Can We Measure Who Loses Most from Public Service Spending Cuts?

Cormac O'Dea
Ian Preston

BUDGET PERSPECTIVES 2015
PAPER 2

June 2014

ESRI

Can We Measure Who Loses Most from Public Service Spending Cuts?

Cormac O'Dea
Ian Preston

BUDGET PERSPECTIVES 2015
PAPER 2

June 2014

Available to download from *www.esri.ie*

The ESRI

The *Economic Research Institute* was founded in Dublin in 1960, with the assistance of a grant from the Ford Foundation of New York. In 1966 the remit of the Institute was expanded to include social research, resulting in the Institute being renamed the *Economic and Social Research Institute* (ESRI). In 2010 the Institute entered into a strategic research alliance with Trinity College Dublin, while retaining its status as an independent research institute.

The ESRI is governed by an independent Council which acts as the board of the Institute with responsibility for guaranteeing its independence and integrity. The Institute's research strategy is determined by the Council in association with the Director and staff. The research agenda seeks to contribute to three overarching and interconnected goals, namely, economic growth, social progress and environmental sustainability. The Institute's research is disseminated through international and national peer reviewed journals and books, in reports and books published directly by the Institute itself and in the Institute's working paper series. Researchers are responsible for the accuracy of their research. All ESRI books and reports are peer reviewed and these publications and the ESRI's working papers can be downloaded from the ESRI website at www.esri.ie

The Institute's research is funded from a variety of sources including: an annual grant-in-aid from the Irish Government; competitive research grants (both Irish and international); support for agreed programmes from government departments/agencies and commissioned research projects from public sector bodies. Sponsorship of the Institute's activities by Irish business and membership subscriptions provide a minor source of additional income.

The Authors

Cormac O'Dea is a Senior Research Economist at the Institute for Fiscal Studies. Ian Preston is Professor of Economics and Deputy Head of Department, University College London.

Acknowledgements

We are grateful to Tim Callan, Carl Emmerson, Claire Keane, Barra Roantree and an anonymous reviewer for helpful comments and to Hannah Mills and Jonas Nystrom for discussions around the UK Treasury's distributional model. This paper draws from and contains material from two previous pieces by the same authors: a (reasonably technical) working paper: 'The distributional impact of public spending in the UK' (http://www.ifs.org.uk/publications/6076) and a non-technical chapter in the Institute for Fiscal Studies' 2011 'Green Budget': Measuring the distributional impact of public service cuts' (http://www.ifs.org.uk/budgets/gb2011/11chap8.pdf). The views expressed here are those of the authors and not those of the Institute for Fiscal Studies, which has no corporate views.

Table of Contents

Can We Measure Who Loses Most from Public Service Spending Cuts?

Abstract

The fiscal tightening currently under way in Ireland, as in many other countries, comprises cuts to spending on public services (transfers in kind) as well as cuts to benefits and increases in taxation (transfers and charges in cash). Unlike with changes to taxes and cash benefits there is no standard methodology for assessing which groups in society lose most from changes to spending on public services. Therefore distributional analyses of government 'cash' decisions are common while those relating to 'in-kind' decisions are rare. This paper considers the reasons for this, sets out some issues which must be tackled in modelling the impact of changes to public service spending and discusses some approaches that are being undertaken in other countries. We highlight the fact that such a modelling exercise will yield only imprecise result. Finally, we assess whether it is worth developing a model of the distributional impact of changes to public services in Ireland.

1. INTRODUCTION

The fiscal tightening ongoing in Ireland comprises a combination of net increases in taxation, cuts to spending on cash benefits and 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, an understanding of how the impact of each is distributed across the population is important.

Distributional analyses of changes to taxation and the payment of cash benefits are frequently produced (in Ireland and elsewhere)[1] and follow an established methodology[2]. For reasons that we discuss in this paper, distributional analyses of changes to spending on public services (health, education, physical infrastructure, government administration etc.) are much less common. However, the recent large cuts to departmental budgets have sparked a good deal of interest in this recently, in the UK, Ireland and elsewhere. The issue is not, however, new: the absence of any modelling of 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 expanding public expenditure.

