Investigating the potential impacts of Time of Use (TOU) tariffs on domestic electricity customers

Report to Ofgem

Joshua Thumim
Head of Research & Analysis
Centre for Sustainable Energy
3 St Peter's Court
Bedminster Parade
Bristol BS3 4AQ

t. 0117 934 1439
f. 0117 934 1410

josh.thumim@cse.org.uk
www.cse.org.uk

Registered charity no. 298740
Registered company no. 2219673

Centre for Sustainable Energy, March 2014

CONTENTS

1 Introduction

The rollout of smart meters will enable electricity suppliers to deploy new time-of-use (TOU) based approaches to pricing electricity. These have the potential to stimulate customers to shift demand from peak to off-peak periods or simply to reduce peak demand, in both cases reducing system costs and improving system efficiency.

However, there will be variations in the ability of different customers to take advantage of TOU tariffs, depending on their current pattern of electricity usage and their willingness or capacity to respond. Understanding the scale and nature of these variations will be important both to the design of effective TOU tariffs and to considerations of how consumer interests may be protected as such tariffs become available.

To support its initial consideration of the potential impacts of TOU tariffs on domestic electricity consumers in the UK, Ofgem commissioned the Centre for Sustainable Energy (CSE) to undertake analysis and modelling of temporal domestic electricity use patterns.

Given (a) the limited availability of data which combine domestic TOU electricity consumption patterns with consumer socio-demographic characteristics, and (b) the limited evidence of how different consumers might respond to TOU tariffs (by changing their demand patterns), this analysis is inevitably constrained in its scope. It has therefore been focused primarily on developing analytical and modelling techniques which have the potential to be applied subsequently to more comprehensive data and evidence as they become available.

The analysis uses the half-hourly smart electricity demand dataset collected during the Energy Demand Research Project (EDRP). The EDRP was an approximately two-year trial of domestic smart meters and energy demand reduction interventions which ran from 2008-2010. It was part-funded by DECC, managed by Ofgem, and run by four energy suppliers. It involved approximately 60,000 households, of which some 15,000 had smart electricity meters installed, with just over half also having smart gas meters. CSE supported Ofgem in its evaluation of the EDRP, and managed the collection of the raw smart meter data from the trial.

This document presents the results of CSE's analysis and modelling using the EDRP data. Although this is a rich dataset, it does suffer from certain drawbacks: the dataset is not statistically representative of UK domestic electricity consumers, and detailed sociodemographic data were not available for the purposes of the work presented here.

The project had two main aims:

1. **Cluster analysis**

   Clustering techniques (essentially grouping cases together based on similar electricity demand patterns) were used to investigate the potential for subdividing Elexon Profile Class 1 - Domestic Unrestricted (PC1) customers into a number of clusters of similar usage patterns within the single PC1 profile which might provide additional useful insight into typical domestic electricity consumption patterns. To achieve this we tested various approaches to clustering the EDRP cases, all using a standard data clustering technique known as K-Means.

2 TOU Tariff simulation and analysis

A tariff model was created to calculate the electricity bills of the EDRP cases based on a set of three TOU Tariffs defined by Ofgem, and to compare this with a representative standard non-TOU tariff. The aims of this phase of the work were to:

- analyse the distribution of differences between the annual bill under (a) the TOU tariff, and (b) a standard flat-rate tariff specified by Ofgem;
- examine the average daily profiles of households gaining or losing significantly under the TOU tariffs.
- Pass the new “profiles” identified in the cluster analysis through the tariff model and calculate the resulting impacts on the bills of households with each profile.

It should be noted that this analysis is limited by the fact that it did not attempt to model household behavioural responses to the tariffs, which was beyond the scope of this initial study. Any distributional impacts revealed in the analysis are therefore those that would occur if consumers changed nothing in response to a TOU tariff.

2 Data inputs

Before the EDRP data could be used for TOU modelling and cluster analysis a number of preparatory steps were required, to maximise the relevance and quality of the data used.

a) Removal all cases which were exposed to time-of-use trial interventions in the EDRP, including any Profile Class 2 (Economy 7) customers, to reduce risk of confounding results.

b) Identification of duplicate meter advances (same case, same date, same time, same advance value) and retention of only one from each set, to maximise quality of demand data, and reduce risk of contaminating results through redundant inputs.

c) Identification of multiple meter advances with different values (same case, same date, same time, different advance value) and removal of these advances from the dataset, to maximise data quality.

d) Identification of meter advances with null or zero values and removal of these from the dataset: on the assumption that zero/null electricity demand values can only be a result of bad data.

e) Selection of the best single continuous year across the 2008-2010 period of the EDRP. That is, the 365-day period with the largest number of valid meter advances. This would define the data to be used in both the cluster analysis and the tariff modelling, and was found to be the year commencing at midnight on 20 April 2009. During this period 4,639 cases had electricity advances which survived the filters applied in (a) to (d). On average each case had 17,366 advances where 100% coverage would require 17,520 advances. This equates to 99.1% valid data for the 4,639 cases over the selected period. Figure 1 below illustrates the distribution of data over the selected year. This was done to maximise the sample size used in the subsequent analysis.

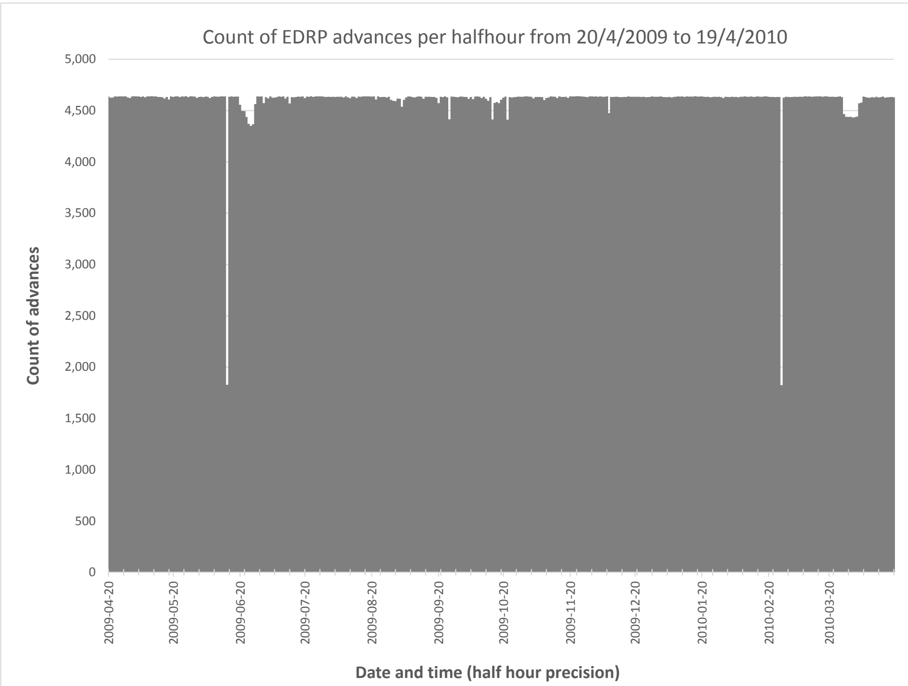

Figure 1: data covering in optimal EDRP year

3 Cluster analysis

3.1 Objectives

The objective here was to use cluster analysis to search for a set of "archetypal" demand profiles within the selected year of EDRP smart electricity meter data. The hypothesis being that PC1 could be divided into a set of sub-profiles using the k-means classification algorithm (see below for an explanation of this).

3.2 Method

A standard clustering algorithm known as k-means was used in this analysis. This approach groups input observations (here: EDRP cases) into a user-specified number (= k) of separate clusters (essentially, groups of cases) using an iterative process:

1. Randomly assign the locations of the centres of the k clusters (known as the cluster "centroids") in the data space. For example if each observation is described using two numbers, this would be 2 dimensional space and could be visualised using a scatterplot. In our case each observation is described using up to 17,520 numbers (i.e. a year's worth of half-hours) so the data cannot be conveniently visualised this way. However the process is the same.
2. Assign each observation to the centroid it is closest to in the data space.
3. Now move each centroid to the position which is the average (the mean) of the locations of the members assigned to it in (2).
4. Repeat steps 2 and 3 until the memberships of the clusters stop changing (aka convergence: in practice the stopping criterion might be a minimum marginal reduction in distortion – see below)

Note that this algorithm requires the user to specify two things: (a) the number of clusters required, and (b) the method for calculating the distance between any pair of observations. The optimal number of clusters cannot be known at the beginning (this is a limitation of k-means), so we ran the analyses with values of k from 1 to 15 (see below). The method we specified for calculating the distance between two observations was the straight-line distance between them. That is, the square root of the sum of the squares of the differences between each pair of values for the two observations being compared: $\sqrt{(difference\ at\ halfhour\ 1)^2+\ldots+(difference\ at\ halfhour\ 17520)^2}$. In effect this is Pythagoras' formula for the length of the hypotenuse in >2 dimensions.

K-means can converge on a suboptimal solution (i.e. a local rather than a global optimum), so most implementations of the algorithm iterate the entire process multiple times before choosing the best result. This includes the implementation we used[1].

We experimented with a number of different input permutations before identifying a preferred solution. This involved using different combinations of:

- absolute consumption values for each half hour in the timeseries
- normalised consumption values for each half hour in the timeseries
- taking the entire 17,520 half-hour timeseries as the description for each case
- taking the average week (336 half-hour timeseries) as the description for each case
- taking the average day (48 half-hour timeseries) as the description for each case

[1] SciPy - http://docs.scipy.org/doc/scipy/reference/cluster.vq.html

There are various ways of selecting the best value of k. The objective is to maximise the explanatory power of the resulting classification while minimising the number of clusters. At one extreme there is a single cluster (no explanatory power), and at the other there is a cluster for every observation (no classification value). For each experiment we ran k-means with all values of k from 1 to 15, analysing the results based on the distortion[2] value, as well as by inspecting visualisations of the average daily profile for each centroid. We also discarded centroids with very small populations on the assumption that these contained outlier observations.

