking respondents how much they would be willing to pay for some hypothetical expansion of government spending on some public service.

There are issues, which we do not discuss here, that make each of these strategies far from perfect. However, if the aim is to value the overall distributional impact of public spending, their greatest weakness is their limited applicability. None of these methods can easily be applied to place a value on every public service, nor can they be applied to assess the distributional impact of a package of spending cuts at the level of detail announced in the recent Spending Review. Spending Reviews typically set out the budgets available to departments rather than the programmes and services on which those budgets will ultimately be spent. However, it is those programmes and services that can (potentially) be valued using a willingness-to-pay measure, rather than the budgets themselves.

Value as the cost of provision

In light of the absence of any method to estimate willingness to pay for all public services, it is perhaps unsurprising that all the studies that have recently attempted to assess the value of spending on public services equate the value to the cost of provision. The approach typically taken is as follows. An estimate of which types of households use a particular public service is made. The cost is then divided equally between the users; the resulting cost per user is assumed equal to the value per user. This approach, while certainly feasible, is problematic. We now outline two problems with it.

Value and cost diverge

The fact that value and cost diverge can be illustrated by stark examples. At one extreme, there is likely to be some government spending that is considered to be worthless or even destructive by some households. For example, there has been much discussion recently about the possibility of cutting 'wasteful' expenditure.[6] Since waste is costly, defining value based on cost would lead to the conclusion that cutting waste in the provision of public services reduces their value to the end-user. In this case, the notion of value as willingness to pay seems more sensible as, presumably, households will have little or no willingness to pay for waste, and it would therefore be assigned a low value.

At the other extreme, some government spending could have a value far in excess of the cost of providing it. There is a range of types of spending that could be considered to fall into this category. For example, in spending money to correct market failures, governments can generate value far in excess of the cost. Natural monopolies are a case in point – in many cases, some essential service (e.g. piped water or the national grid) might either not be provided or be provided at much higher cost in the absence of the government either providing it directly or regulating the market so that private provision can occur. Once again, in this context, the notion of value as willingness to pay seems more sensible. In the absence of the provision of piped water, households would be willing to pay quite a lot for it – likely more than the average cost of providing it to them through public intervention.

[6] We can distinguish between two different types of waste. First, there is 'pure waste' (e.g. leaving the lights on in government offices at night), which benefits nobody. Second, there is paying people to do things that are not considered socially valuable. The latter type of expenditure is often labelled 'waste' and, while it does not create any value for users of public services, it does benefit the recipient of the payment. Its effect, therefore, is rather like that of a transfer payment. The discussion in the text here relates primarily to 'pure waste'.

The value of public services will not be the same for all users

The second concern with directly equating cost and valuation is that it implies that the value is the same to all those who use a public service of similar cost. In fact, there is good reason to believe that the valuation of certain public services will vary with household characteristics and, in particular, household income. This is a crucial issue when the question at hand relates to the extent to which government spending has a different impact on those at different points of the income distribution. If users with different incomes value the (similar) service they receive differently, then assigning everyone the same value will yield misleading results. This can be illustrated with two examples that highlight why intuitive and apparently innocuous assumptions can lead to results that seem questionable.

First, consider the case of military spending. Without any obvious way of differentiating 'usage' of the military, the approach under discussion would naturally assume that everyone benefits to the same extent. If everyone is assigned the same cash valuation, those at the bottom of the income distribution get more value from the military as a proportion of their income. Military spending therefore will seem to be progressive[7] and any cut in military expenditure will appear regressive assessed against conventional criteria. The combination of this conclusion (military spending is progressive) and the initial assumption (all individuals benefit equally) highlights the fact that equal cash valuation of some public service across the income distribution is *not* the same as saying that the impact of spending on that service is distributionally neutral. On the contrary, distributional neutrality is more conventionally associated with proportionality to income.[8]

The second example is one where the concept of usage is easier to define than in the case of spending on the military. Consider two families with different income levels, both with a child in a local maintained school that is facing a funding cut. Both families can be considered to be 'using' the school to the same extent. However, it is quite possible that, even if both sets of parents have the same sense of the importance of education to their children, the richer family will be willing to pay more than the poorer family to avoid the cut in the school's funding, simply because it can afford to.