The comments above are in no sense particular to Ireland. Elsewhere, on the cash side, tax and benefit models are also used in policy formation by governments and in policy evaluation by those outside government. Also, formal analysis of the distributional impact of the recent cuts to spending on public services have been rare internationally, just as they were absent in times of increasing spending. Recently though, there have been some models developed and we discuss one such model (that of the UK Treasury). We do not discuss this by way of a recommendation that an exact replica should be produced for Ireland – indeed there is much that we would like to see done differently – but as an example of what could be done, what issues would have to be considered in designing a model, what the results would look like, what caveats would have to be attached to these results and what insights could be drawn from them. We hope this can contribute to a discussion about whether such an endeavor would be a worthwhile use of resources in Ireland.

This paper proceeds as follows. Section 2 outlines the considerable difficulties in evaluating precisely the distributional impact of public services that are provided

1 For a recent example in Ireland, UK and a selection of European countries (though not including Ireland) respectively see Callan *et al*. (2013), Section 7.5 of Joyce and Phillips (2013) and Avram *et al.* (2013).

2 For an overview of this methodology see Hancock and Sutherland (1992) and Redmond *et al.* (1998).

in kind rather than in cash. Section 3 turns to a description of the model that the UK Treasury have developed that attempts to look at the distributional impact of *changes* in such spending. Section 4 concludes. [3]

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 a new motorway? Or the Department of Finance? 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, roads or any government department.

The impact of public service spending decisions is substantially harder to document precisely than 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 health service?

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 euro per week, for example) that the household has lost. In assessing the effect of a cut in spending on the health service, on the other 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 to 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 (like the value of public services) that has no simple quantitative

[3] A more detailed treatment of the same issues is considered in O'Dea and Preston (2010).

interpretation.[4] 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 service spending (e.g., closing libraries compared with closing hospitals). Putting these together, it would allow a distributional analysis of an *entire* fiscal consolidation or expansion, 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). If all else is equal then modelling the impact of the full set of government decisions is better than only a particular subset of them – doing the latter could give a misleading impression of the true impact of these decisions on households.

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 the value to the user with the cost of provision.

2.1 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 good school, 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 examples of this type of research see Gibbons and Machin (2003 and 2008).

[4] 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 be often better than some number, see: Diamond and Hausman (1994).

- 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 which aims to value public health care.
- 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 or whether they would be willing to pay a given amount. For an example see Hall and Preston (1998).

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 on Budget day. The latter typically sets out the budgets available to departments rather than on which 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.

2.2 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. We briefly discuss a selection of these papers below. 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 and 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 this approach.

2.2.1 Value and Cost Diverge

The fact that value and cost diverge can be illustrated by a stark example. 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, governments

often state an intention to cut 'wasteful' expenditure.[5] 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 no willingness to pay for waste, and it would therefore be assigned a zero value.

2.2.2 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 spending on the environment. Without any way of assessing 'usage' of the environment, perhaps the most obvious approach is to 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 environment as a proportion of their income. Spending on the environment therefore will seem to be progressive,[6] and any cut in environmental expenditure will appear regressive. While this may or may not be true, the combination of this conclusion (environmental 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.[7]

---

5 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 a transfer payment. The discussion in the text here relates primarily to 'pure waste'.

6 Where the definition of progressivity is taken to be that the value of the benefits as a proportion of income falls as incomes rise, that 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.

7 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.

The second example is one where the concept of usage is easier to define than in the case of spending on the environment. Consider two families with different income levels, both with a child in a local school which 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 share the same sense of the importance of education to their children the richer of the two families will be willing to pay more than the poorer family to avoid the cut in the school's funding, simply because they 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 stage 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 state school – then their valuation of spending on state schools could well be zero.