Following this analysis we selected the following permutation as the preferred solution:

- Normalised consumption values (so that the shape of the demand profile influenced cluster membership, but the size did not: this is consistent with the way profile classes are used in the electricity industry generally).
- Using the average week in the year commencing midnight on 20 April 2009.
- K-means clustering using the above parameters, with k=3: that is, we forced the clustering process to output 3 new groupings of cases.

3.3 Results

3.3.1 Summary of clustering results

The following section presents the results of the clustering using the inputs summarised above.

Figures 2 and 3 below show the average week for PC1 and the EDRP data sample respectively. It is interesting to note how similar these are – the main difference appears on first inspection to be that the EDRP sample shows a smoother curve, possibly because of the larger sample size of just under 5,000. In addition the average consumption of the 4,760 selected EDRP cases over the chosen year was 4,353 kWh, versus a figure of 4,208 kWh for the standard Class 1 Profile provided by Ofgem. This similarity suggests (but does not prove) that the sample of EDRP cases used in for this project may in fact be reasonably typical households in terms of their temporal electricity use patterns. That said, there are differences between the EDRP sample and PC1 which reinforce the caveats outlined here that, unlike PC1, the EDRP is not a representative sample of GB electricity consumers.

[2] Distortion is a measure of the overall quality of each clustering result, defined as the sum of the squared distances between all observations and their centroids. Higher values indicate more dispersed clusters, lower values indicates tighter clusters. The units of distortion are distance squared, where distance is as defined in Section 3.2.

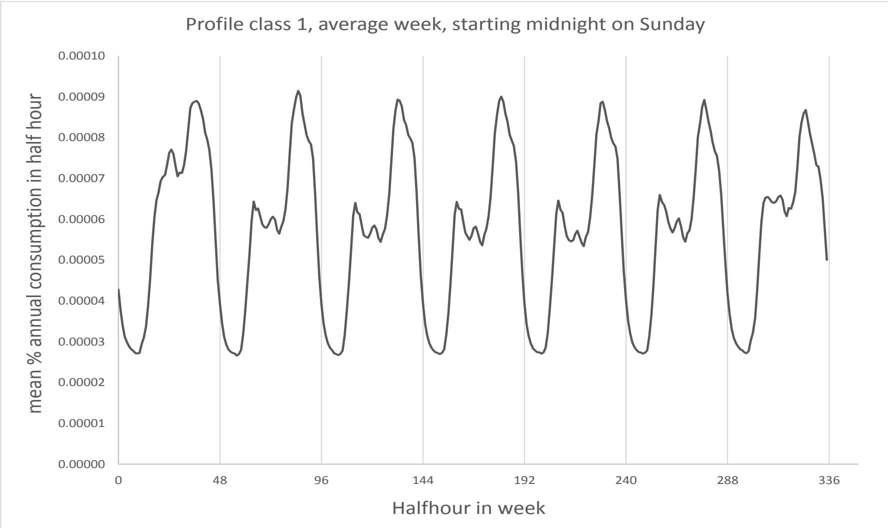

Figure 2: PC1 average week

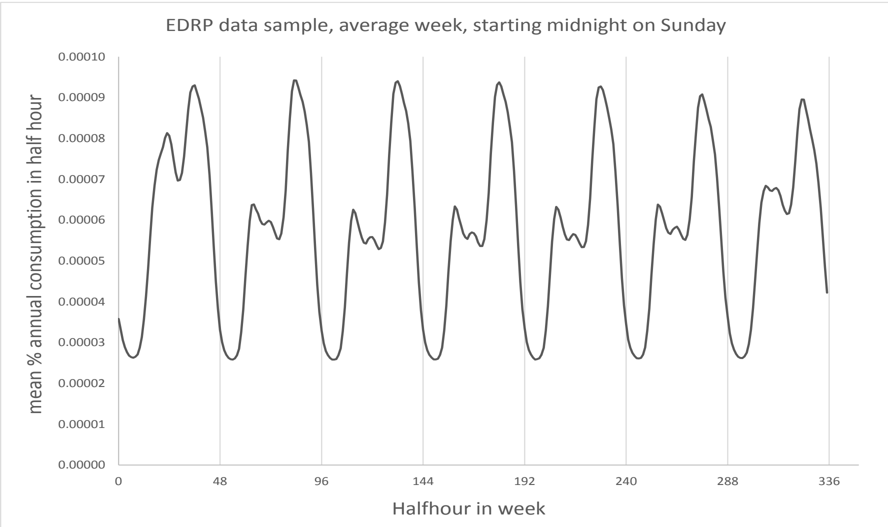

Figure 3: EDRP sample average week

Figure 4 below repeats this comparison, this time overlaying the average day for the PC1 timeseries with that of the EDRP sample used in the analysis. The grey shaded areas show the percentage difference between the curves – this is largest where the consumption values are smallest, an artefact of comparing percentages. In Figure 4 a further subtle but important difference is also visible: the evening peak for the average day from EDRP data sample is slightly earlier, and slightly higher, than that of the PC1 timeseries. This leads on average to a slightly different interaction with the TOU tariffs (see Section 4).

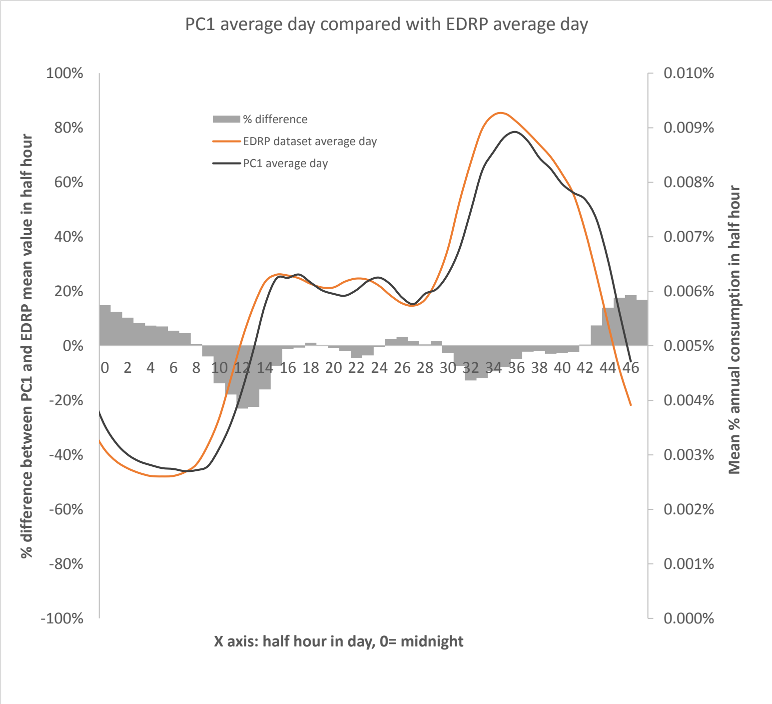

Figure 4: Average days overlaid - PC1 and EDRP sample

Figure 5 shows the total cluster distortion against the number of clusters. As discussed earlier, the total distortion falls with the number of clusters, but beyond a certain point the increase in number of clusters starts to deliver lower reductions in distortion.

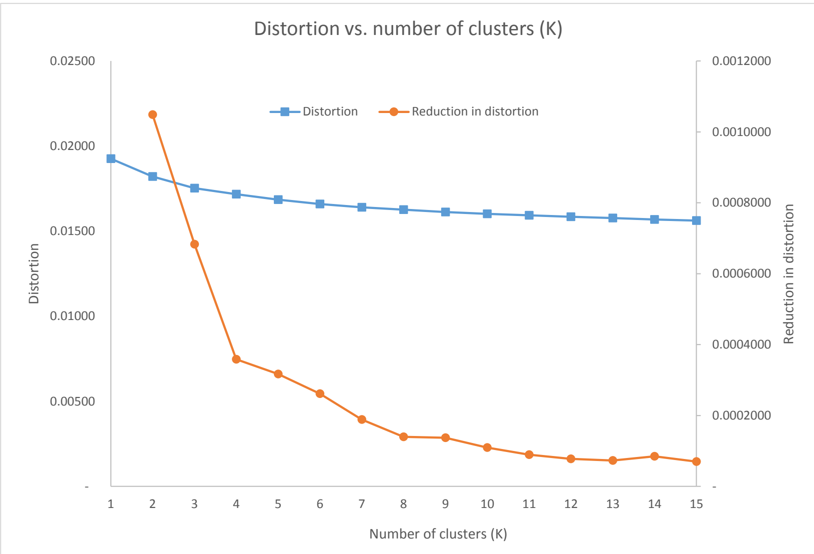

Figure 5: Distortion versus number of clusters

Figure 6 shows how the cluster centroids evolve from K=1 to K=3. The centroids represent the average normalised week over a year calculated over the membership of each cluster. That is: for every case, each half hour's consumption is divided into the annual total for that case. The result is a percentage value, the "normalised" consumption. These sum to 1 (or 100%) for each case.

Note that as the dataset is divided into more clusters, the population of each cluster drops, as does the overall distortion value. The K=3 results were chosen for further analysis based on the inflection in Figure 5 above: we identified this as the point beyond which there would be diminishing returns from increasing the number of clusters.