These examples suggest that if one is satisfied to put a value on public service spending that is informed by the notion of willingness to pay, then cash valuations should rise with income, at least up to some particular level of income. This is not to suggest, however, that valuation increases with income throughout the income distribution. As incomes continue to rise, individuals become more likely to supplement public provision with private provision of a close substitute and, in a more extreme case, opt out of public provision altogether and rely entirely on privately provided alternatives. Once incomes get to the level that individuals start supplementing public provision with private alternatives, valuations will stop rising. If individuals opt out of public provision altogether – say by sending their children to a private school rather than the local maintained school – then their valuation of spending on maintained schools could well be zero.

---

[7] Where the definition of progressivity is taken to be that the value of the benefits as a proportion of income falls as income rises. The definition of neutrality is that the value of the benefits as a proportion of income is constant, and that of regressivity is taken to be that the value of the benefits as a proportion of income rises as income rises.

[8] Another way to think about this is to note that 'progressivity' in spending is not the same as 'equal benefits', any more than progressivity in taxes is the same as equal tax payments.

To summarise this discussion: we have outlined that cash valuations placed on public services are likely to vary over the income distribution; even with the same preferences, those with more income will be willing to pay more. Assuming that value and cost are equal will mean that this variation in value will be missed.

Some 'mitigating factors' to the equating of cost and value

The previous discussion outlined many reasons *not* to use cost as the basis for valuation. However, the practical difficulty of adequately ascertaining willingness to pay means that if the aim is to value either the current level of, or the change in, spending on a variety of different public programmes (such as the decisions contained in the Spending Review), cost is perhaps the only feasible starting point.[9]

Given the fact that the previous discussion was quite clear on the problems associated with using the cost of provision as a measure of value, here we note three of the ways in which the problems that are associated with such an approach might be mitigated.

First, studies that take this approach typically correct for differential usage of public services across the income distribution. Differential usage is one (though only one) of the reasons why value might differ across the income distribution.[10]

Second, under certain assumptions, there is likely to be a point in the income distribution where the cost of provision is equal to the willingness to pay.[11] This insight, combined with evidence gleaned from other sources on how valuations vary with income, could, in principle, be used to estimate a valuation for all households that, while depending on the cost of provision, is not exactly equal to that cost. However, the studies that we have seen and that we discuss in Section 8.3 do not attempt anything like this. They typically divide the cost of provision by the number of users and assign the same cash valuation to each user.

Third, some studies attempt to estimate the value of a *change* in spending on public services rather than the current level of public spending. Unless the political system is considered to be completely dysfunctional, it is probably the case that the average valuation of a particular change is not too far from the average change in the cost of provision. Information from other sources on the relationship between cost and value can then be used in conjunction with this (roughly correct) average figure to estimate the distributional impact.

The value of spending on public services versus the value of *changes* to spending on public services

Before discussing results, it is worth emphasising one further point on how the methods of valuation discussed above can – or cannot – be used to evaluate the distributional impacts of *changes* in spending on public services, such as those announced in the recent

[9] At least the cost of everything is relatively amenable to measurement. The economist may be in danger of being someone who miscalculates the value of everything because he is not satisfied with knowing the cost of everything.

[10] Defining and estimating usage of public services is itself not free from difficulties. Issues include whether to define the benefits of certain public services (healthcare, for example) as accruing to those who access them in a particular period or as insurance benefits that accrue to everyone in each period, and whether to measure use of public services over the life cycle or in a particular period. We do not discuss these issues here, but a detailed discussion is contained in section 2.5 of C. O'Dea and I. Preston, *The Distributional Impact of Public Spending in the UK*, 2020 Public Services Trust, London, 2010 (http://www.ifs.org.uk/publications/5234).

[11] See section 2.1 of C. O'Dea and I. Preston, *The Distributional Impact of Public Spending in the UK*, 2020 Public Services Trust, London, 2010 (http://www.ifs.org.uk/publications/5234).

Spending Review. If the distributional impact of a change in spending on public services is of primary interest, then every possible effort should be made to estimate the impact of that change directly, rather than it being informed by estimates of the distributional effects of the existing level of spending on public services. It may be possible to determine a reasonable estimate of the distributional impacts of aggregate spending on a particular public service. However, it will not necessarily be the case that the distributional impact of the *change* in spending on that service will mirror the distributional impact of the existing *level* of public expenditure on it. For example, simply because health spending tends to be progressive does not mean that every conceivable reduction in health spending is regressive. There are likely to be individual parts of health spending that are regressive, or at least less progressive than health spending as a whole. It is therefore clearly possible for the distributional impact of a *change* in health spending to bear little relation to the distributional impact of the *current level* of health spending.