It is very important to emphasise that the fact that value (as defined by willingness to pay) may increase with income over some range does *not* imply that services provided to richer households are more valuable from the perspective of society or should carry more weight in policymakers' decisions. In policymaking, a social welfare function will be applied (albeit not formally) to any distributional analysis that will likely place a greater weight on transfers and services on lower income households than on higher income households.

To summarise this discussion we have outlined the following: 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.

2.2.3 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 Budget), cost is perhaps the only feasible starting point.[8]

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.[9]

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. [10] 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 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 services rather than the current level of public spending. Unless the political system is considered to be completely dysfunctional, it is probably the case for some items of expenditure (such as those where the relationship between value and cost is reasonably smooth) 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.

---

8 At least the cost of everything is relatively amenable to measurement. However, 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.

9 Defining and estimating usage of public services is itself not free from difficulties. Issues include whether to define the benefits of certain public services (health for example) as accruing to those who access them in a particular period, or as insurance benefits which accrue to everyone in each period and whether to measure use of public services over the lifecycle or in a particular period. We do not discuss these issues here but a detailed discussion is contained in Section 2.5 of O'Dea and Preston (2010).

10 See section 2.1 of O'Dea and Preston (2010).

Finally, while it should be clear that the cost of provision is not the same as the value to the recipient, ascertaining the areas of the income distribution which receive greater or lesser shares of public spending is arguably a well-formed question and one to which the answer is informative, as long as it is interpreted appropriately.

2.2.4 Studies Evaluating Distributional Impact of the Public Services Using Input Cost Approach

Many existing and ongoing studies evaluate the distributional impact of spending on public services (the term Social Transfers In Kind (STIKS) is sometimes used to describe such spending) using the approach of assuming that input cost is equal to the value to the user. These include an annual publication in the UK by the Office for National Statistics (the most recent edition of which was published in 2013), a paper that implements a similar procedure in Ireland Callan and Keane (2009) on health and some non-compulsory education expenditures, and an ongoing programme of work at the World Bank looking at a developing and middle income countries (see Lustig *et al.* 2013). Cross national studies on OECD countries, including Ireland, have been undertaken by Verbist *et al.* (2012) and on a selection of European countries include Paulus et al (2010) and Aaberge *et al.* (2010) (though neither paper considers Ireland). Also, the Australian Survey of Income and Housing microdata is now published with estimates of the value of STIKS at the individual household level.

These studies typically only model the distributional impact of public services where differential use across households can be identified (so they, for example, typically include health but not environmental spending).[11] The studies also typically do not take into account the fact that greater spending focused on a particular part of the income distribution might be as a result of differing needs. For example, poorer people tend, in many countries, to use public health care to a greater extent than richer households, but this is, at least partly, as a result of poorer people tending to be unhealthier. Ignoring the fact that some greater spending is simply as a result of pre-existing inequalities whose effect is not adequately represented by income statistics will over-state the well-being of more needy recipients relative to those less in need of that particular spending.

Notwithstanding these caveats, and the caveats around assuming that value is equal to cost, the redistributive impact of the government interventions

[11] We are aware of two attempts to allocate the remainder of government expenditure in the UK to households. These are Volterra Consulting (2009), and Horton and Reed (2010).

modelled in these studies is clear. Those at the bottom of the income distribution receive a substantially greater share of their total access to resources from the state. Those at the top of the income distribution, on average, receive only a very small proportion of their total income in terms of services in kind from the state.[12]

The studies referenced above mean that, if we are willing to accept cost of provision as a valid metric for the value to a recipient of a service, there is quite a lot that is known (and more that can be known) about the distributional impact of in kind transfers.

2.3 The Value of the Level of Spending on Public Services Versus the Value of Changes to Spending on Public Services

It is worth emphasising one further point on how both 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 a Budget statement. 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 either 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.

