Grid Edge Policy

Regulation · Energy · Consumers

Ofgem's Approach to Distributional Impacts: A Technical Assessment

30 September 2020

By Maxine Frerk and Joshua Kenway

The Authors

Maxine Frerk is Director Grid Edge Policy and an Associate with Sustainability First. Until 2016 she was Senior Partner Networks at Ofgem responsible, among other things, for network charging. Earlier in her time in Ofgem she led Ofgem's work on vulnerable customers (including commissioning the original CSE archetype work) and was a member of the government's Fuel Poverty Advisory Group.

Contact: maxine.frerk@gridedgepolicy.com

Joshua Kenway has a Masters in International Public Policy from Stanford University

Introduction

Ofgem has recently produced updated Guidance on its approach to Impact Assessments[1] and in particular has set out a new methodology that it intends to use going forward to look at the distributional impacts of its policies. It has also published an updated set of archetypes produced by CSE which Ofgem will use going forwards as part of its assessment.

To inform Sustainability First's response to Ofgem on the Guidance this paper looks in technical detail at the methodology proposed. In particular it explores:

- the features of the ONS Living Costs and Food (LCF) survey used by Ofgem in its distributional impact assessments as compared to the BEIS dataset, used by Ofgem for its Typical Domestic Consumption Value (TDCV) analysis;
- potential issues with the use of average price to convert from expenditure (in the LCF survey) to consumption;
- potential reasons why the results from using the LCF dataset are out of line with conventional wisdom among consumer groups and previous analysis by Grid Edge Policy using the BEIS data;
- the case for the equivalisation adjustments that Ofgem has made (and which are one reason for the apparent difference in results);
- the limited variation in consumption levels between socio-demographic groups in the Ofgem analysis and the need to look at variations within groups, including being clear on the distinction between mean and median consumption;
- the statistical technique used by CSE in creating their updated archetypes and why that has led to a narrower range of consumption values across the archetypes than in their 2012-14 analysis.

Finally, we comment briefly on the distributional impact analysis produced by Ofgem as part of the IA for the RIIO2 Draft Determinations which reinforces a number of the concerns identified with Ofgem's proposed approach in the IA Guidance.

[1] https://www.ofgem.gov.uk/publications-and-updates/impact-assessment-guidance

Choice of dataset - overview

Both the Ofgem and the CSE analysis use the ONS Living Costs and Food Survey (LCF). There is no discussion of why this dataset was chosen over the BEIS NEED data, but it is assumed that it is because of the much wider range of socio-demographic data available. However, there are other important differences between the two data sources in terms of their robustness, scale and coverage as shown in the table below.

	ONS Living Costs and Food Survey	*BEIS sub-national energy statistics feeding into the NEED data base*
Sample size	7000 pa (archetype analysis uses 4 years data so 28000 – unclear if decile analysis uses multiple years)	BEIS data = population (summary tables) NEED = 4.3 million
Geographic	UK wide	BEIS data = GB NEED = England and Wales
Income	Self-reported	Based on Experian data (several years old)
Consumption metric	Energy expenditure as reported in questionnaire response. Asked to consult bill or statement where possible but if not then asked for best estimate of last bill. Converted to consumption using average price.	Annual energy consumption (estimated from meter reads) provided by suppliers as part of BEIS national statistics
Other considerations	Link to wide range of socio-demographic factors	Used as basis of Ofgem TDCV. Link to narrower set of socio-demographic factors in NEED. Will include empty and second homes (but eg near zero consumption properties dropped from TDCV analysis).
Age of data	2014-17	BEIS: 2018 TDCV: 2016-17

Calculation of energy consumption from the ONS energy expenditure figures

The CSE document describing the archetypes states that energy expenditure is converted to consumption using an average price from BEIS taking account of standing charge and payment method.

Clearly getting a robust consumption estimate from the expenditure data (which is itself a reported figure) will depend on having an accurate unit price. There are a number of reasons why the average price may well not be the same across different demographic groups or income deciles.

One obvious factor is payment method which CSE take account of. The archetypes analysis includes data on payment method mix (though figures aren't given in all cases). It is unclear if the income decile analysis also adjusts for payment method mix.

However, what the report does not make any reference to having allowed for is the lower average tariff for Economy 7 customers where off-peak usage is charged at a lower price. Without adjusting for this, the consumption for this group of customers would be understated, which an informed reading of the figures for electric heated homes suggests is likely the case.