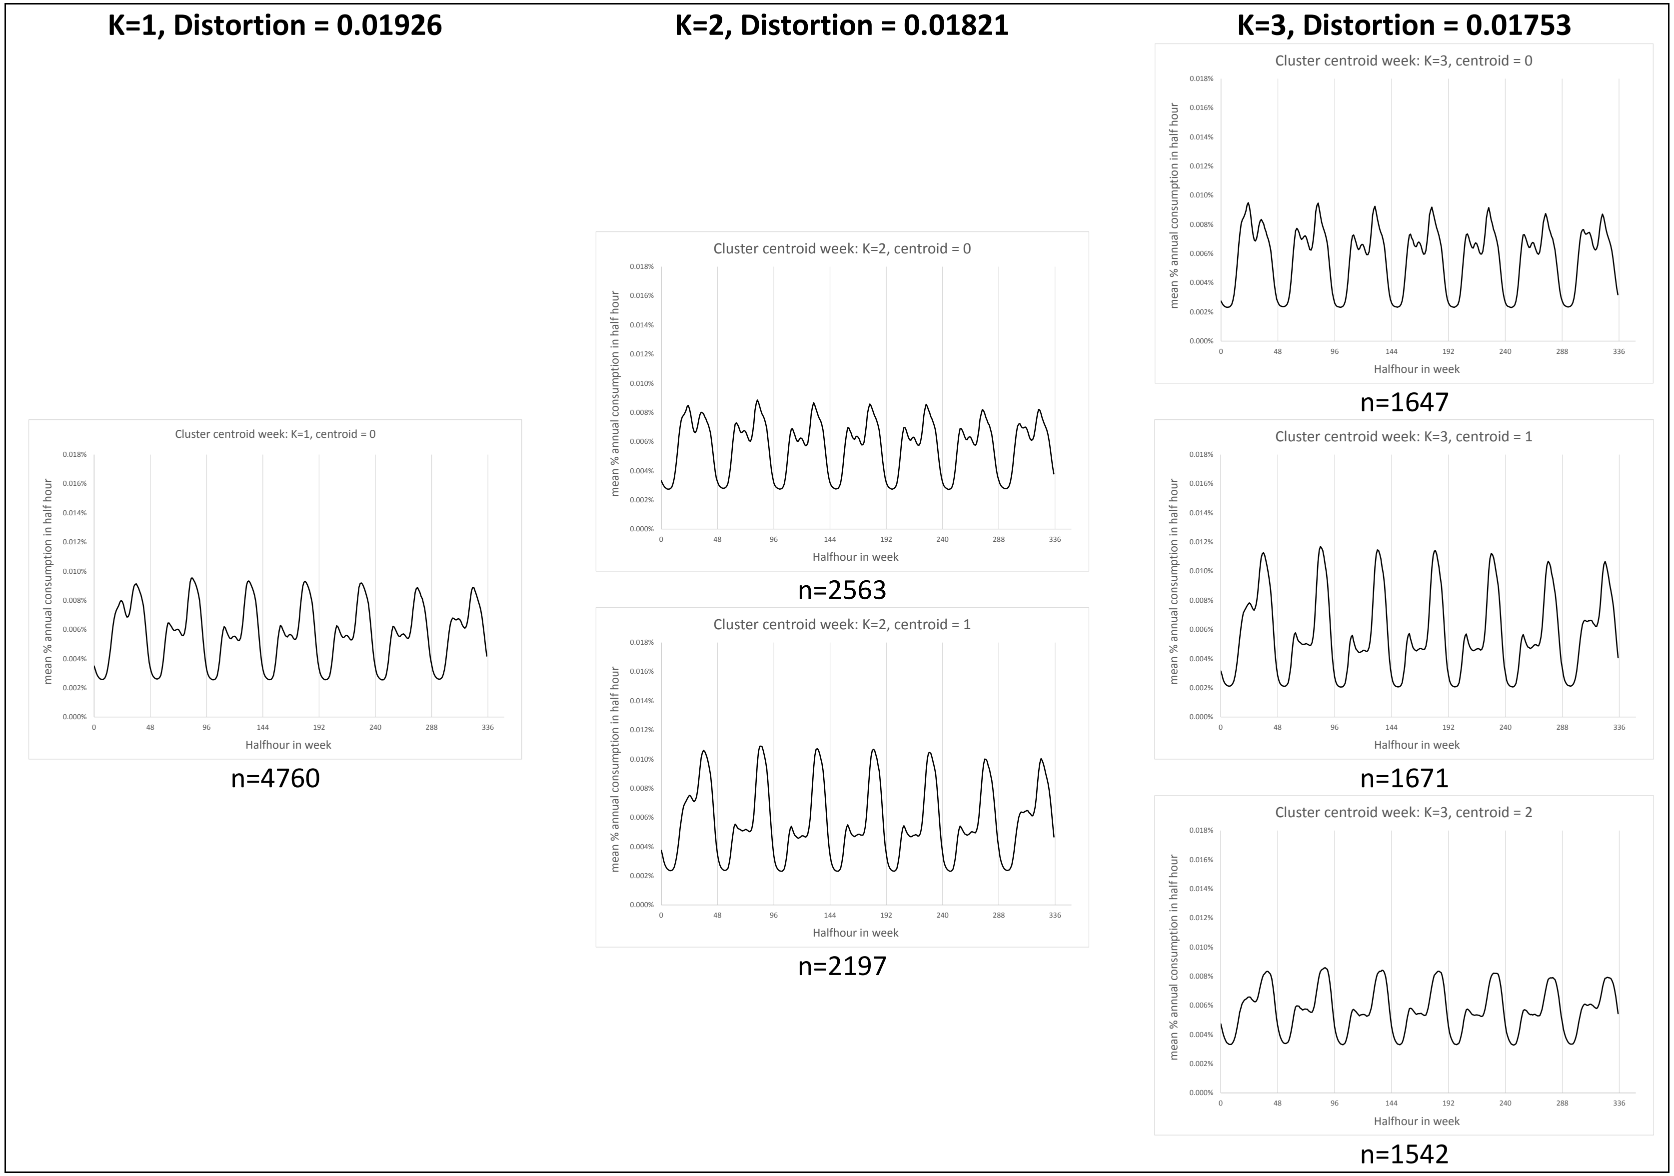

Figure 6: average week for cluster centroids from K=1 to K=3

3.3.2 Analysis of clustering results (K=1,2,3)

Figure 6 above illustrates the way the cluster centroids change as the number of clusters increases. From left to right the overall sample is split into 1, 2 and 3 clusters. As expected, the average shapes of these clusters changes with their membership. For a single cluster (K=1) we are simply looking at the average week (starting at midnight on Sunday) for the whole sample of 4760 cases. This is therefore identical to Figure 3, and very similar to the standard PC1 curve.

Results for a single cluster (K=1)

The following summaries should be read with reference to the k=1 average week (the leftmost graph in Figure 6).

K=1 Weekdays

The average weekday for the single cluster result shows a morning peak at around 0800. Demand then falls until around 1500, with the exception of a small peak at noon. The largest peak of the day occurs around 1800, and demand then falls until around 0430 the next morning, when it starts increasing towards the morning peak at 0830. As discussed, this is very similar to PC1.

K=1 Saturdays

The average Saturday for the single cluster is similar to the average weekday, except that the morning peak is sustained until noon. Demand then drops until around 1500, subsequently climbing steadily to an evening peak at around 1800. Again this is similar to PC1.

K=1 Sundays

As with PC1, the average Sunday for the single cluster result is significantly different from either the average Saturday or the average weekday. The first peak is significantly higher, and later, than that seen on Saturdays or weekdays. Demand peaks once at around noon, falls until around 1430, and then climbs towards the evening peak which occurs around 1730 - slightly earlier than on Saturdays or weekdays.

Results for two clusters (K=2)

The following summary should be read with reference to the middle pair of graphs in Figure 6.

The results for K=2 are included so show how the cluster profiles evolve as the overall sample is split into further clusters. Here we can see that a group of cases with a high evening peak has been grouped together (see centroid 1). We also see that the relationship between weekdays, Saturdays and Sundays observed for K=1 (and for PC1) holds, even though the profile shape is very different: that is, weekdays and Saturdays are similar, and Sundays are very different. This is not surprising if we consider typical social and economic behaviour patterns across a week.

Results for three clusters (K=3)

The following summaries should be read with reference to Figures 7-11 below.

The 3-cluster results have been selected as the main output from the clustering process described here. There are three distinct patterns associated with each of the three clusters (centroids 0,1 and 2). For each centroid the between-day relationships are very similar to those described above. Weekdays have a morning peak, demand drops until a smaller lunchtime peak, and then rises again towards an evening peak. Saturdays are similar to weekdays, except that the first peak is largely sustained until the midday peak occurs. Sundays again consist of two peaks rather than three, the first of which is later than the weekday morning peak, and the second of which is earlier than the weekday evening peak.

Despite this similarity, the patterns differ for each cluster centroid, and this has a direct bearing on the way they interact with the time of use tariffs (see Section 4 for more detail on this). The different patterns shown below reflect in aggregate the different usage patterns of the members of each cluster. It is worth noting that each centroid represents the average week across a whole year, from over 1500 households: this means that they are an aggregation of around (52 *1500) = 78,000 different household-weeks. This means that the average patterns may conceal more detailed differences: without more detailed sociodemographic information for each case it is therefore probably not worth speculating about the specific drivers behind each of the centroid profiles (although this is tempting).

Figure 7 shows the average Sunday, weekday and Saturday patterns for the first cluster (Centroid 0) in the k=3 results set.

Centroid 0 summary: fairly even demand over the waking day with similar sized morning and evening peaks, and little difference between weekdays and weekends. As discussed above, Sundays are different in having two distinct peaks, with the morning peak higher and later than that seen on the other days.

Sunday **Weekday** **Saturday**

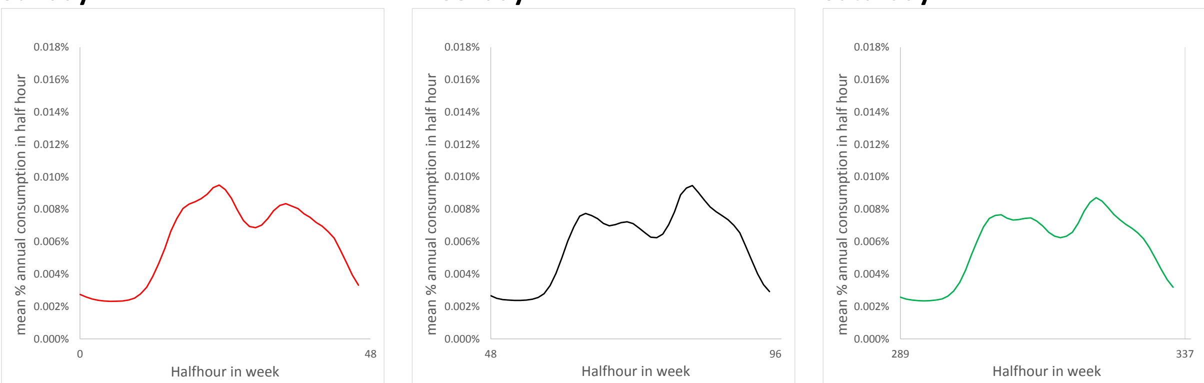

Figure 7: Centroid 0 for K=3

Figure 8 extracts only the sizes and times of the peak demands for the average weekdays, Saturdays and Sundays, along with the minimum load over the average week. This view reveals the detailed differences in the sizes and timings of the peaks: for example the morning peaks are 0800 (weekdays), 0900 (Saturdays) and 1130 (Sundays). Figure 8 also shows that the peaks are similar in size.