Take a specific example: the potential introduction of free GP care for the under-six age group would likely have a very different distributional impact to the distributional impact of the current level of public spending on GP care. The fact that the poorest households are those exempt from GP fees makes current public spending on GP care in Ireland progressive. Free GP care to the under-sixes would not benefit such households – and the beneficiaries would be located further up the distribution.

[12] The studies referenced above model the impact of spending on several public services. Additionally, there is a wealth of research that looks at the distributional impact of individual public services (e.g., health, education, police etc.). For a summary of this literature see Section 3 of O'Dea and Preston (2010).

The precise composition and manner of implementation of a package of spending cuts will be of crucial importance in determining how progressive or regressive they are. This level of detail will not generally be included in a fiscal statement such as the Budget, 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 these estimates the distributional impact of a package of changes to spending on public services. 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 being informed by estimates of the distributional effects of the existing level of spending on public services. The devil will be in the detail.

3. The UK Treasury's Model of the Distributional Impact of Spending on Public Services

In this section, we summarise the methodology and results coming from the UK Treasury's (reasonably) new model which they use during the policy formation process and which is used to publish their assessment of the distributional impact of fiscal statements.[13] The Treasury have for many years had a model *(IGOTM*, similar to the ESRI's *SWITCH* and the IFS's *TAXBEN*) that attempts to measure the distributional impact of changes to taxes and cash benefits. What is new is the facility to assess the distributional impact of changes to spending on public services.

Such a model is rare internationally (indeed possibly unique – we are not aware of any other government that has undertaken to develop a comparable model) in the sense that it aims to look at the distributional impact of *changes* to spending on public services. However, at its core is actually a model of the distributional impact of the *level* of spending on public services. The approach taken, for expenditure on health care (as an example), is as follows:

1. A variety of data sources are used to estimate the extent to which those in different parts of the income distribution use publicly-provided health care services.
2. The total benefit from health care spending (which is set equal to total cost spending) is assigned across the income distribution in proportion to that estimated differential usage.

[13] Fiscal statements in the UK are the annual budget (which takes place in March), the 'Autumn Statement' (which despite its name most recently was delivered in December 2013) and the (approximately tri-annual) Spending Review at which departmental spending settlements are determined for the subsequent number of years.

3. Changes in spending on health care are assumed to affect households in direct proportion to the extent to which they currently benefit.[14]

As discussed above, there is absolutely no guarantee that this last step will give a reasonable estimate of the change in valuations, even if we set aside concerns about equating value and cost.

The first results from the model were presented alongside the Spending Review in 2010 and each fiscal statement since has contained an updated set of results.

The approach taken by the Treasury shares many features with that taken in the studies 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 just over half of (current and capital) spending on public services[15] were included in the analysis carried out, over 80 per cent of which was spending on either health or education.[16] The modelled components, however, account for a much smaller proportion of the change in spending on public services. The fact that the proportion of the *change* in expenditure modelled is substantially less than the proportion of the current *level* of expenditure modelled comes from the fact that health expenditure makes up over half of the level of expenditure that *is* modelled but was an area that has been protected from real cuts since the current government in the UK came into office (2010). At the time of the Spending Review in 2010 the modelled expenditures represented about one-third of the *changes* in spending on public services. The Treasury have not published sufficient detail for us to update this figure[17] – though it is likely not to have greatly changed.

Figure 1 shows the Treasury's analysis of the distributional impact of the modelled cuts to spending on public services. It gives the proportionate fall in value obtained from the modelled components as a proportion of their total

[14] There a few small exceptions to this – i.e., cases where a particular change in expenditure has been explicitly modelled. A recent extension to free school meals is an example of this.

[15] We are using the terms spending on public services here as a shorthand for 'Departmental Expenditure Limits' (DELs). UK government spending is divided into Annually Managed Expenditure (largely social security payments, payment on debt interest, net pension payments to public servants and some local government expenditure) which, as the name suggests, are determined annually and Departmental Expenditure Limits (the cash available for departments to provide public services).