Another reason for expecting that customers on lower incomes may be paying a higher unit price for their energy is the lower switching rates that Ofgem find among this group both as a result of lower engagement but also debt blocking. This would tend to mean that they pay a higher rate – unless the tariff cap is considered to have fully addressed this issue. If they do pay a higher price, the analysis will overstate the consumption by lower income households.

Understanding the tariff paid by decile is a crucial insight as Ofgem looks to think about the impacts of removing the tariff cap. Absent any other information a comparison of the ONS and BEIS data could potentially provide some insight on this question.

Different patterns of consumption by income decile from the two datasets

Analysis carried out by Grid Edge Policy on the TCR,[2] which was the subject of a Sustainability First roundtable, made use of BEIS NEED data and showed a very clear link between income and electricity consumption that resonated with the experience of consumer groups over many years. The analysis presented in Figure 1 below shows a difference of a factor of 2 between consumption in the bottom and top income deciles and indicates that consumption increases monotonically across the income deciles, which gives the sense that the relationship is robust.

In contrast, the relationship in the ONS survey data as presented by Ofgem, and replicated in Figure 2, shows a much weaker link. Even for the blue "unadjusted" figures the energy expenditure is shown as being largely flat across the income deciles apart from a higher figure for the top decile. Overall there is only a 30% difference between the bottom and the top deciles. There is some random variation between deciles in the middle of the range which is assumed to just be noise in the data, noting as well that this is just reported expenditure based on the last bill.

Ofgem comment on the fact that the data does not rise monotonically and say "*In particular, there are some low-income households with high energy expenditure relative to their household size. This may be, for example, because they are living in poorly insulated homes or have electric storage heating."* This does not really make sense as we are not looking at individual expenditure but at the average for a decile.

We would have expected Ofgem to want to try to reconcile these different sources of data given the analysis of the NEED data was highlighted by a number of consumer groups in their responses on the TCR, given also that the ONS data is clearly less robust as a source of energy consumption (based as it is on a survey) and given the reliance placed elsewhere in Ofgem on the BEIS data to inform the Typical Domestic Consumption Value (TDCV) analysis.