The precise composition and manner of implementation of a package of spending cuts will be of crucial importance in determining how progressive or regressive it is. This level of detail will not generally be included in a fiscal statement such as the Spending Review, in part because much of the requisite detail might not be decided for some time. This fact should warn against estimating the value of the current level of public expenditure and then inferring from this estimate the distributional impact of a package of changes to spending on public services.

8.3 What can existing studies tell us about the distributional impact of planned cuts?

This section discusses two recent analyses of the distributional impacts of spending on public services. The first is the most recent edition of an annual study carried out by the Office for National Statistics (Barnard, 2010)[12] that evaluates the impact of government taxes, cash transfers and certain transfers in kind on household income. The second is an analysis published by the Treasury alongside the Spending Review in October 2010, which also estimated the distributional impact of some of the measures that were announced at that time.[13] The ONS analysis is of the distributional impact of the current level of spending on public services, while the Treasury analysis is of the impact of the planned cuts in spending on public services announced in the Spending Review.

There are a number of other studies that we do not discuss that similarly attempt to value elements of public expenditure. These include Sefton (2002), Volterra Consulting (2009) and Horton and Reed (2010),[14] which report findings broadly similar to those discussed below. Additionally, there is a wealth of research that looks at the distributional impact of

[12] A. Barnard, *The Effects of Taxes and Benefits on Household Income, 2008/09*, Office for National Statistics, 2010 (http://www.statistics.gov.uk/CCI/article.asp?ID=2440).

[13] Annex B of HM Treasury, *Spending Review 2010*, Cm 7942, 2010 (http://cdn.hm-treasury.gov.uk/sr2010_completereport.pdf).

[14] T. Sefton, *Recent Changes in the Distribution of the Social Wage*, CASEpaper 62, London School of Economics, 2002 (http://sticerd.lse.ac.uk/dps/case/cp/CASEpaper62.pdf); Volterra Consulting, *The Fiscal Landscape: Understanding Contributions and Benefits*, 2020 Public Services Trust, London, 2009 (http://www.2020publicservicestrust.org/publications/item.asp?d=1528); T. Horton and H. Reed, *Where the Money Goes: How we Benefit from Public Services*, Trades Union Congress, London, 2010 (http://www.tuc.org.uk/extras/wherethemoneygoes.pdf).

individual public services (e.g. health, education and police). For a summary of this literature, see section 3 of O'Dea and Preston (2010).[15]

ONS study (Barnard, 2010): the effects of taxes and benefits on household income, 2008/09

This (annual) study investigates what the distribution of income would look like if the cost of provision of certain transfers in kind used by a particular household is taken to be effectively part of that household's income. Households are divided into 10 equally sized groups on the basis of their equivalised income (these groups are called 'deciles') and the extent of the net transfer from the state to households in each of these deciles is compared.

Taxes, benefits and spending on public services are included in the study if the household that is the payer of the tax, the recipient of the benefit or the user of the service can be satisfactorily identified. So, for example, corporation tax is not included since the extent to which its ultimate incidence falls on different households is unclear. Similarly, expenditure by the Ministry of Defence is not included as differential usage of the military cannot be identified. Those taxes that *are* allocated account for 55% of general government expenditure and the cash and non-cash benefits that are allocated account for 52% of general government expenditure.[16]

This tax and spending is allocated to households across the income distribution (using the Family Resources Survey and the Living Costs and Food Survey – surveys of representative samples of the household population) on the basis of their payment of taxes, their receipt of benefits and their estimated usage of government services in kind. The vast majority (97%) of the spending on transfers in kind that is included is comprised of spending on health and education.[17] Health service usage for each household is estimated on the basis of the age and sex of its members. Education spending is allocated to households according to how many members of that household are currently enrolled in state-funded primary, secondary or tertiary education.

Figure 8.1 presents a decile chart using data from Barnard (2010) that illustrates some of the findings. For each decile, the left bar shows the average transfer of benefits in kind provided to the households in that decile as a percentage of income, where income is net of all modelled state transfers (whether in cash or in kind). The bar to the right shows the average net transfer of the total of taxes, cash benefits and benefits in kind as a percentage of income. Note that while the measure of income used in the denominator (for both bars) here *includes* benefits in kind, the income deciles are constructed using a measure of income that *excludes* benefits in kind; a similar analysis that ranked households based on their income inclusive of the value of benefits in kind would have been informative but was not included in Barnard (2010).