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

[17] However, one could carry out a back of the envelope calculation get an approximation to the answer.

income (defined as net cash income plus the value of modelled public spending) relative to a baseline of no real change in DELs.

The biggest proportionate losers are those in the bottom income quintile. The proportionate falls are roughly similar for those in each of the four upper income quintiles.

FIGURE 1 UK Treasury Analysis of Impact of Planned Cuts in Spending on Public Services Between 2010/11 and 2015/16

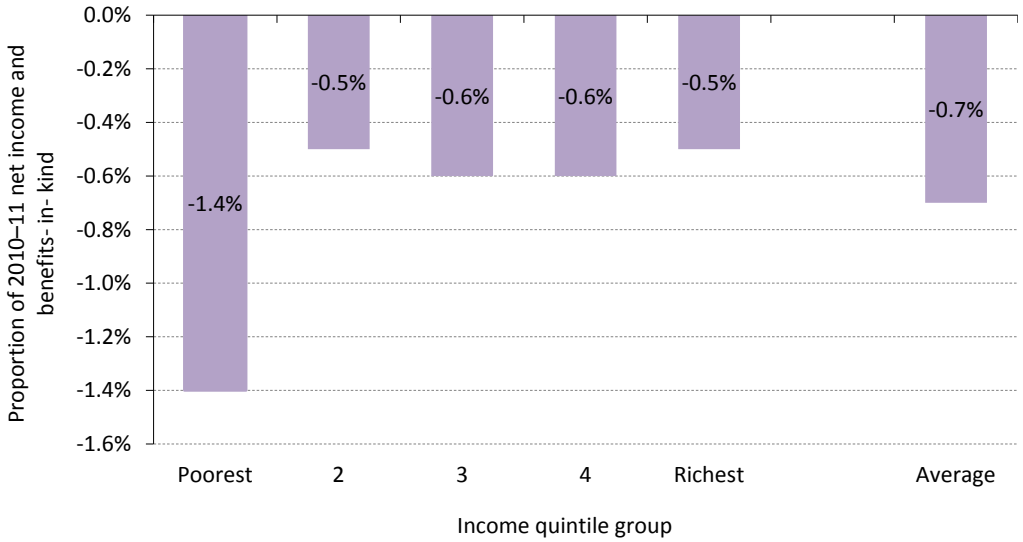

Note: Income quintiles groups are derived by dividing all households into 5 equal-sized groups according to income adjusted for household size using the McClements (before housing costs) equivalence scale. Quintile group 1 contains the poorest fifth of the population, quintile group 2 the second poorest, and so on up to quintile group 5, which contains the richest fifth.

Sources: Chart 2I, *UK Budget 2014*: Impact on households: distributional analysis to accompany *Budget 2014*, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/293738/budget_2014_distributional_analysis.pdf

Figure 2 show a similar graph with data presented using the first published set of results – those released alongside Spending Review 2010 which showed the distributional impact of the planned spending cuts up to 2014/15 – a horizon that ends one year before that in the most recent results. We reproduce this figure from O'Dea and Preston (2011) for two reasons. First, to make the point that while the figures show qualitatively different results (the second quintile suffer the greatest proportionate loss in the 2010 version), we cannot tell whether this is due to changes in the modelling methodology, to policy changes between 2010 and 2014 or to the extension of the modelling horizon. It would be useful if the UK Treasury, every time it produced a new version of their model, generated a set of results that show the old set of policies modelled using the new methodology. This would allow consumers of the research to assess properly the distributional impact of any new policy announcements.

The second point we wish to make in comparing the two graphs is that after the Spending Review in 2010, the Treasury released sufficient information for us to calculate what proportion of the changes in Departmental Expenditure Limits were being modelled and what were not. This allowed us to augment the 'average' bar in Figure 2 with the impact of unmodelled expenditure. The Treasury no longer releases information that allows us to do this. This is unfortunate. Most of the spending items that are omitted from the model are not included as different households do not use the items differentially[18] (e.g., spending on the environment) or because the extent of differential usage across household income groups either cannot be measured (e.g., for capital expenditure). In neither case is there a good reason to exclude the impact of change in spending on these items in the 'average' bar. Including it here means that the magnitude of the omitted items is clear, even if their distributional impact is not. At the very least, if the Treasury analysis is not to include this, it would be useful if sufficient information were published to allow external researchers to do so (as we previously were able to do).