[2] Available here

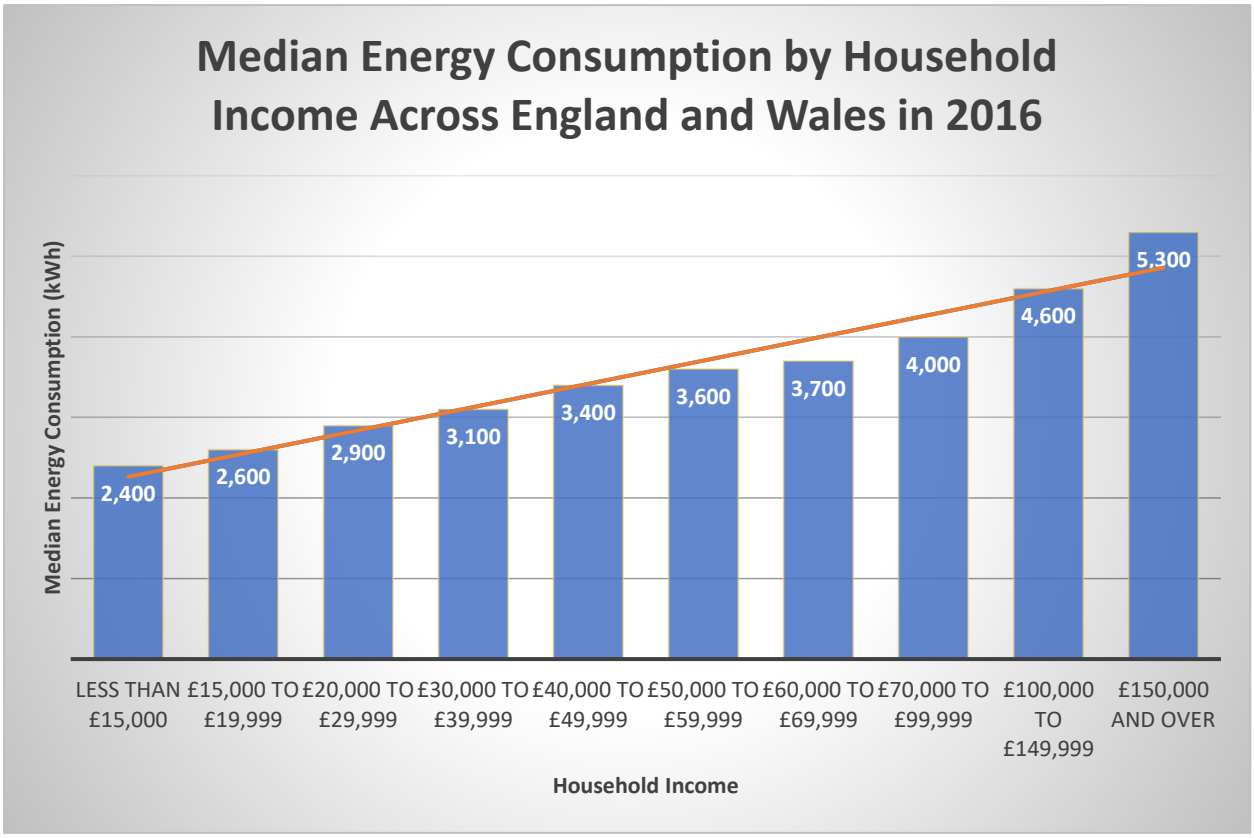

Figure 1: Analysis of electricity consumption by income based on BEIS NEED dataset (source Grid Edge Policy)

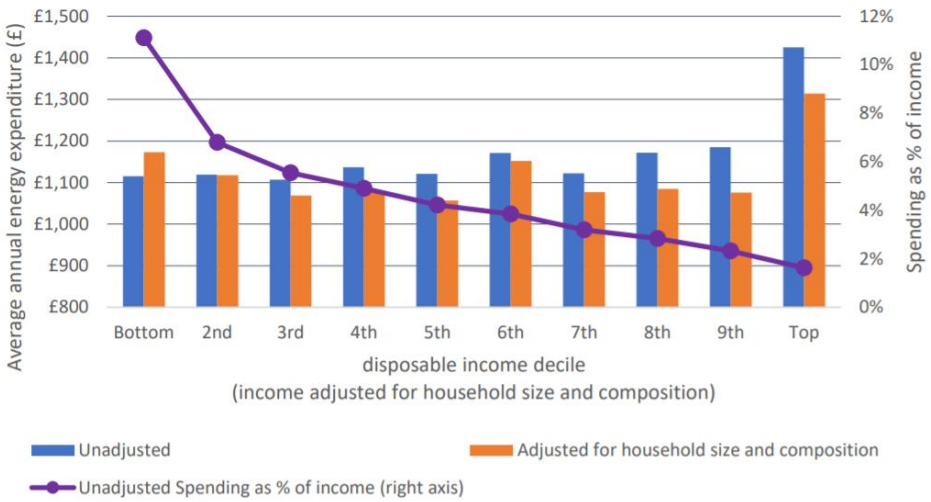

Source: Ofgem analysis of ONS (2020) "Disposable income and energy expenditure for different fuel type households and household types, UK: financial year ending 2018."

Figure 2: Analysis of energy expenditure from Ofgem distributional analysis Annex

Our assumption is that there are two main factors responsible for this difference in results:

1) If, as seems likely, low-income customers pay a higher unit price, then in terms of comparing the two charts one would expect expenditure to be flatter than consumption. However, this reinforces the point that using a common price to convert from expenditure to consumption risks overstating the level of consumption for low income households.
2) The use of equivalised incomes and (shown separately in Figure 2) adjustments to consumption to take account of household size. While the Ofgem paper does not allow one to see how much of an effect income equivalisation has, it is potentially quite significant. Using equivalised income will create a flatter profile.

What is the case for equivalising income and adjusting expenditure?

The BEIS NEED analysis looks at actual consumption by income decile. It is "real". It shows that there are 10% of households with income below £15k who have median consumption at a low level of 2400kWh. It is important to understand the impacts of any policy change on that group, and to be able to describe clearly what the actual bill impact will be on these customers. Currently, Ofgem skips this first "descriptive" step which is important for transparency (and arguably is necessary given Ofgem's duty to have regard to the interests of those on "low incomes").