[15] C. O'Dea and I. Preston, *The Distributional Impact of Public Spending in the UK*, 2020 Public Services Trust, London, 2010 (http://www.ifs.org.uk/publications/5234).

[16] See page 3 of Barnard (2010) and additionally table 13.

[17] Other than health and education spending, other smaller items included are some housing expenditure, transport subsidies and spending on school meals and welfare milk.

Figure 8.1. An estimate of the distributional impact of some taxes, benefits and benefits in kind in 2008–09

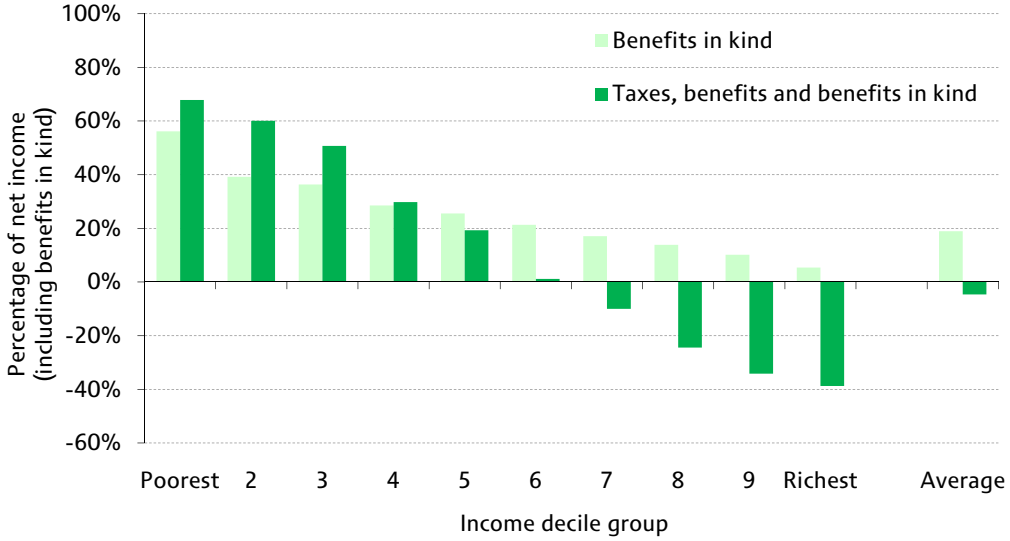

Notes: Income decile groups are derived by dividing all households into 10 equal-sized groups according to income adjusted for household size using the McClements (before-housing-costs) equivalence scale. Decile group 1 contains the poorest tenth of the population, decile group 2 the second poorest, and so on up to decile group 10, which contains the richest tenth. Columns represent net transfer from state expressed as a percentage of a measure of income that *includes* that net transfer.
Source: Table 14 of Barnard (2010).

The fact that the modelled benefits in kind have a marked progressive effect is a result of three factors. First, those at the bottom of the income distribution are more likely to have children than those further up, so are more likely to be using and benefiting from state spending on education.[18] Second, even conditional on having children, those at the bottom of the income distribution are more likely to be using the state education system than those further up, who are more likely to opt out and use private education. Third, the heaviest users of the health services are the elderly, who are more likely than younger people to be in the bottom half of the income distribution.[19]

Together, these factors result in the absolute (i.e. pounds) value of modelled benefits in kind falling over the income distribution, from approximately £6,600 per year on average for those in the bottom income decile to £3,600 on average for those in the top income decile. Expressing these values (or costs) of services received as a proportion of income, this decline is even more marked.

Notwithstanding the caveats noted in the previous section about interpreting distributional analyses that rely on assuming that value is equal to cost, the redistributive impact of the government interventions modelled here is clear. Those at the bottom of the income distribution receive, on average, a substantially greater share of their total access to resources from the state. Those at the top of the income distribution receive, on average, only a very small proportion of their total income in terms of services in kind

[18] This effect partly comes from the fact that households tend to have lower income when they have young children than later in life. If the income measure were a broader one that sought to measure income over the life cycle, then this effect, if evident at all, would be much weaker. It also comes from the use of equivalised incomes, which means that a household's assumed standard of living falls when children arrive.

[19] Once again, this effect is partly driven by the measure of income being a current rather than lifetime income measure. Clearly, everyone who is old was young once, and everyone who is young has an expectation of being old. So the distributional impact of health expenditures could look very different once a life-cycle perspective is brought to the analysis.

from the state. Of course, adding in the effect of taxation (which is mildly progressive on average) and cash benefits (which are strongly progressive on average) makes the overall picture even more progressive than that found by looking at the effect of benefits in kind in isolation.