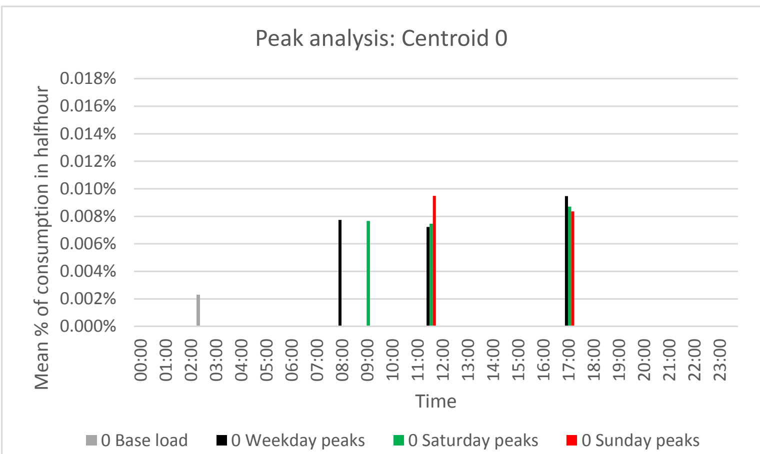

Weekday peaks:
0800, 1130, 1700
Saturday peaks:
0900, 1130, 1700
Sunday peaks:
1130, 1700
Peak:base = 4.1

Figure 8: Peak analysis for Centroid 0

Centroid 1: as Figure 9 shows, centroid 1 differs markedly from centroid 0. It has much lower morning peaks, much higher evening peaks, and relatively low base load. Consequently the peak:base ratio is high, with a value of 5.7:1.

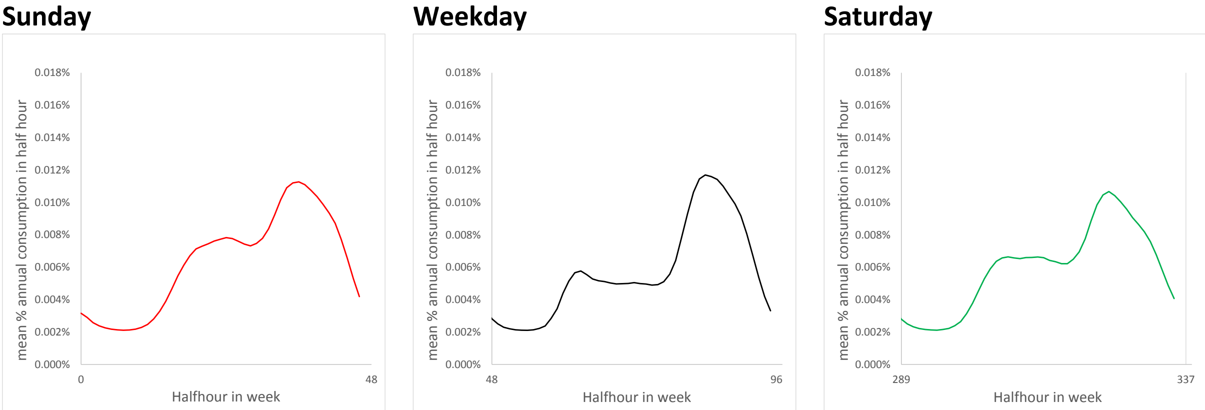

Figure 9: Centroid 1 for K=3

Figure 10 summarise the peaks for centroid 1. Relative to centroid 0 the evening peak is an hour later at 1800, and the weekday morning peaks an hour earlier at 0700. There is also a noticeable difference in peak timings for weekday mornings and weekend mornings, while the evenings are much more similar.

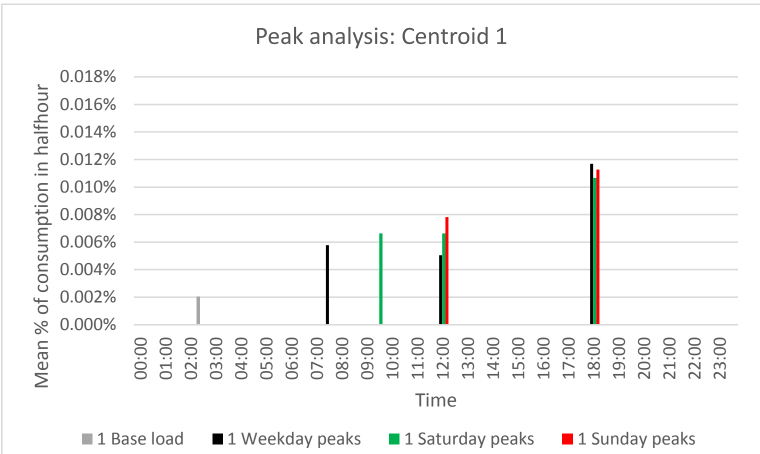

Weekday peaks:
0700, 1200, 1800
Saturday peaks:
0930, 1200, 1800
Sunday peaks:
1200, 1800
Peak:base = 5.7

Figure 10: Peak analysis for Centroid 1

Centroid 2: Figure 11 shows that centroid 2 is a very different profile from the previous two. The peak:base ratio shown in Figure 12 is far lower at 2.6:1, and this is reflected in the much flatter shape of the profile. Nevertheless the same basic pattern obtains: morning and evening peaks on all days, and (in this case very slight) midday peaks. In addition, for centroid 2 there is very little difference between weekdays and Saturdays, while Sundays stand slightly apart.

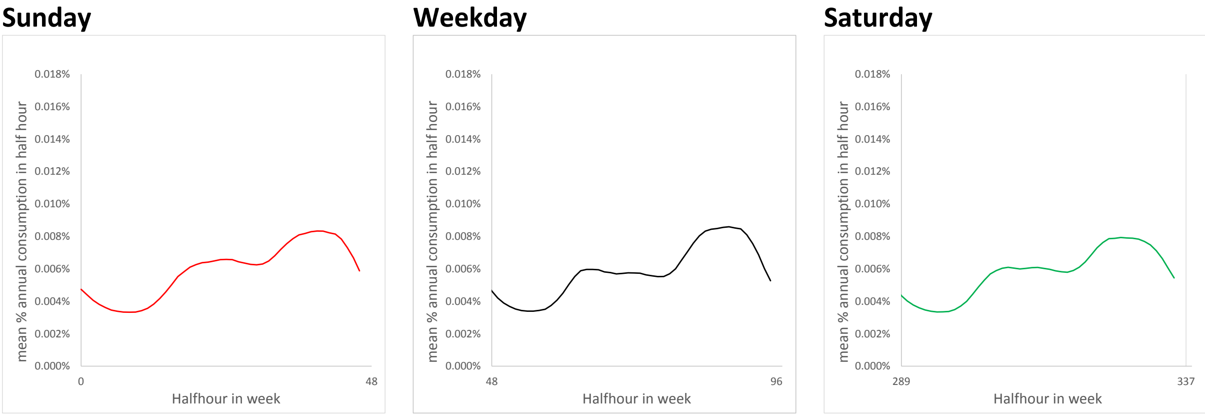

Figure 11: Centroid 2 for K=3

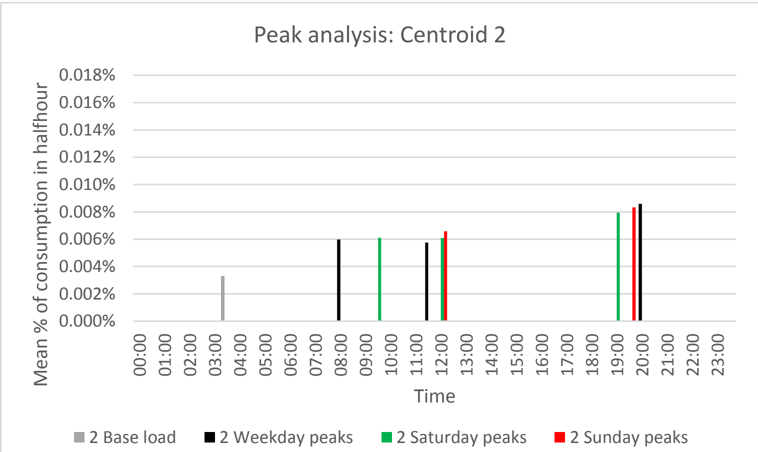

Figure 12: Peak analysis for Centroid 1

Weekday peaks:
0800, 1130, 2000
Saturday peaks:
0930, 1200, 1900
Sunday peaks:
1200, 1930
Peak:base = 2.6

3.3.3 Interpretation

The main conclusion to be drawn from the results presented above are that cluster analysis using a fairly straightforward algorithm and distance metric may have successfully identified three new profile classes "hidden" within PC1. This is an interesting finding. The clusters identified can be summarised as:

(0) similar to PC1, albeit with a higher Sunday AM peak
(1) high variance, with very high evening peaks and very low base load
(2) low variance, with the typical diurnal patterns present at much lower amplitudes than for (0) and (1).

Further work could be done to verify and improve these results. This might include testing other clustering algorithms, and certainly other distance metrics more specifically tailored to timeseries/waveform data. It would also be valuable to repeat the analysis on other smart meter datasets to check whether the above clusters are generally present, or merely artefacts of the EDRP dataset used here.

3.3.4 Caveats

It is worth emphasising at this point that PC2 (Economy 7 tariff) households were specifically excluded from the dataset from which the clusters were produced, as were all households which experienced time-of-use tariff based interventions. It seems likely that their inclusion would have led to further clusters with different characteristics from those produced here. In addition, there may well be other perhaps less common demand patterns with very particular characteristics - for example those associated with night shift working patterns, multiple occupancy, and university accommodation. It is also probable that the EDRP dataset, given its relatively small size and somewhat ad-hoc sampling strategy, does not offer a representation of all the significant patterns of demand which would be found in a larger sample of halfhourly electricity consumption data.