The point of equivalisation is then effectively to help Ofgem in assessing how far any differential impact is a matter for concern.

The use of equivalised income is a standard approach in looking at poverty-related issues. Smaller households need less income to live on and hence it is reasonable for Ofgem to look at the impacts by equivalised income group. Inevitably that will lead to a flatter profile (as the households with the lowest actual income tend to be smaller and have lower consumption as a result). Once the effect of household size is stripped out through equivalising incomes then the underlying effect of income on consumption will be more muted – although it seems unlikely that this would remove the effect altogether as the Ofgem analysis suggests.

What is much less clear is why Ofgem have also presented figures with an adjusted energy consumption. Indeed, it is unclear whether they are proposing to use the adjusted or unadjusted figures in their analysis[3]. Intuitively if income has been adjusted to reflect the fact that smaller households need less income to cover all their living costs – including energy – then it does not make sense to also scale up the energy consumption for those households in assessing the impacts. This feels like double counting. Ofgem have not provided an adequate explanation for why this is necessary.

Ofgem are possibly applying the adjustment as it is used in fuel poverty calculations, but that is done for a very specific reason linked to the definition of fuel poverty adopted by government of "low income – high cost". To be categorised as being in fuel poverty a household must have both a low income (on an equivalised basis) and have a relatively high level of required energy expenditure (on a normalised basis adjusting for household size).

What is not clear is why the use of this double adjustment is considered relevant to Ofgem's approach to assessing distributional impacts which is a different question. Indeed, the fuel poverty

[3] Looking at the pattern of impacts in Ofgem's Figure 2 for a 10% price change it would appear they are using the unadjusted figures based on the profile across the deciles – but it is not made clear.

methodology[4] is different in a number of other ways, reinforcing the point that a direct read across is not necessarily appropriate. In particular, the fuel poverty definition is based on "required" energy expenditure to maintain a suitably warm home (modelled using BREEM data), not actual energy expenditure – which is known to be lower for those on low incomes as a result of underheating. Moreover, the fuel poverty methodology then makes a further adjustment to the income threshold in that it adds the energy costs to the income for the household in calculating the income threshold to reflect the fact that households with larger energy costs need a greater level of income to meet this greater cost. It might also be noted that there was some considerable debate about the energy cost adjustment in the context of the fuel poverty definition when it was first introduced as acknowledged in the final Hills Review report[5].

Aside from the technical debate, returning to the net effect of these various adjustments as shown in Figure 2 we are left with the counter-intuitive view that income level has no bearing on energy consumption. This sits at odds with the on-the-ground experience of consumer bodies who are all too aware of the problems that many low-income households face in being able to afford their energy bill. It is also out of step with the picture painted by the archetypes where there is a broad correlation between income and consumption. This suggests there is an issue with the methodology being used for the decile analysis.

Of course, Ofgem do, positively, look at impacts as a percentage of income and also apply equity weights to their analysis so that a negative impact on a low-income household is viewed as having a bigger effect pound for pound than on a high-income household. However, it would seem counterintuitive given Ofgem's stated intentions and detailed modelling that equity weights would end up doing all the work and that, for example, increases to standing charges offset by lower unit rates would not apparently have any differential impact on low income households.

Given that this runs counter to the accumulated wisdom of consumer groups it needs a much fuller explanation and validation than Ofgem has provided.

Limited variation between groups and need to understand variation within groups

As shown in the chart below, all but one of the 13 updated CSE archetypes have annual electricity consumption estimates that sit within the TDCV inter-quartile range (ie between the low and high TDCV consumption figures) – despite the fact that 50% of the population has consumption that falls outside that range. What this shows is that despite its attempt to move away from looking at the "average" customer Ofgem has defined the groupings that it looks at in such a way that they still present a highly averaged picture and indeed one that is less differentiated than in the 2012-14 CSE archetypes.

It is also striking that the spread of the archetypes seems to be towards the upper half of the TDCV range. It is not clear why this should be but reinforces the point made above about needing to try to reconcile the data from the different sources.