HM Treasury's distributional analysis of decisions announced in the 2010 Spending Review

The Spending Review announced substantial cuts in spending on public services to take place over the next four years. Annex B to the 2010 Spending Review document contained estimates of the distributional impacts of some of the measures that were announced as part of the Review. The decisions for which impacts were modelled included some changes to welfare benefits, some changes to taxes and, of primary interest here, some changes to spending on public services.

The approach taken by the Treasury shares many features with that taken in the ONS study discussed above. In particular, value is assumed to be equal to the cost of provision, and only expenditure on items where it was considered that the end-user could be identified were modelled. As a result of the latter restriction, approximately half of Departmental Expenditure Limits (DELs) were included in the analysis carried out, over 80% of which was spending on either health or education.[20] The modelled components, however, account for only about one-third of the *changes* in DELs.[21]

Figure 8.2. HM Treasury analysis of the distributional impact of planned cuts in spending on public services, by 2014–15

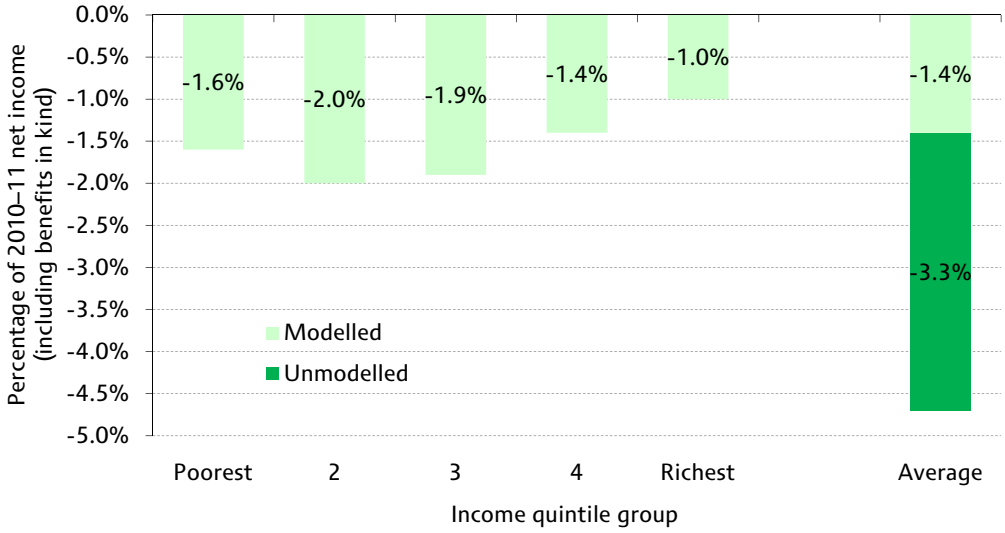

Notes: Income quintile groups are derived by dividing all households into five equal-sized groups according to income adjusted for household size using the McClements (before-housing-costs) equivalence scale. Quintile group 1 contains the poorest fifth of the population, quintile group 2 the second poorest, and so on up to quintile group 5, which contains the richest fifth.
Sources: Data on modelled expenditure are from chart B.6 of Spending Review 2010. Unmodelled expenditure is from authors' calculations using data from annexes A and B of Spending Review 2010 and private communication with HM Treasury officials.

[20] Excluded were almost all capital expenditure, spending on 'pure public goods' such as defence and environmental protection, and central government administration costs.

[21] The proportion of the *change* in expenditure modelled is substantially less than the proportion of the current *level* of expenditure modelled because health expenditure makes up over 50% of the level of expenditure that *is* modelled but was an area that was protected from real cuts in the Spending Review.

Figure 8.2 shows the Treasury's analysis of the distributional impact of the modelled cuts to spending on public services. It gives the fall in value obtained from the modelled components as a percentage of households' total income (defined as net cash income plus the value of modelled public spending) relative to a baseline of no real change in DELs. The importance of the components whose distributional impact has not been modelled is clear from the size of the 'unmodelled' bar in the column labelled 'average'.

The biggest proportionate losers are those in the second income quintile. The proportionate impact of the spending cuts then falls as income increases, making the modelled cuts appear broadly regressive over the richest 80% of the population. The poorest fifth, however, lose less on average than those in the second and third income quintiles.