FIGURE 2 UK Treasury Analysis of Impact of Planned Cuts in Spending on Public Services, Between 2010/11 and 2014/15: Model and Spending Plans as of Spending Review 2010

Note: As Figure 1.
Sources: Data on modelled expenditure is from Chart B6 of Spending Review 2010 (https://www.gov.uk/government/publications/spending-review-2010). Unmodelled expenditure is from authors' calculations using data from Spending Review Annex A and Annex B.

Having undertaken the analysis lying behind the previous two graphs, the UK Treasury add those distributional impacts to those coming from their tax and cash benefit microsimulation model. Figure 3 shows the combined impact of

[18] This is perhaps a questionable rationale for omitting items. We return to this below.

changes to both modelled and in-kind net transfers. The lightest two bar sections here represent the output from the standard tax and benefit microsimulation model. The darker bar segments come from the public service spending module while the black line gives the net effect.

The distributional impact of all modelled changes follows an inverted 'U' shape. Those in bottom quintile and the top quintile lose the most in proportionate terms, with those in the middle of the income distribution losing less. The composition of the distributional impact at either end of the income distribution differs – those at the bottom are losing mainly due to falls in spending on public services while those at the top are losing mainly due to increases in taxation.[19]

FIGURE 3 UK Treasury Analysis of Impact of Cuts in Taxation, Benefits and Spending on Public Services, by 2015-16

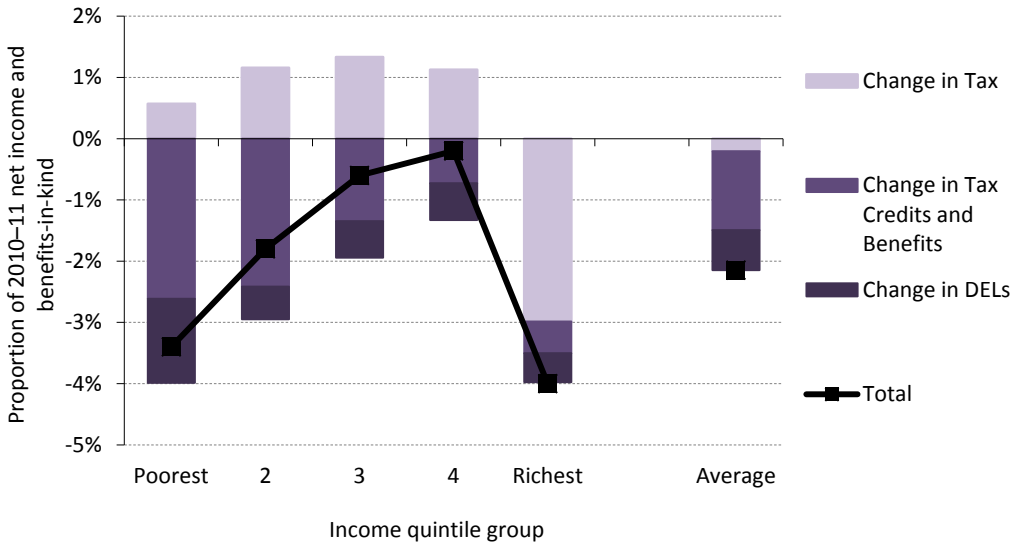

Note: As Figure 1.

Sources: Chart 2I, *UK Budget 2014*: Impact on households: distributional analysis to accompany *Budget 2014*, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/293738/budget_2014_distributional_analysis.pdf

3.1 A questionable Rationale for Omitting Certain Expenditure?

The model of the Treasury is incomplete as the distributional impact of much of government activity is not modelled. In particular, one important restriction is that they only value expenditure where usage of the service varies across the income distribution (though they are clear and transparent about this fact).