4 Tariff modelling

4.1 Objectives

The tariff modelling presented in Section 4 formed a major component of the project. The aim was investigate potential financial impacts of certain TOU tariffs on domestic electricity consumers. Financial impact being defined as the difference in the annual electricity bill under a notional TOU tariff, when compared with a baseline flat-rate tariff. The detailed objectives were to:

- model, illustrate and analyse the distribution of change in household annual electricity cost between a standard flat-rate tariff, and three illustrative time-of-use tariffs.
- develop Ofgem's understanding of the types of electricity consumption pattern which would be likely to lead to gain or loss under the TOU tariffs.
- investigate in more detail the demand patterns of cases experiencing significant gain or loss (the top 5% of the dataset on each criterion).

4.2 Method

The TOU tariffs specified by Ofgem (see Section 4.3) were implemented as a scripted database application running over all EDRP cases surviving the filtering described in Section 2. The output of this is a table[3] showing the baseline electricity bill for each case, and the difference in bill compared to that baseline for each of the three time of use tariffs. This table includes the cluster centroids which were also run through the tariff model. The timeseries for the centroids were created by cloning the 336 half-hour (1 week) sequence 52 times to create the 17520 half-hour (1 year) sequence required for the tariff model – that is, the year of electricity demand fed into the tariff model consisted of 52 identical weeks with normalised values summing to 1, scaled to the PC1 average annual consumption of 4207.9 kWh.

4.2.1 The tariffs

Ofgem designed three TOU tariffs to be revenue neutral for the standard PC1 demand curve using goal seeking in MS Excel, providing the example static tariff structures and levels for the retail tariff model. Their methodology is summarised below. To note, the scope of this research is limited to static tariffs only. This is because static tariffs are likely to come to market before dynamic tariffs. Examining the impact of dynamic tariffs will be a natural extension to this research as more dynamic tariff trials are undertaken and greater volumes smart metering data become available.

Tariff structure

There are many possible TOU tariff structures that could be developed to meet the specific needs of suppliers and their customers. For this research exercise Ofgem focused on a likely evening peak structure which is reflected in the daily pattern of wholesale prices. Figures 13 and 14 illustrate the tariffs, and their relationship with the spot market price for electricity. Figure 13 shows the variable component of all three tariffs in pence/kWh, alongside the spot price for electricity in £/MWh. Figure 14 shows the variable component of Tariff 3 alongside the spot price for electricity, with both values in units of £/MWh. It is evident that the peak/ off-peak ratio within the example tariff is noticeably greater than the the equivalent ratio in wholesale prices. It is worth reiterating that the tariffs developed in this research are purely illustrative. TOU tariffs developed in the future will need to be carefully designed to strike the balance between wholesale cost-reflectivity and providing strong enough price signals to consumers to encourage behaviour change.

[3] Provided separately as a .csv text file

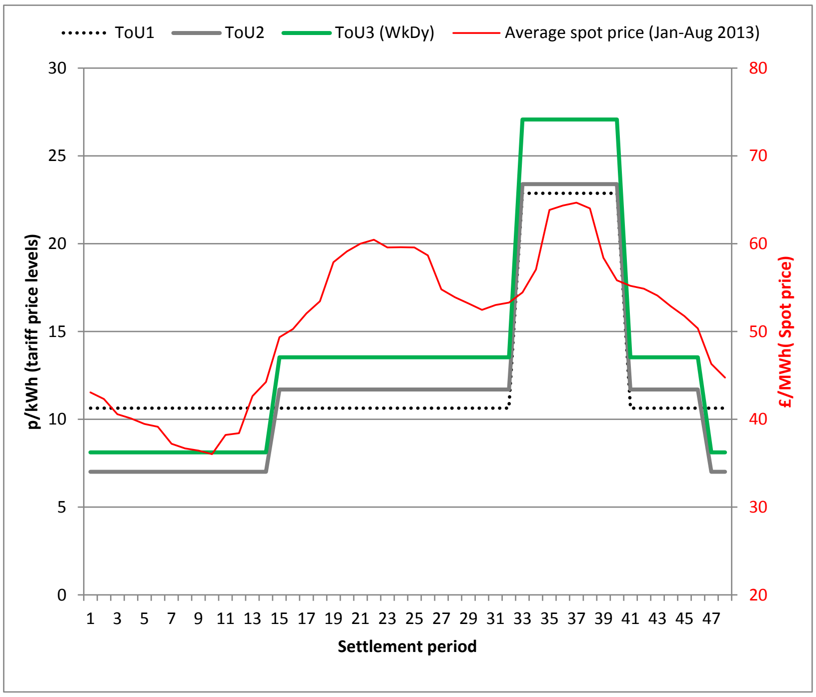

Figure 13: TOU tariffs vs Spot price for electricity

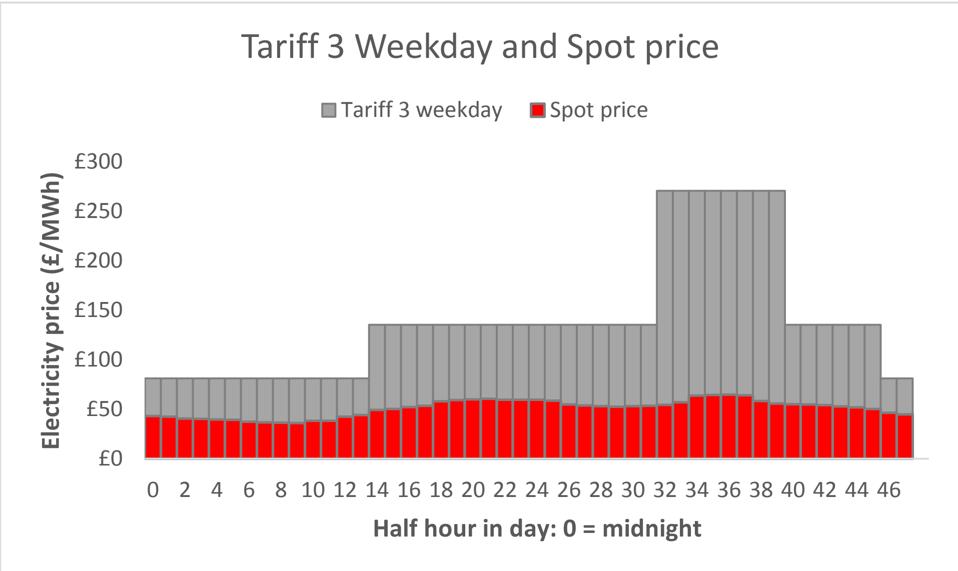

Figure 14: Tariff 3 weekday and spot price, same scale

The time bands and differentials between the price levels were set based on a review of a range of ToU trials and research from across industry.[4]

Ofgem selected tariffs with the intention of demonstrating the effect of as many characteristics as possible with a limited number of tariffs. This means incrementally adding differentiation to each tariff, so as to isolate the effect of each individual characteristic. In summary:

4 These include Energy Demand Research Project: Final Analysis, Ofgem, June 2011, Low Carbon Network Fund Customer-led Network Revolution (2013), Demand Side Response in the domestic sector- a literature review of major trials, DECC, August 2012 "

- The first tariff isolates the effect of charging a higher tariff at peak, in a two-rate tariff.
- The second tariff incorporates a cheaper night rate in addition to a higher peak rate. This brings it up to a three-rate tariff.
- The third tariff incorporates a weekend rate in addition to the three rate tariff, exploring the impact of weekend pricing variations.

Tariff levels

Ofgem set the price levels for each TOU tariff so that the reference case PC1 demand profile would experience no change in annual bill compared to a standard tariff.[5] The price ratios between the different price levels in each of the three ToU tariff were fixed, and a goal-seek function was used to determine the price levels that would give the reference case the same annual bill as the standard tariff. In other words, the price levels of the three ToU structures are set such that costs are fully recovered for the reference load shape (but not all load profiles found in the EDRP sample: see Section 4.4)

This revenue-neutral approach allows for the isolation of the impact of the tariff structure on customers' overall bills.

Figure 15 below illustrates the three TOU tariffs alongside the baseline tariff.

[5] Ofgem used the average standard tariff from our supply market indicator in August 2013.

On each of the graphs in Figure 10 below, the Y axis is the unit price for electricity in £/kWh. The X axis shows the time of day in half hours from midnight, where midnight is zero. All tariffs included a daily standing charge of 22p in addition to the time-of-use unit pricing.

<table>
<tr><td>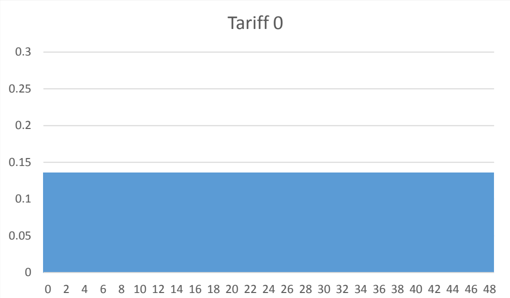
</td><td>This is the baseline tariff against which the bills from the time of use tariffs were compared.

It consists of a flat rate of 13.6 pence with a daily standing charge of 22 pence.</td></tr>
<tr><td>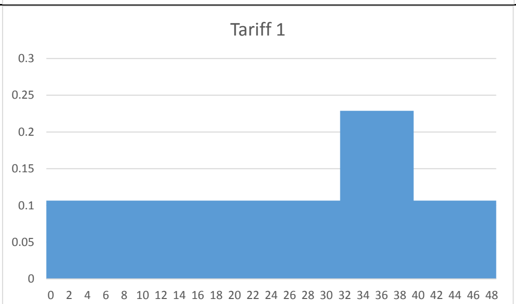
</td><td>Tariff 1 applies across all 7 days of the week and has a daily standing charge of 22 pence.