As an analysis of the distributional impact of the package of planned cuts to spending on public services, this picture is, of course, incomplete since the effect of two-thirds of the cuts in DELs is missing. The modelled elements include some clearly progressive elements – for example, the introduction of the pupil premium (which benefits households with poor children) and the deep cut in taxpayer subsidies for the teaching budgets of higher education institutions (which are one of the areas of public spending where the 'users' of the service tend to be better off).[22] But without any analysis of the progressivity or otherwise of the unmodelled elements, we cannot make any concrete assessment of the overall distributional impact of the cuts to spending on public services announced in the Spending Review.

In addition to its estimation of the distributional impact of the planned cuts to spending on public services, the Treasury estimated the distributional impact of changes to tax and cash benefits. Figure 8.3 shows the combined impact of the changes to spending in kind

Figure 8.3. HM Treasury analysis of the distributional impact of changes in taxation, benefits and spending on public services, by 2014–15

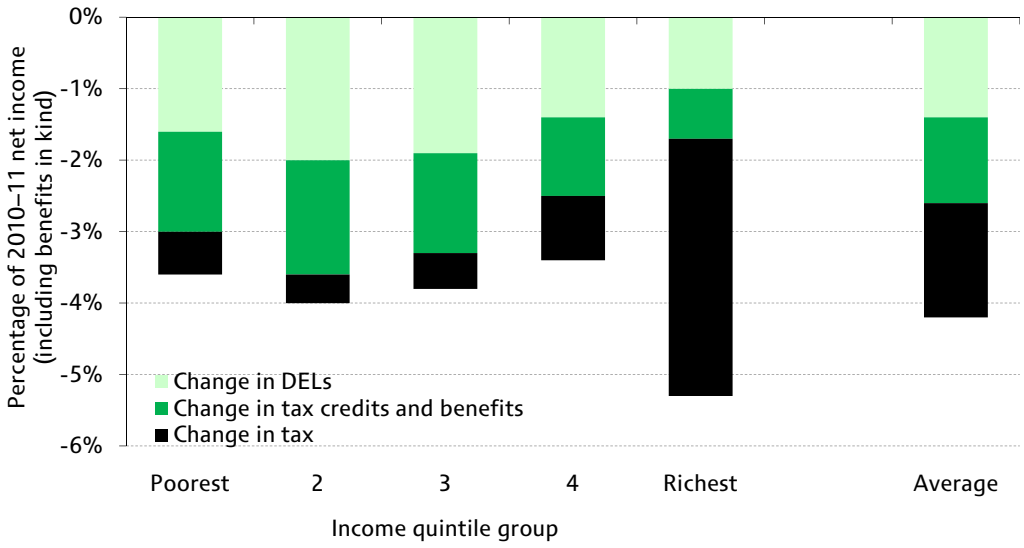

Notes: As Figure 8.2
Source: Chart B.6 of Spending Review 2010.

[22] For an evaluation of the reforms to the funding of higher education, see H. Chowdry, L. Dearden and G. Wyness, *Higher Education Reforms: Progressive but Complicated with an Unwelcome Incentive*, IFS Briefing Note 113, 2010 (http://www.ifs.org.uk/publications/5366).

discussed above and the changes to taxation and social security payments. The latter include measures announced in the Spending Review, those announced in the June Budget and measures pre-announced by the previous government on which the current government is intending (broadly) to legislate. Adding in these changes to those whose impact was shown in Figure 8.2 does not substantially alter the relative positions of the quintiles that make up the bottom 80% of the population. Here, however, the income quintile that, on average, loses the most is the richest quintile. This fact is largely due to a reduction in the generosity of tax relief on pension contributions, a measure pre-announced by the previous government (although the coalition government has altered the precise form this restriction will take) that will only affect a small number of rich individuals who contribute more than £50,000 a year to a private pension.