[19] This modelled benefit and taxation decisions includes some measures that were announced by the previous government but implemented by the current UK government.

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 which 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 who use the service) is any more reliable for public services which are used to a different extent 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 who use them.

4. CONCLUSION

This paper has outlined the considerable difficulties in establishing precisely the distributional impact of changes to spending on public services as distinct from changes in rates of taxation and cash benefits. 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 is 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 impacts 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.

Despite these difficulties, existing research on public spending can inform an understanding of what effect the planned spending cuts might have. In countries where studies have been undertaken, 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.

This paper started by noting the importance of attempting to assess the distributional impact of changes to spending on public service. In light of the discussion around the work undertaken by the UK Treasury it is natural to ask whether we would recommend the development of a similar model for Ireland. The answer to this is a qualified 'yes'. The 'yes' part of that answer comes from the fact that such a model would unambiguously increase the evidence base on the effect on households of Irish government action. It would facilitate a deeper and more complete understanding of the impact of successive budgets. However, the qualification to our positive answer comes from the fact that developing such a model would be a reasonably large job. Given the inevitable imprecision of any results (coming from the issues discussed in this paper) and the necessary caveats to their interpretation, it might not be the best use of scarce resources. A definitive assessment of whether such a model should be developed in Ireland depends on the opportunity cost of those particular resources.

If such a model was to be developed, we make five specific recommendations relevant to the analysis.

- First, we agree wholeheartedly with the UK Treasury Select Committee's (2010) 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" (paragraph 83). The type of analysis carried out is not one that has 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 and on the publication of same.
- Second, 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 UK Treasury only model the distributional impact of public services where differential usage across households can be ascertained. The case has not been made why their method of valuation is any more reliable for those services which are used differentially than those which are used to the same extent by all residents (such as spending on environmental protection).

- Third, in modelling the distributional impact of the provision of a particular public service (say health), use should be made of what research exists on the matter internationally. Results from other countries will not be exactly and immediately generalisable to Ireland, but, especially where that research relates to developed countries with similar demographics and a similar institutional framework, it will contain lessons that travel.
- Fourth, in the event that a set of results is to be published that differs from the previous set of results in both the range of policies modelled and the underlying methodology, an evaluation of the old policies under the new methodology should also be published. This would allow the user of the research to distinguish what are the effects of the new set of policies and what are the effects of the methodological revisions.
- Fifth, the many issues discussed in this paper 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 government or other authors, should be interpreted with care, and, while informative, should not be considered definitive. Alongside the publication of such results it would be good to see an assessment of how sensitive the headline result is to changes in the underlying assumptions.

REFERENCES

Aaberge, R., A. Langørgen, and P. Lindgren (2010). "The impact of Basic Public Services on the Distribution of Income in European Countries", in A. B. Atkinson, and E. Marlier (eds.), *Income and Living Conditions in Europe*. Eurostat, Luxembourg.

Avram, S., F. Figari, C. Leventi, H. Levy, J. Navicke, M. Matsaganis, E. Militaru, A. Paulus, O. Rastringina and H. Sutherland (2013). "The Distributional Effects of Fiscal Consolidation in Nine Countries," EUROMOD Working Paper No. EM 2/13, Essex: Essex University Institute of Social & Economic Research.

Callan, T. and C. Keane (2009)."Non-cash Benefits and the Distribution of Economic Welfare", *The Economic and Social Review*, Vol. 40, No. 1, pp. 49-71.

Callan, T., C. Keane, Savage, M. and Walsh, J.R. (2013). "Distributional Impact of Tax, Welfare and Public Service Pay Policies: Budget 2014 and Budgets 2009-2014", Special Article in Quarterly Economic Commentary, WInter 2013,

Diamond, P. and J. Hausman, (1994). "Contingent Valuation: Is Some Number better than No Number?", *The Journal of Economic Perspectives*, Vol. 8, No. 4, pp. 45-64.