The unit prices are:
22.9 pence from 16:00 - 20:00
10.6 pence the rest of the time.</td></tr>
<tr><td>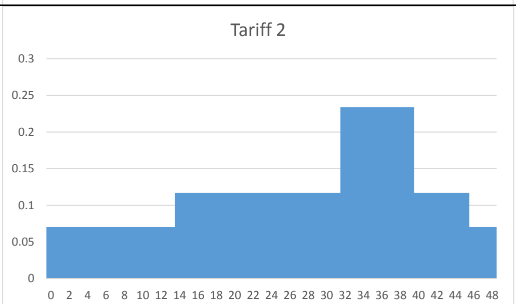
</td><td>Tariff 2 applies across all 7 days of the week and has a daily standing charge of 22 pence.

The unit prices are:
23.4 pence from 16:00 - 20:00
11.7 pence from 14:00 - 16:00 and from 20:00 – 23:00
7.0 pence the rest of the time</td></tr>
<tr><td>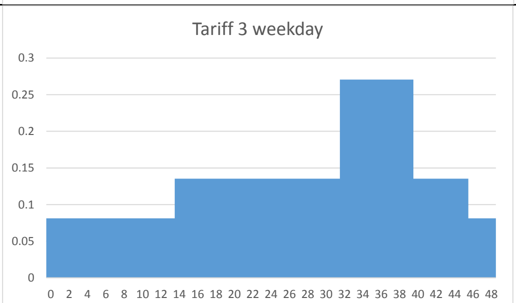
</td><td>Tariff three has a separate weekday and weekend component. The weekday component is shown to the left, and has a daily standing charge of 22 pence.

The unit prices are:
27.1 pence from 16:00 - 20:00
13.5 pence from 14:00 - 16:00 and from 20:00 – 23:00
8.1 pence the rest of the time</td></tr>
<tr><td>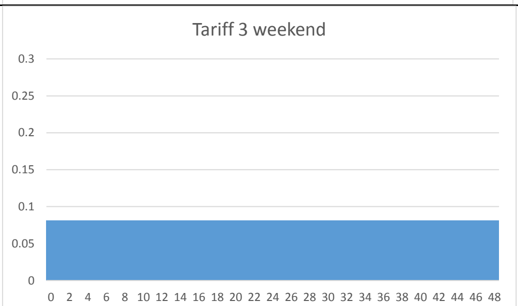
</td><td>The weekend component of Tariff 3 has a daily standing charge of 22 pence and a flat unit rate of 8.1 pence.</td></tr>
</table>

Figure 15: The Tariffs

4.4 Results

This section presents the results of the tariff modelling. Figure 16 shows the distribution of the baseline bills applied to the annual consumption of the EDRP cases selected for analysis as per Section 2. This shows a right-skewed distribution which is consistent with other datasets on the distribution of energy expenditure/consumption (for example the Living Costs and Food Survey). This is the distribution of bills against which the impacts of the TOU tariffs are calculated.

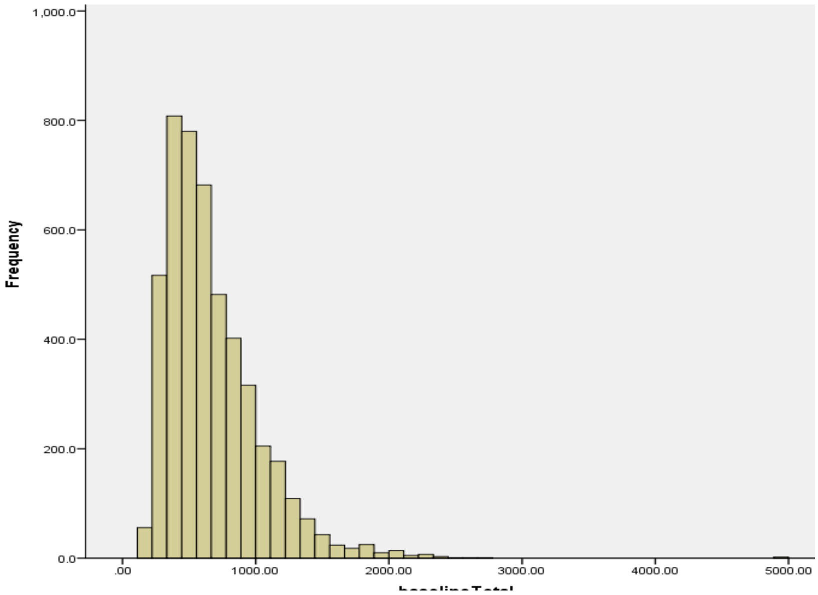

Range: £126.87 - £4,905.91

Mean: £672.33

Median: £596.62

Standard Deviation: £354.95

Figure 16: Distribution of bills under baseline flat tariff

The following section presents the effects of the tariffs on bills. **It is essential to note that the analysis of change in bills presented hear assumes no change in usage pattern in response to the introduction of a time of use tariff.** In reality, the whole point of introducing time of use charging is to encourage the shifting of demand away from peak times, to reduce system costs. However modelling of demand response was beyond the scope of this initial work.

Figure 17 gives the distributions of gain and loss for each of the tariffs, in percentage and absolute terms. These are relatively normal (i.e. the mean, median and modal values are close), with means slightly higher than zero. However the asymmetry in the ranges (that is, the minimum and maximum values in the distributions are not equally spaced around zero) suggests that a small number of extreme outliers are experiencing larger changes in bills. This is consistent with (a) the skewed distribution of baseline bill shown in Figure 16, and (b) the observation that the average EDRP household day has a higher and earlier evening peak than the PC1 demand curve (see Figure 4).

This observation is reinforced by Table 1 and Figures 18-20: here we can see that for deciles 2-9 there is a symmetrical and relatively tight distribution around the mean, but that in deciles 1 and 10 (which by definition contain the outlier cases), there appear to be a small number of outliers experiencing larger changes in bill – the averages for these deciles are being influenced by these outliers. For tariff 1 this effect is larger for cases experiencing increases in bills, while for tariffs 2 and 3 the effect is larger for cases experiencing a decrease in bills. However, on the whole there is a slight bias towards larger increases, which is explained by the higher average evening peak in the EDRP dataset when compared with the PC1 demand curve (see Figure 4).

Note that these patterns result directly from the interaction between the halfhourly usage patterns in the EDRP sample, and the structures of the tariffs modelled in this work. While different tariff structures would have produced different sets of outlier cases, it is possible that similar distributions would arise whatever the structure of tariffs modelled. As shown in Figure 4 there is a slightly earlier and higher average evening peak in the EDRP sample dataset on which the tariffs were modelled, compared with the PC1 data upon which the tariffs were designed to be revenue neutral. This would explain the very slight overall increase in bills seen here.

Tariff 1: Absolute change in bill £/yr 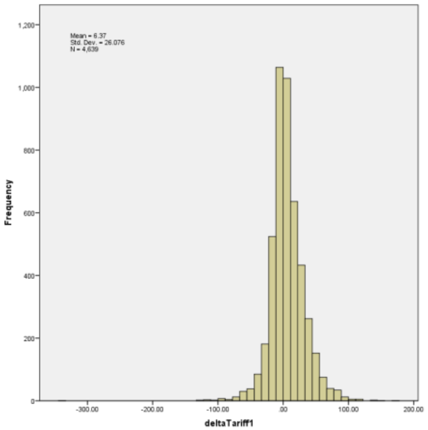 Range: £-342.52 to £169.16 Mean: £6.37 Median: £3.45, SD: £26.93	Tariff 2: Absolute change in bill £/yr 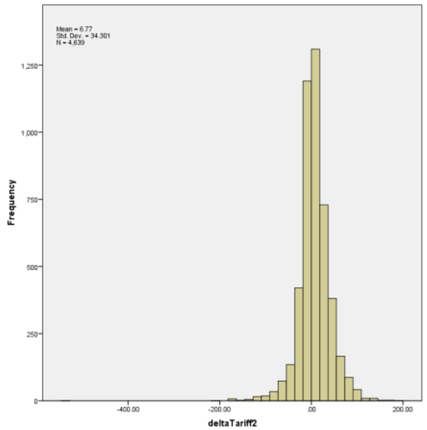 Range: £-536.96 to £194.10 Mean: £6.79, Median: £4.75, SD: £34.07	Tariff 3: Absolute change in bill £/yr 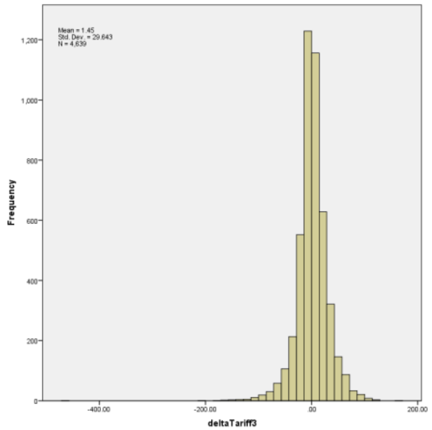 Range: £-465.26 to £171.22 Mean: £1.59, Median: £0.82, SD: £29.54
Tariff 1: % change in bill £/yr 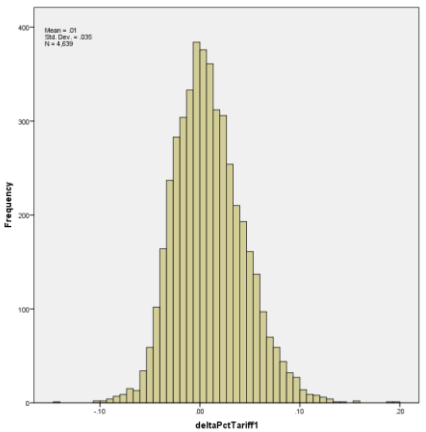 Range: -15% to +20% Mean: 0.1%, Median: 0.6% SD: 3.5%	Tariff 2: % change in bill £/yr 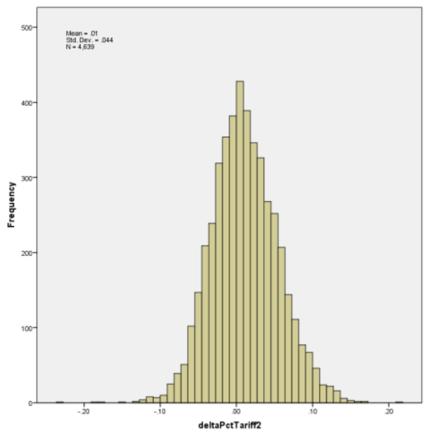 Range: -23% to +21% Mean: 1.1%, Median: 0.8% SD: 3.4%	Tariff 3: % change in bill £/yr 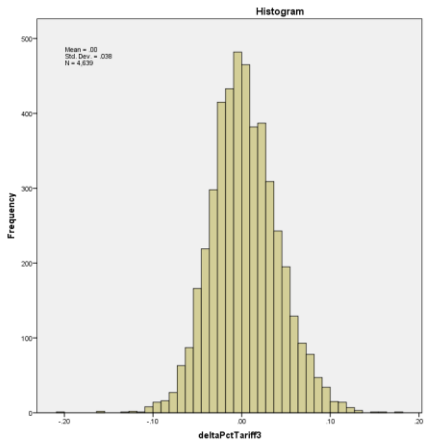 Range: -21% to +18% Mean: 0.4%, Median: 0.2% SD: 3.7%