One further comment is worth making on Figure 8.3. The Treasury only includes 'measures where there is sufficiently robust data available to attribute changes in tax, tax credits or benefits to individuals'.[23] For this reason, recent changes to Housing Benefit, among other decisions, are not considered. Neglecting to include the changes to Housing Benefit makes the package of changes recently announced look more progressive, as those who have been affected by changes to Housing Benefit are located largely towards the bottom of the income distribution.[24] It is certainly harder to estimate precisely the distributional impact of changes to Housing Benefit than that of many of the other forthcoming changes to taxes and benefits, but it is most certainly easier to do so than to estimate the distributional impact of changes in spending on public services. It is odd, therefore, to produce an analysis that *includes* the distributional impact of some of the changes to spending on public services but that *excludes* the distributional impact of measures that are substantially easier to model. We suggest, therefore, that the next time the Treasury publishes an analysis showing the distributional impact of decisions that transfer resources in cash and those that transfer resources in kind, it models a more comprehensive set of the former type of decisions than previously. At the very least, for those decisions for which it is felt that a distributional breakdown simply is not possible, we suggest that that the Treasury should be very clear about their size. One way to do this is to include them in the 'average' column so that their magnitude is clear, even if their distributional impact is not (along the lines of the 'unmodelled' bar shown in Figure 8.2). Further detail should also be provided: for example, breaking down the percentage of the overall fiscal consolidation package that has been modelled into the percentage of tax rises, the percentage of welfare cuts and the percentage of cuts to public service spending that have been included.

Finally, it is important to note that the Treasury has revealed very little detail on precisely how these estimates were obtained. Much of the variation in distributional impact is coming from its estimates of differential 'usage' of public services by those with different incomes. The manner in which this 'usage' is estimated has not been forthcoming, making any evaluation of the robustness, or otherwise, of the distributional analysis impossible.

[23] Paragraph B.34 of Spending Review 2010.

[24] For a distributional analysis of recent changes to taxes and benefits that includes the changes to Housing Benefit, see J. Browne, 'Distributional analysis of tax and benefit changes', presentation at IFS 2010 Spending Review Briefing, 21 October 2010 (http://www.ifs.org.uk/publications/5313).

A questionable rationale for omitting certain expenditure?

Of course, the distributional impact of public services estimated both by Barnard (2010) and by the Treasury is incomplete as the distributional impact of much of government activity is not modelled. Both model only the value of expenditure where usage of the service varies across the income distribution (though both are very clear about this fact).[25]

The rationale for omitting items where usage does not differ across the population is not clear. There are two principal requirements for evaluating the distributional impact of public services. First, differential usage must be identified, where this is relevant. Second, conditional on usage, valuation must be assessed. The difficulties inherent in the first step (measuring usage) are avoided completely among those services such as defence and environmental protection where usage is (presumably) the same across the population. These are exactly the items that are omitted from the distributional analyses discussed here. Of course, the second requirement (measuring valuation) must be addressed, but there is no reason to believe this step is harder for items where usage is uniform than where it is not.

In short, there is a danger in assuming that the methodology for valuation applied here (which assumes value equals cost, and that value is the same across all households that use the service) is any more reliable for public services that are used to different extents by different households. Allocating defence spending or spending on environmental protection uniformly across all households could be just as good (or just as bad) as allocating other spending uniformly across all households that use them.

How progressive or regressive will the planned cuts to public service spending be?

The issues outlined in Section 8.2 show that there are too many difficulties to allow a precise, quantitative assessment of how progressive or regressive the planned cuts to spending on public services will be. However, we can draw on the research discussed earlier to make some tentative statements about the distributional impact of the spending cuts that have been announced and will be implemented over the current parliament.

Given that the use of public services is concentrated on those with lower incomes, it will be uncontroversial to assert that government spending (or at least spending on those public services where use can be satisfactorily defined) is, broadly speaking, progressive. A cautious initial presumption has to be, then, that cutting government expenditure is more likely to be regressive than progressive.

Even this tentative assessment, however, merits a couple of qualifications. First, the precise distributional impact of the planned cuts will largely be determined by the manner in which the cuts in Departmental Expenditure Limits are dealt with by individual departments – and it will be some time before this will be clear. The extent to which the pain falls largely on services that are used to a greater extent by those further

[25] There have been two recent attempts to allocate the remainder of government expenditure to households: Volterra Consulting, *The Fiscal Landscape: Understanding Contributions and Benefits*, 2020 Public Services Trust, London, 2009 (http://www.2020publicservicestrust.org/publications/item.asp?d=1528); T. Horton and H. Reed, *Where the Money Goes: How we Benefit from Public Services*, Trades Union Congress, London, 2010 (http://www.tuc.org.uk/extras/wherethemoneygoes.pdf).

up the income distribution will limit the presumption towards the regressivity of the spending cuts. Second, we focused here on the extent to which *use* of public services is concentrated towards the bottom of the income distribution. As discussed in Section 8.2, differential usage is only one reason that the value of public services might be different at different points of the income distribution. If those users higher up the income distribution have a higher willingness to pay for public services, this will once again limit the presumption towards the regressivity of the spending cuts.