Gibbons, S. and S. Machin (2003). "Valuing English primary schools", *Journal of Urban Economics,* Vol. 53, pp. 197-219.

Gibbons, S. and S. Machin (2008). "Valuing school quality, better transport and lower crime: evidence from house prices", *Oxford Review of Economic Policy* Vol. 24, pp. 99-119.

Hall, J. and I. Preston (1998). “Tax price effects on attitudes to hypothecated tax increases”, *Journal of Public Economics,* Vol. 75, pp. 417-438.

Hancock, R. and H. Sutherland (eds.) 1992. *Microsimulation Models for Public Policy Analysis: New Frontiers,* STICERD Occasional Paper 17, London: London School of Economics.

Horton, T. and H. Reed (2010). *Where the money goes: How we benefit from public services*, http://www.tuc.org.uk/sites/default/files/extras/wherethemoneygoes.pdf

Joyce, R. and D. Phillips (2013). “Tax and Welfare Reforms Planned for 2013-14”, Chapter 7 in C. Emmerson, P. Johnson and H. Miller (eds.), *The IFS Green Budget* 2013, London: Institute for Fiscal Studies (http://www.ifs.org.uk/budgets/gb2013/GB2013_Ch7.pdf).

Lustig, N., F. Amabile, M. Bucheli, G. Gray Molina, S. Higgins, M. Jaramillo, W. Jimenez Pozo, V. Paz Arauco, C. Pereira, C. Pessino and M. Rossi (2013). "The Impact of Taxes and Social Spending on Inequality in Argentina, Bolivia, Brazil, Mexico, Peru and Uruguay: An Overview," Working Paper 1316, Tulane University, Department of Economics. (http://econ.tulane.edu/RePEc/pdf/tul1313.pdf).

O'Dea, C. and I. Preston (2010). “The distributional impact of public spending in the UK", IFS Working Paper 12/06. London: Institute for Fiscal Studies.

O'Dea, C. and I Preston (2011). “Measuring the distributional impact of public service cuts", Chapter 8 in M. Brewer, C. Emmerson, H. Miller, (eds.), *The IFS Green Budget 2011*, London: Institute for Fiscal Studies (http://www.ifs.org.uk/budgets/gb2011/11chap8.pdf).

Office for National Statistics (2013). *The effects of taxes and benefits on household income, 2011/12*, (http://www.ons.gov.uk/ons/rel/household-income/the-effects-of-taxes-and-benefits-on-household-income/2011-2012/index.html).

Paulus, A., H. Sutherland and P. Tsakloglou (2010). “The Distributional Impact of In Kind Public Benefits in European Countries”, *Journal of Policy Analysis and Management*, Vol. 29, pp. 243-266.

Redmond, G., H. Sutherland and M. Wilson (1998). *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: Cambridge University Press.

Treasury Select Committee (2010). *Spending Review 2010, Sixth Report of Session 2010-11,* http://www.publications.parliament.uk/pa/cm201011/cmselect/cmtreasy/544/544i.pdf

Verbist, G., M. F. Förster and M. Vaalavuo (2012). "The Impact of Publicly Provided Services on the Distribution of Resources: Review of New Results and Methods", *OECD Social, Employment and Migration Working Papers*, No. 130, OECD Publishing. http://www.oecd-ilibrary.org/social-issues-migration-health/the-impact-of-publicly-provided-services-on-the-distribution-of-resources_5k9h363c5szq-en

Volterra Consulting (2009). *The fiscal landscape: Understanding contributions and benefits*, http://www.2020publicservicestrust.org/publications/item.asp?d=1528.

The Economic & Social Research Institute
Whitaker Square
Sir John Rogerson's Quay
Dublin 2, Ireland
+ 353 1 863 2000 www.esri.ie