---

4 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/882233/fuel-poverty-methodology-handbook-2020.pdf

5 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/48297/4662-getting-measure-fuel-pov-final-hills-rpt.pdf

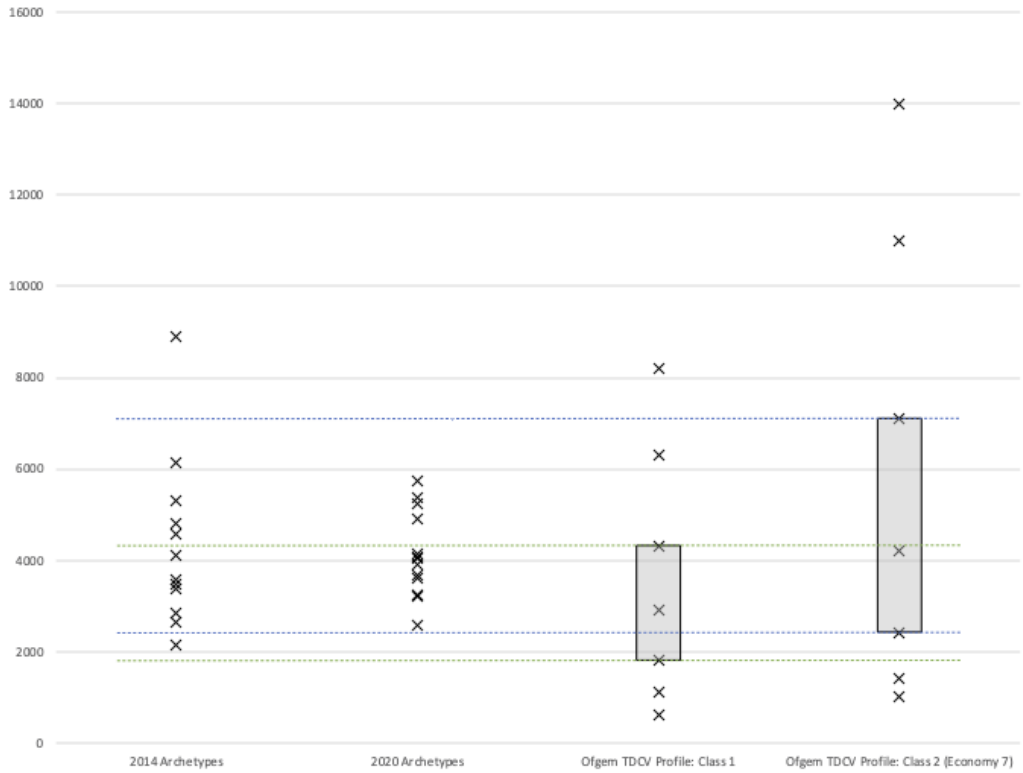

Figure 1: Comparing Electricity Consumption Across Archetypes and TDCV by Profile Class (kWh)
Note: TDCV values for 5th, 10th, 25th, 50th, 75th, 90th and 95th percentiles are shown[6]. Green dashed line indicates Inter-Quartile Range (IQR) for Profile Class 1; blue dashed line indicates IQR for Profile Class 2.

While consumption figures are not provided for the deciles, based on the very flat pattern of expenditure across the deciles it seems clear that at most the top decile will be outside this inter-quartile range.

Hence, Ofgem are still not looking at customers with relatively higher or lower consumption – which is the key to understanding who it is that is impacted by changes in the structure of charges.

As discussed above, this is assumed to be in part because the equivalisation process removes much of that variation between income deciles, which further reinforces the need for Ofgem to be able to clearly articulate the effects of the adjustments it is making.

While there is little variation in consumption across income deciles according to Ofgem's analysis they have nonetheless proceeded to break down the impacts for the different demographic groups within each decile. This provides very little insight as it is hard to detect any real patterns in the data. Simply turning the handle to produce these figures will not help Ofgem make more informed decisions on these difficult questions.

Given the limited variation between groups (on the Ofgem figures) what is also needed is an understanding of the distribution of consumption overall and, ideally, also within each income decile or archetype. One feature of energy consumption data is that the distribution is skewed with a small number of very high users pushing up the mean – hence Ofgem's use of the median in the TDCV.

[6] Taken from Ofgem's latest TDCV review

It is notable that throughout its distributional analysis Ofgem simply refers to "average" consumption. It is assumed that this is the **mean** figure. For electricity for example it is given as 3980 kWh. In contrast the **median** figure given in the TDCV is 2900 kWh for Profile class 1 and 4200 kWh for Profile class 2. In considering the distributional impacts this distinction between mean and median is clearly important.

