SDRC 2.3

CUSTOMER MODEL

SAVE
Solent Achieving Value from Efficiency

Scottish & Southern
Electricity Networks

Solent Achieving Value from Efficiency

Solent Achieving Value through Efficiency (SAVE) is an Ofgem funded project run by Scottish and Southern Electricity Networks (SSEN) and partnered by the University of Southampton (UoS), DNV GL and Neighbourhood Economics (NEL). The innovative programme evaluates the potential for domestic customers to actively participate in improving the resilience of electricity distribution networks and thereby defer the need for traditional reinforcement. The government has forecasted an increase in electricity demand of 60% by 2050 meaning peak demand is likely to grow to six times higher than what the network was designed for.

CONTENTS

List of Tables

List of Figures

EXECUTIVE SUMMARY

This Successful Delivery Reward Criteria (SDRC) report presents the outcomes of the Customer Modelling Framework conducted in the Solent Achieving Value through Efficiency (SAVE) project. The report builds on two previous SDRC reports: SDRC 2.1: SAVE *Customer Model Framework Specification*, and SDRC 2.2: *Updated Customer Model*. This report provides details of the final implementation of the Customer Model within the SAVE project, and documentation of the work carried out to successfully meet the project objectives.

The Customer Model has demonstrated a method to generate household electricity demand profiles under 'baseline' and intervention conditions, and to provide network planners with half-hourly electricity demand profiles for a range of Customer Types: households defined by occupancy, dwelling size and primary heating fuel. Profiles have been generated from the representative SAVE project sample, providing household demand data through the implementation of a customer typology that represents a greater diversity in household demand during peak hours and daily demand profiles than existing industry standards.

Evaluation of the interventions trialled under the SAVE project was conducted as part of the Customer Model Framework. The analysis presented in this report provides evidence on the effectiveness of a number of distribution network operator (DNO) led interventions aimed at reducing load during the peak hours of domestic demand (between 16:00 and 20:00 hours), giving estimates of the impact of each intervention.

For the roll-out of LED lighting upgrades, an estimated maximum average demand reduction per household during peak hours of 7 percent was observed, equivalent to 47 Watts (90% CI, -96 to 7 W)[1].

The evaluation of the data-informed and price signals treatments tested during the second trial period estimated average demand reduction per household during peak hours as follows:

- Data-informed and price signals: 7 percent reduction (35W) observed during Challenge 4[2]
- Data-informed only: 3.8 percent reduction (21W) observed during Challenge 1

In the dynamic pricing trial held during the third and final trial period, the estimated maximum average demand reduction per household, averaged over the peak hours, were as follows:

- Opt-in recruitment: a 2.6 percent reduction (-17W, 90% CI -71 to 43);
- Opt-out recruitment: a 7.1 percent reduction (-44W, 90% CI -97 to 15).

1 90 percent confidence interval around the estimated treatment effect. Negative values are reduction in load.
2 Challenge 4 targeted only the central two hours of the peak period, 17:00 to 19:00 hours.

The load reductions observed during the SAVE trials have been represented within the demand profiles created using consumption data measured under intervention conditions. The Customer Model has also demonstrated a method that enables network planners to model the change in half-hourly demand under intervention scenarios by representing the change in electricity demand for each customer type under a range of interventions using intervention ‘impact profiles’.

Finally, the wider application and future development of the Customer Modelling framework has been addressed. The Customer Model provides a replicable and scalable model, allowing other distribution network operators to apply customer load profiles specific to the mix of households within their own geographic region for simulation work. The model is also expandable, allowing for suitable high-quality data to be input into the framework to expand and refine both the customer typology and output load profiles.

Within the SAVE project, the Customer Model is a module of the Network Investment Tool (NIT). Outputs from the Customer Model provide the NIT with a diverse set of customer type demand profiles under both ‘baseline’ and ‘intervention’ conditions. These profiles are allocated to individual network topologies, according to the mix of customers connected specific to individual assets, allowing the simulation and evaluation of a range of investment options. Details of the interaction of the Customer Model and other modules within the NIT are provided in SDRC 8.2: *Network Investment Tool* and SRDC 8.5 and 8.6: *Pricing Model, Customer Model and Network Model*.

1

INTRODUCTION

This Successful Delivery Reward Criteria (SDRC) report presents the outcomes of the Customer Modelling Framework conducted in the Solent Achieving Value through Efficiency (SAVE) project. The report builds on two previous SDRC reports: *SDRC 2.1: SAVE Customer Model Framework Specification* and *SDRC 2.2: Updated Customer Model*.[3] This report provides details of the final implementation of the Customer Model within the SAVE project, and documentation of the work carried out to successfully meet the project objectives.

1.1 Background

Network planners have generally employed standard load profiles for domestic customers that capture very little variation in the load profile between individual, or groups of, customers. The use of such profiles imposes a significant constraint on the simulation of low-voltage (LV) network infrastructure. In order for distribution network operators (DNOs) to make sound investment decisions, more accurate modelling of the thermal and voltage profiles within network assets is required. Providing realistic representation of the demand profiles of different household types is central to this task. These demand profiles are required to provide realistic loads for network asset modelling under both 'baseline' and 'intervention' scenarios in order to evaluate cost-effective solutions and to mitigate traditional reinforcement.

Box 1 Overview of Customer Model

The Customer Model provides household electricity demand profiles under 'baseline' and intervention conditions. Profiles are generated from the household demand data collected from the SAVE sample for a number of Customer Types.

The Customer Types have been developed to represent the differing levels of household demand associated with a number of characteristics (household size, dwelling size and primary heating type). Intervention 'impact' profiles are generated to represent the treatment effects (change in electricity demand) observed under a number of SAVE trial conditions.