Figure 17: Histograms showing distribution of absolute and % gain/loss by tariff

	% Change under Tariff 1			% Change under Tariff 2			% Change under Tariff 3		
Decile	**Min**	**Max**	**Mean**	**Min**	**Max**	**Mean**	**Min**	**Max**	**Mean**
1	-15%	-3%	-4%	-23%	-4%	-6%	-21%	-4%	-6%
2	-3%	-2%	-3%	-4%	-2%	-3%	-4%	-3%	-3%
3	-2%	-1%	-2%	-2%	-1%	-2%	-3%	-2%	-2%
4	-1%	0%	-1%	-1%	0%	-1%	-2%	-1%	-1%
5	0%	1%	0%	0%	1%	0%	-1%	0%	0%
6	1%	1%	1%	1%	2%	1%	0%	1%	1%
7	1%	3%	2%	2%	3%	3%	1%	2%	2%
8	3%	4%	3%	3%	5%	4%	2%	3%	3%
9	4%	6%	5%	5%	7%	6%	3%	5%	4%
10	6%	20%	8%	7%	21%	9%	5%	18%	7%

Table 1: % gain/loss by decile for the 3 TOU tariffs

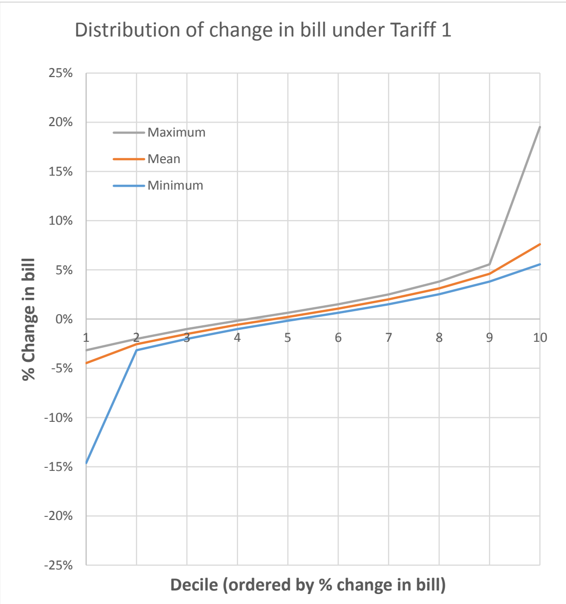

Figure 18: % gain/loss by decile for tariff 2

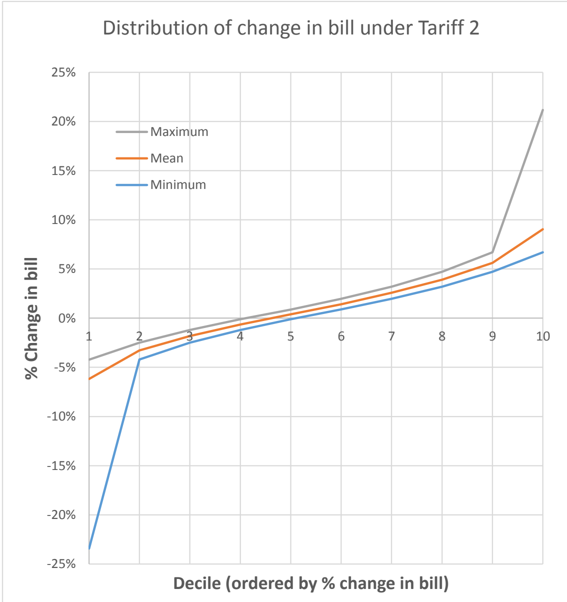

Figure 19: % gain/loss by decile for tariff 3

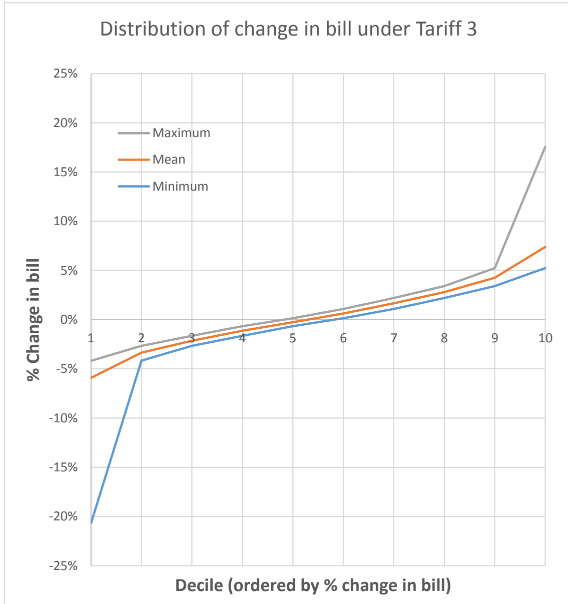

Figure 20: % gain/loss by decile for tariff 1

Figure 21 summarises the impacts of the TOU tariffs on the 3 cluster centroids developed in Section 3. For each tariff, the absolute and % change in bill are shown, along with the change in demand (kWh per year) required to offset those changes. This calculation is based on moving demand from the most expensive to the cheapest period for each tariff.

Cluster centroids for k=3: Effects of Tariffs (baseline bill = £652.58)		**1**	**2**	**3**
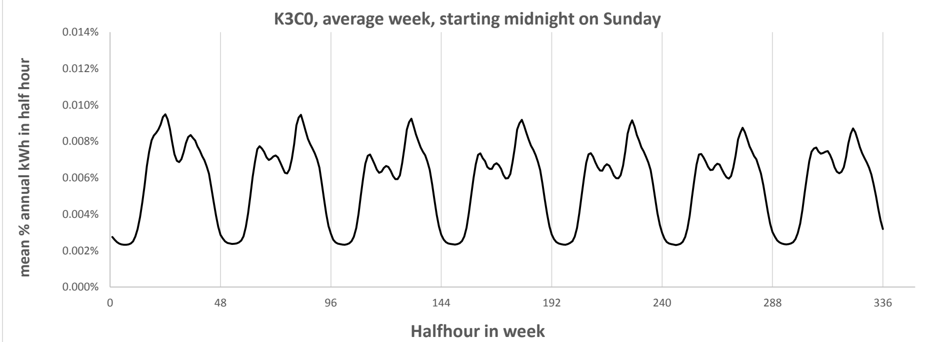	£ Change in bill	-£0.57	+£3.76	+£3.83
	% Change in bill	-0.1%	+0.6%	+4.5%
	kWh shift to offset	+4.7 kWh	-23kWh	-20 kWh
	% kWh shift to offset	+0.7%	-3.5%	-3.1%
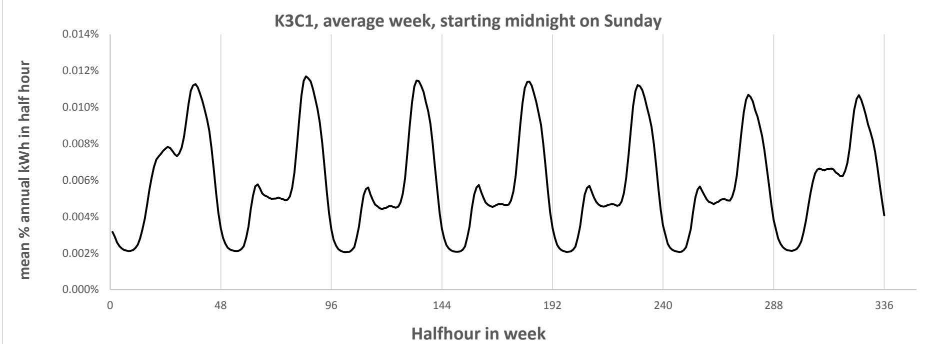	£ Change in bill	+£28.79	+£33.90	+£29.45
	% Change in bill	+4.41%	+5.20%	+4.51%
	kWh shift to offset	-235 kWh	-207 kWh	-155 kWh
	% kWh shift to offset	-36.1%	-31.7%	-23.8%
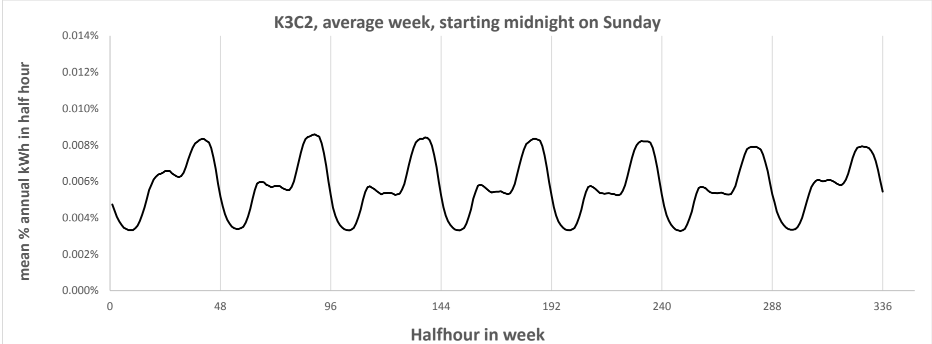	£ Change in bill	-£9.82	-£16.80	-£13.46
	% Change in bill	-1.51%	-2.57%	-2.06%
	kWh shift to offset	+80 kWh	+103 kWh	+71 kWh
	% kWh shift to offset	+12.3%	+15.7%	+10.9%

Figure 21: Effects of tariffs on cluster centroids for K=3

Figure 22 gives the same summary, but this time for the top 5% of EDRP cases in terms of gain and loss under the three TOU tariffs. The graphs show the average week for the cases in the top 5% of losers and winners respectively. Here it is very clear why these cases have lost and gained: the losers have an extremely high evening peak, a very low morning peak, and a peak:base ratio of 10.3. The winners are the opposite: the morning peak is actually higher than the evening peak, and demand on the whole is much less variable, with a peak:base ratio of 2.7. The scales of the two graphs are the same.