The other reason that it is important to look at the distribution within groups is that the range is actually very wide which can result in some extreme winners / losers as the structure of charges changes. It is important that Ofgem understands such effects which even if not directly linked to vulnerability may be seen as unfair if the changes are significant and not explained well.

Worked examples

Ofgem provide two worked examples to show how they would apply their methodology to policy decisions. Both raise serious questions about Ofgem's approach.

The first example is one which Ofgem says shows the effect of a uniform % change in price (and hence where equity weights are important). However, the example they cite of a change in the Default Tariff Cap (DTC) is an odd choice. The impact of the DTC will vary across groups depending how many households are in scope in each decile and on consumption level (if there is a shift in the standing charge / unit cost balance). This change will not, as implied in the report, result in a uniform percentage change across income groups. Ofgem has to be able to articulate clearly what the effects of a policy will be on different groups before it can plug the effects into its model.

The second example is a change that Ofgem assumes varies with consumption level. The explanation given is not completely clear but it appears that an assessment of the policy is made for two consumption levels (low and high consumption taken from the TDCV), which is then in turn used to interpolate what the impact would be for the particular consumption level for each decile, archetype, etc. It is striking that this is the one context in which TDCV values are proposed to be used and it is unclear whether the TDCV based interpolation (based on unadjusted consumption) can be used to determine the impacts for different levels of adjusted consumption (if indeed Ofgem are using adjusted consumption in their analysis).

Adopting this approach also assumes that the impacts of the change are linearly related to consumption. While this will be a reasonable assumption in most cases, a step change could apply if, for example, there was a new category of high consuming customers who would be subject to a higher standing charge (a proposal now dropped from the TCR but a useful test case).

Finally, the worked example highlights again the extent to which the process of averaging and the noise within the data leads to results that look highly precise but actually convey very little useful meaning. In the example presented, Ofgem assumes that a customer on the low TDCV would pay £30 less than previously. This means that 25% of households will save more than that. However, in its granular analysis of "winners and losers" by vulnerability and by decile none of the groups save even £10. When we are talking about a saving this is not such an important issue, but if this were an additional cost then Ofgem would be assuming that there was no disproportionate impact on low income, low consumption customers when there could well be.

What the Ofgem analysis does show is how viewing the results as a percentage of income or using equity weights magnifies the impacts on low income households. However this reinforces the point that it is important not just to understand the average impacts on low income households but also

the range of those impacts. Once magnified in this way the impacts on relatively low (or high) consuming low-income customers could be exceptionally harsh.

The methodology used for the CSE archetypes analysis and why that yields a narrower range of consumption

As shown in the chart above the range of electricity consumption across the updated archetypes is noticeably narrower than it was in the 2012-14 analysis. This is assumed to be as a result of a different statistical technique being used this time with no explanation provided as to the reasons for the change in approach.

Different statistical techniques have different strengths and weaknesses and it is helpful to be clear how the choice of methodology will have influenced the results. We have set out below our understanding of the methodologies used.

In broad terms, while the 2012-14 technique was aimed at explaining the variation in consumption patterns (and in particular at understanding the characteristics of low-income high-usage customers), the updated analysis is seemingly focussed on developing archetypes that capture a range of different socio-demographic characteristics as prioritised by Ofgem. The fact that it presents a more averaged picture of consumption should not therefore be a surprise.

More specifically, CSE's 2012 archetype development methodology used CHAID ('Chi-square Automatic Interaction Detection'), a popular analytic technique for performing classification or segmentation analysis. It is an exploratory data analysis method used to study the relationship between a dependent variable [household gas and electricity consumption levels] and a set of predictor [in this case socio-demographic] variables.[7] CHAID modelling selects a set of predictors and their interactions that best predict the variability in the dependent measure. The resulting CHAID model is a classification tree that shows how major 'types' formed from the independent variables differentially predict a criterion or dependent variable. A series of Chi-square tests of statistical independence are used to test for the most relevant characteristics determining variation in the dependent variable.

As noted in the 2012 report, "energy usage is … a key dimension of the difference observed between the segments (and similarities within them) underpinning the archetypes". Thus, this approach allows for an exploration of the key factors explaining energy usage habits among different consumer segments; "energy usage is therefore a key dimension of the difference observed between the segments (and similarities within them) underpinning the archetypes".