8.4 Conclusion

This chapter has outlined the considerable difficulties in establishing precisely the distributional impact of changes to spending on public services such as those announced in the Spending Review. The difficulties come from three sources in particular:

1. It is often not clear to what extent people with different incomes use public services differentially.
2. It is not clear how much people with different incomes value the services they receive from the state, or to what extent they would be affected if the service were removed or curtailed.
3. The link between any change in funding available to individual departments and exactly which individual services will be removed or curtailed will not always be clear at the time of the announcement of the cuts to departmental budgets.

Each of these difficulties is either not relevant or substantially less acute when the aim is to establish precisely the distributional impact of changes to taxation or cash benefits. This means that an evaluation of the distributional impact of changes to cash transfers can be undertaken with substantially greater precision than a similar evaluation of changes to spending on public services.

In spite of these difficulties, existing research on public spending can inform an understanding of what effect the planned spending cuts might have. Use of public services (or at least those public services where 'use' can be defined and can be thought of as different among different households) is more concentrated at the bottom of the income distribution. As a result, while the ultimate distributional impact will depend on precisely how the planned spending cuts are implemented, it is fair to say that those towards the bottom of the income distribution have more to lose from spending cuts. Indeed, the Treasury's distributional analysis shows the planned cuts to public service spending having a regressive effect over the 80% of the population with the highest incomes (including assessed benefits in kind), with those in the poorest 20% losing a little less in proportionate terms than those who are slightly richer. However, that analysis only includes approximately one-third of the change in spending on public services.

The challenge of limiting the impact of cuts to public service spending on those at the bottom of the income distribution would have been faced by *any* government planning to bring borrowing back to sustainable levels through large cuts to spending on public services. It is worth noting that the plans in the final Budget of the previous government also suggested substantial falls in departmental spending.[26]

[26] See Office for Budget Responsibility, *Pre-Budget Forecast: June 2010* (http://budgetresponsibility.independent.gov.uk/d/pre_budget_forecast_140610.pdf) and R. Chote, R. Crawford, C. Emmerson and G. Tetlow, *Filling the Hole: How Do the Three Main UK Parties Plan to Repair*

This chapter started by noting the importance of attempting to assess the distributional impact of changes to spending on public service. Given the magnitude of the spending cuts announced in the recent Spending Review, an awareness of their distributional consequences is essential. As such, it is to be welcomed that the Treasury published a distributional analysis of some of the changes in expenditure announced in the Spending Review. However, we make three specific recommendations relevant to its analysis.

- First, we agree wholeheartedly with the Treasury Select Committee's recommendation that 'the Treasury publish not just the sources but additional information on the calculations underpinning their distributional analysis to provide further transparency and encourage debate on how the methodology of such analysis might be improved'.[27] Such calculations were not published alongside the distributional analysis in the Spending Review. The type of analysis carried out does not have an established methodology (unlike the distributional analysis of changes to taxation and cash transfers) and, as a result, the credibility of the published results relies crucially on how they are derived.
- Second, the many issues discussed in this chapter imply that the results on the impact of benefits in kind should be interpreted cautiously. Any single result on the overall distributional impact of a package of spending cuts, whether produced by the Treasury or other authors, should be interpreted with care and, while informative, should not be considered definitive. In future official distributional analyses, it would be good to see an assessment of how sensitive the headline result is to changes in the underlying assumptions.
- Third, more care should be taken and explanation given regarding why certain elements of public spending (either on cash transfers or public services) are excluded from the distributional analysis. In particular, the Treasury only models the distributional impact of public services where differential usage across households can be ascertained. The case has not been made for why their method of valuation is any more reliable for those services that are used differentially than for those that are used to the same extent by all residents (such as spending on environmental protection). Additionally, in its broadest distributional analysis, the Treasury *excludes* the consideration of certain decisions (changes to Housing Benefit, for example), on the basis that the distribution of valuations is hard to model, but *includes* other decisions that are certainly harder to model (changes to education spending, for example). We recommend that future analyses of this type take a more consistent approach to determining which measures are included.

the Public Finances?, IFS 2010 Election Briefing Note 12 (IFS Briefing Note 99), 2010 (http://www.ifs.org.uk/publications/4848).

[27] Paragraph 83 of House of Commons Treasury Committee, *Spending Review 2010*, Sixth Report of Session 2010–11, 2010 (http://www.publications.parliament.uk/pa/cm201011/cmselect/cmtreasy/544/544i.pdf).