5% biggest losers: effects of tariffs

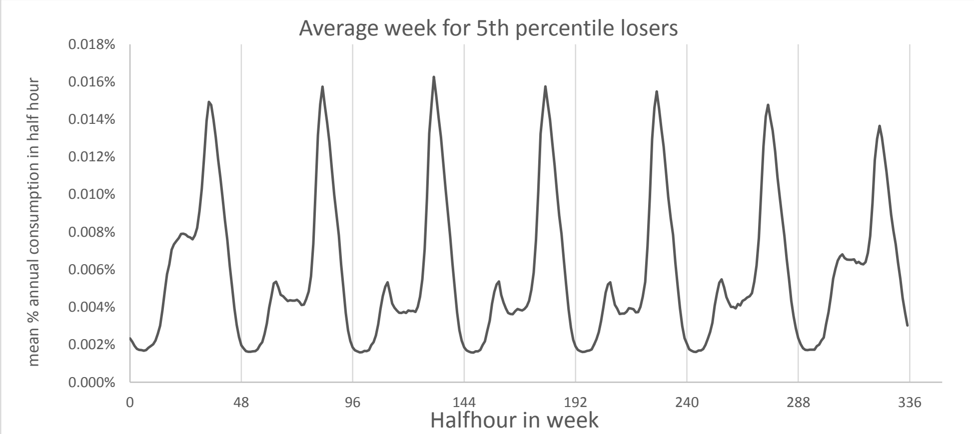

Tariff:	1	2	3
Average consumption	4141 kWh	4324 kWh	4107 kWh
Average baseline bill	£643.41	£668.34	£638.87
£ Change in bill	+£57.29	+£70.10	+55.71
% Change in bill	+9.0%	+10.6%	+8.7%
kWh shift to offset	-466	-427	-293
% kWh shift to offset	-11.2%	-9.9%	-7.1%

5% biggest winners: effects of tariffs

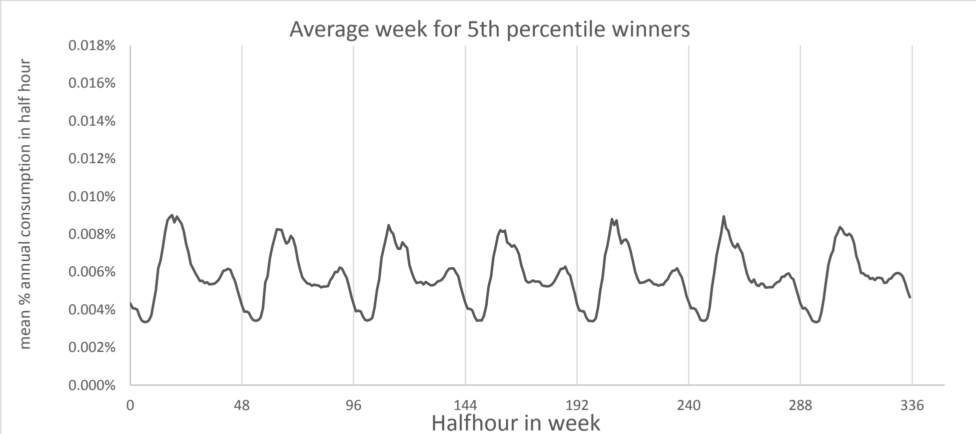

Tariff:	1	2	3
Average consumption	4387 kWh	5036 kWh	5292 kWh
Average baseline bill	£676.90	£765.18	£799.98
£ Change in bill	-£36.78	-£58.69	-£56.71
% Change in bill	-5.4%	-7.5%	-7.1%
kWh shift to offset	+299	+358	+298
% kWh shift to offset	+6.8%	+7.1%	+5.6%

Figure 22: Effects of tariffs on the 5th percentile winners and losers

Table 2 contextualises the demand shifts required to offset the changes in bill resulting from the three TOU tariffs, for the three cluster centroids shown in Figure 21, and for the 5% biggest winners and losers shown in Figure 22.

Demand profile	Approximate shift from peak to off-peak or peak reduction required to offset bill change	Example of behavioural responses required
Centroid 0	+20 kWh per year	Automated response of fridge or freezer to avoid running at daily peak half-hour (peak shift)
Centroid 1	+200 kWh per year	Replace halogen and CFL light bulbs with LED equivalents (peak reduction)
Centroid 2	-100 kWh per year	No response required
5% biggest losers	+400 kWh per year	Shift use of electric hob/cooker by one hour from peak demand (peak shift)
5 % biggest winners	-100 kWh per year	No response required

Table 2: Examples of behavioural responses which would offset changes in bill resulting from TOU tariffs

4.5 Discussion

Three different kinds of tariff modelling result have been presented here: descriptions of the distribution of results across all selected EDRP cases; detailed results for the three cluster centroid profiles, and; detailed results for the 10% most affected cases from the EDRP (the latter split into the cases with the 5% largest increases and reductions in bills). It is worth noting again here that this analysis is based on an assumption that there is no change in demand patterns in response to the TOU tariff. Of course, in reality, the whole point of TOU tariffs is to encourage changes in demand.

All these results must be considered in the context of the specific three TOU tariffs used in the modelling. Many different TOU tariff structures are possible, and different structures could lead to different distributions of gain and loss. Furthermore the EDRP dataset cannot be assumed to be representative of UK domestic electricity consumption. Finally, the EDRP sample used here was explicitly filtered to remove cases which were already experience TOU related interventions (e.g. PC2 customers, and those in TOU trials). These cases may have been more likely to benefit from the TOU tariffs.

That said, in each case the results are not surprising: the overall distribution of results makes sense when we take into account the distribution of baseline bills for EDRP cases, the difference in the average day for the EDRP cases versus PC1, and the fact that the tariffs were designed to be revenue neutral for PC1. Equally, the three cluster centroids presented in Section 3 interact with the tariffs presented in Section 4 in three different but predictable ways: centroid 0 is the most similar to PC1, and experiences the least change in bill. Centroid 1 is the profile with the highest evening peak, and consequently experiences the largest increase in bill. Centroid 2 is has the lowest peak:base ratio, and the lowest evening peak, and experiences a decrease in bill. The 5th percentile of cases with bill increases have an extremely high peak:base ratio of 10.3, while the 5th percentile of cases with bill decreases have the a relatively flat profile, with a peak:base ratio of 2.7. This is comparable to that of centroid 2, but here the highest peak occurs in the morning rather than the evening, leading to a larger reduction in bill.

TOU tariffs are a price instrument for encouraging a shift away from peak times. Their structure has to reflect this. In the EDRP data sample there are some households which contribute disproportionately to peak loads as defined in the three tariffs. The modelling shows that they would experience very high increases in bills if they did not shift their demand. However if this were not the case, then the tariffs could not be expected to perform their function.

The modelling done here did not consider the potential for households to respond to the TOU signal by shifting demand to cheaper times of day. The key question from a distributional impacts point of view is when different sets of households, including vulnerable households, use energy and how easily they could shift their demand to gain (rather than lose) from the TOU price signals. This needs further work covering (1) the sociodemographics associated with different usage patterns, and (2) the kinds of energy service use which lead to demand patterns which coincide with peak electricity usage across the domestic sector.

5 Conclusions

The rollout of smart meters will enable energy suppliers to deploy new time of used based approaches to pricing electricity for households. The ability of consumers to engage positively with TOU tariffs depends on how households currently use electricity, their ability to shift demand over the day/week/season, and the range of TOU products on offer to suit the needs of different households. The work presented here shows how certain groups of households might be made better or worse off under a set of illustrative TOU tariffs, assuming no behaviour change.

The cluster analysis identified three main patterns of electricity use within the set of EDRP cases it was based on. There will be other customers who display different usage use patterns, such as nightworkers, houses in multiple occupancy, etc. Furthermore the analysis excluded PC2 customers who could be expected to have higher demand over the night, and who effectively are already on a TOU tariff, albeit a fairly simple one.

By definition, those who will benefit from TOU tariffs either already use more electricity at cheaper times of day, or are willing and able to shift their demand accordingly. In the illustrative TOU tariffs used in this work, those households that are best off are those with below average consumption at peak times. Other potential tariff structures, such as free Saturdays/weekends, tariffs focused on electric vehicle charging, or network-led tariffs, would benefit other groups.

It is essential to ensure that consumers are well-informed about their energy use before considering switching to a TOU tariff. Smart meters and related energy services have the potential to support this, but education will be crucial to ensure consumers are moving onto the right tariff for them, in a way that enables the industry to realise increases in overall efficiency by reducing the aggregate peak:base load.

It is important to reiterate that this research represents a contribution to the very initial stages of Ofgem's consideration of TOU charging for domestic electricity consumption. While providing some interesting findings, it also has limitations which CSE considers would need to be addressed in any future work. Initial thoughts on how these might be addressed include:

1. Further investigation of the outputs shown here, including but not limited to: more detailed analysis of the strength of the clustering results; focus on outlier demand patterns; investigation of other approaches to clustering; and investigation of a range of other TOU tariff structures.
2. Repetition of the analysis done here using larger and/or more representative samples of half hourly smart meter data.
3. Introduction of sociodemographic variables and investigation of whether there are interesting relationships between sociodemographic patterns and time-of-use patterns.
4. Consideration in detail of the characteristics of households more and less likely to be able to adapt their electricity demand to respond to TOU tariffs and develop model to include range of likely consumer responses.