CSE pointed to three distinct advantages of this approach in developing archetypes that "capture key different groups of domestic energy consumers". These were:

1. Revealing directly the statistical extent of the differences between different consumer characteristics with respect to energy consumption levels;
2. Providing a basis for an iterative process with Ofgem in making informed choices of meaningful archetypes for analysis and communication purposes;

[7] https://www.cse.org.uk/downloads/reports-and-publications/policy/beyond_average_consumption_2012.pdf

3. Ensuring the archetypes represent groups of households with much in common as consumers of electricity and/or gas (and are hence fundamentally 'energy consumer archetypes').

In contrast, clustering analysis — including Ward's hierarchical clustering, which was used in the development of the new archetypes — does not differentiate between dependent and independent variables. Instead, the core of this algorithm is an objective function (defined by the implementer) that is optimised to group data points into clusters based on an iterative optimisation process. The precise version of this chosen objective function is not described in the report but might reasonably be assumed to be a weighted variance minimisation ($Ward_P$).

CSE indicates that this methodology "allows for different predictive fields to be allocated different weightings to enhance or diminish how significant these are in the clustering." However, no specific justification is offered for why this improves the archetypes relative to their previous characterisations. More importantly, though, there are a number of known weaknesses and caveats to the Ward method that should have been addressed in the methodology section of the report. Some of these are worth exploring in greater depth as they may undermine the presumed preferability over the CHAID approach employed previously.

First, as mentioned, the version of Ward's method used allows for weighting of particular variables over others. Those that are more highly weighted will be more influential in shaping the archetypal clusters than those with a lower weight. Naturally, such weighting will introduce an element of subjectivity into the process and it is therefore problematic that no explanation is offered for these weightings. This is particularly true given that some weighting divergences seem quite unintuitive. For example, gas usage is highly weighted for mains gas consumers, meaning that it will have a large impact in determining the composition of those clusters under that umbrella, while electricity usage is weighted as low. Income is highly weighted, but rurality, age of head of household, and dwelling type, all potentially significant determinants of energy usage, are weighted at a low level, meaning that they have a comparatively low impact in shaping the clusters and consequently the archetypes.

Second, one of the aims of the exercise was seemingly to identify a similar number of archetypes to the 2012-14 model. While this is not an unreasonable goal, a sound implementation of Ward's method requires, to an extent, that the iterative process itself determine the number of clusters. That the desire to achieve a pre-determined number of consumer archetypes led to the "adding in or removing [of] different fields in the underlying data set" is potentially concerning in this context and would benefit from further clarification.

Third, one of few clearly stated motivations for using the Ward model of achieving, for each archetype, "dense clusters of similar sizes (i.e. similar number of households)" is similarly unjustified. The largest archetypal group in the 2012-14 schema was over fifteen times larger than the smallest, while in the updated archetypes that ratio is closer to three. Yet there is little reason to assume that the original difference in archetype group sizes was inherently problematic. Identifying smaller groups of potentially vulnerable customers with atypical consumption patterns could be of real value in understanding distributional impacts. Conversely, there is no reason to think that important distinctions and similarities among consumers would be best represented by dividing the population into similarly sized groupings. Again, the data should shape the clusters with minimal and, where necessary, justified subjective decision-making from CSE and Ofgem.

Fourth, this implementation of Ward's algorithm produces clusters with mean values for highly weighted variables that will skew towards their average. This could, in part, explain the lack of

divergence in consumption across the archetypes relative to the 2012-14 analysis, which is visible in the chart above.

Overall then there are a number of features of the methodology adopted for the updated archetypes which mean that the nature and basis of these archetypes, while superficially similar to the earlier ones, is actually quite different. This should be made clear. It is only the absence of a high usage electric heat archetype that prompted us to explore the change in methodology which will not be apparent to most consumer groups but has implications for how the archetypes should be interpreted.

The RIIO2 Draft Determinations IA – the new methodology in practice

As part of its IA for the RIIO2 Draft Determinations Ofgem has produced a distributional impact assessment in line with its published Guidance. While this is helpful to see as an example of how the methodology is being applied in practice, it reinforces a number of the points made above and also raises some additional concerns with wider ramifications:

- The whole analysis is based around the £20 bill impact figure quoted in Ofgem's overview document. The IA makes clear that this figure is based on the TDCV median consumption but it is hard to see how it has been derived including what assumption has been made as to the split between domestic and non-domestic charges (which is an important aspect of distributional impacts that Ofgem's Guidance does not address). Paragraph 3.2 refers to a LiMO model but the footnote reference does not link to a model of that name. Transparency around the basis of numbers used in the assessment is important.
- This use of the TDCV median consumption demonstrates the problem flagged above that Ofgem is using a different dataset for its distributional impacts than it uses for its TDCV. This is highlighted by the fact that the gas bill saving for the TDCV median consumption is £19 whereas in Table 2.8 the average across all customers based on the distributional impacts model is shown as £23. This potentially reflects, at least in part, the difference between the median and mean consumption but Ofgem does not offer any explanation. What is then striking is that the difference between these two figures is actually greater than the difference between the average customer and any vulnerable group in Table 2.8, raising real questions about the value of the analysis.
- The IA says at 3.6 that the distributional impacts model does not distinguish between fixed charges and variable costs and that they do not know what proportion is fixed and variable in this instance. They then, correctly, note that if there is an element of fixed costs this will lead to an under-estimate of savings at lower consumption levels. However, one would have expected that the interpolation method mentioned in the Guidance would allow account to be taken of fixed and variable charges – and that Ofgem would understand these impacts from its oversight of network charging and from its work on the retail price caps. This is not critical in relation to the RIIO proposals (as my understanding is that gas network charges are essentially variable) but raises questions about the adequacy of the model for exploring changes in the structure of charges – which is where distributional impacts are most likely to arise.
- While the analysis finds very little difference in the impact between vulnerability groups, the effect is – unsurprisingly - much greater when equity weights are used. Similarly, there is a much clearer impact when looked at in terms of % income. However as noted above it does

feel odd that all the work is being done by the equity weightings etc and that the effects of different consumption levels is apparently minimal.

- There is much greater variation in the impacts when looking at the archetypes and this does show that those with higher consumption (and typically higher consumption) benefit more. Indeed, although Ofgem cite archetype D7 as saving £27, the greatest savings are actually for A2 (a high-income group) at £34. This then paints a rather different picture from that presented by Ofgem's decile and vulnerability analysis. Ofgem would do well to reflect on the reasons for that difference. Given that the archetypes themselves do not actually reflect the full range of consumption levels as shown in the Figure above, this reinforces our concerns about the very flat pattern of demand that Ofgem is assuming across the deciles.
- While Ofgem has cranked the handle on the distributional impacts in line with its published Guidance it is very hard to see what conclusion or insight one is expected to draw from this analysis. The fact that Ofgem has carried out the full modelling exercise to show how the £1 saving on transmission varies by group suggests that it has lost sight of the wood for the trees. The one significant insight perhaps stems from the fact that £19 of the headline £20 bill saving comes from gas distribution. As Ofgem rather coyly says this means that off-gas customers are "the groups that will overall experience the least savings". If nothing else Ofgem should perhaps make this point clearer in its overview of the impacts.
- While some attempt is made to explain equity weights and it is helpful to see the figures on a raw and equity weighted basis, the use of "equivalised disposable income deciles" in Figures 1 and 2 is not explained. Even linking back to the IA Guidance it is unclear (as noted above) whether Ofgem is using the adjusted or unadjusted expenditure figures.

While the distributional impacts between different groups of domestic customers is not a particularly important element of the RIIO proposals, this example is important in highlighting that Ofgem still has some way to go in making its distributional impact assessments transparent and in road testing the numbers that it is using which we maintain underplay the variability of impacts as a result of looking at what are still highly averaged consumption figures. Triangulating the results with data from the BEIS dataset and the TDCV would help Ofgem and others understand what is going on where results are counter-intuitive.

For the RIIO Draft Determinations it is also worth noting that there are elements of distributional impacts that are more relevant but which are not addressed in Ofgem's framework:

- The analysis presents a static picture based on the £20 bill saving and no acknowledgment is made of the way that figure could vary depending on the use of uncertainty mechanisms through the period;
- No data is given on how the impacts vary by geographical region (which may be important given different company plans);
- No framework is provided for thinking about inter-generational impacts despite Ofgem noting that both its move to CPIH and its change to depreciation on gas transmission have the effect essentially of moving money between price control periods. This is a complex but important issue in the context of the energy transition.