1.2 Project objectives

To meet the project objectives, the requirements of the SAVE Customer Model Framework were set out in *SDRC 2.1 SAVE Customer Model Framework Specification*.

The initial requirements for the SAVE customer model were as follows:

- The ability to produce 'baseline' half-hourly electricity consumption profiles at the individual household level for any day of the year (or aggregation of days) as input to the Network Model;
- To produce similar profiles for trial intervention groups as input to the Network Model, taking account of intervention and community trial effects *where feasible*;
- To produce similar profiles for designated Census areas in the Solent region under a range of demand response scenarios including those trialled by the SAVE project;
- The estimation of electricity consumption increase/decrease at specific times of day that can be attributed to the SAVE intervention trials for overall effect reporting;
- The analysis of the household economic, demographic and behavioural factors that mediate these changes to provide insights relevant to future DNO interventions;
- The ability to estimate changes in temporal (half hourly) demand that might ensue from other (non-trialled) behavioural changes;

These objectives have been met by the outcomes described in this and other project reports relating to the Customer Model. Table 1 provides a summary of the wider project objectives and signposting to the supporting documentation.

3 Available online at https://www.ssen.co.uk/save/

Table 1: Project objectives relevant to Customer Model

Project objective/outcome	Supporting evidence (Section)
Review of existing literature to provide a sound basis to customer model development	See SDRC 2.1 and 2.2.
Customer model development: a model which will allow interrogations of scenarios and undertake simulations	Spatial microsimulation tool to generate area-level (aggregated) demand profiles (see SDRC 2.2). This capability discontinued and replaced through the development of Customer Type demand profiles (Section 3) and Census Interface to apply profiles within network simulations (see SDRC 8.6).
Ability to distinguish between effects caused by selection and by actual measures	Randomised Controlled Trial design (see SDRCs 4 & 2.2).
Ability to detect 'subconscious' behavioural change from observations.	Use monitored electricity consumption to determine response to interventions (Section 2.4).
Produce customer model revealing customer receptiveness to measures	Analysis of variation in treatment effects by socio-demographic variables (Section 2) and application to Customer Model (Section 3.4).
Ability to distinguish between novelty effects and longer-term change. How enduring are the impacts of each measure and what costs if any are associated with sustaining the impacts?	Short-term and longitudinal analysis performed to evaluate impact of trial interventions (Section 2.4, see also SDRC 2.2 for trial period 1 evaluation). Costing information is provided in reports detailing the specific interventions trialled under the SAVE project (see SDRC 8.2, 8.3, 8.4 and 8.7).
Ability to detect statistically significant effects	Sample size of treatment groups determined by statistical power analysis (see SDRCs 4 & 2.2). Use of appropriate statistical tests to evaluate treatment effects (Section 2.4).
Ability to extrapolate results to the general customer population	Sample designed to be representative of wider customer base (see SDRC 2.2). Development of Customer Type demand profiles (Section 3.3).

1.3 Report structure

This report presents the outcomes of Work Package 2: Customer Model to meet the project objectives set out above. The report is structured with content provided in the following sections:

Section 2: provides a summary of the analysis completed for the evaluation of the second and third trial periods. The data sources and methodology used in the evaluation are presented, alongside a summary of the key findings from each of the interventions tested. This section will be of interest to academics, consultants and innovation teams reviewing new evidence of energy efficiency and behaviour change trials in the UK, and also of best-practice in trial design and evaluation;

Section 3: provides details of the development of the Customer Model with a focus on the development of the *Customer Type* demand profiles. This section is relevant to network engineers and technicians looking to understand the methodology and build of the customer model, the development of customer typology and the representation of demand profiles in the model;

Section 4: presents the wider applicability and future development of the customer modelling framework and details of how it can be extended using additional data sources and applied to other network regions. This section will be of interest to wider planning teams and management, concerned with applying the Customer Model framework to their own operations. It is also relevant for DSO teams concerned with future development of network modelling innovation;

Section 5: provides the conclusion and recommendations.

2

TRIAL ANALYSIS

2.1 Overview

This section provides an overview of the data used in the evaluation of the impact of the interventions trialled during the SAVE project; a description of the analysis methods, assumptions and statistical models used; and high-level summaries of results for each of the interventions. In addition, a description of how trial impact analysis fed into the Customer Model is provided.

Project objectives and knowledge gaps evidenced in this section:

Ability to detect 'subconscious' behavioural change from observations.
Produce customer model revealing customer receptiveness to measures
Ability to distinguish between novelty effects and longer-term change. Persistence of the impacts of each measure.
Ability to detect statistically significant effects

2.2 Data

The specification of the achieved data sources used in the evaluation of the SAVE trial interventions is as follows:

Household socio-economic and demographic data from an initial recruitment survey and repeated 'update' surveys implemented during the repeated waves of fieldwork;

Household response person time-use activities recorded at 15-minute intervals during the period of the time-use diary;

Dwelling level electricity consumption data at 15-minute intervals (in watt-hours, Wh/15-min).

2.2.1 Household survey data

This report uses household survey data collected by the fieldwork contractor (BMG Research). This dataset contains the socio-economic and demographic data for the participants in the fieldwork, along with other information about the dwelling occupied and appliances owned by each household. Update surveys were conducted at intervals during the trials where data was over 12 months old to ensure that basic household attributes such as number of occupants were accurate.

The fieldwork contractor also collected data on all LED lightbulb installations completed during the period August 2017 to March 2018. The data included the original (replaced) bulb ratings, ratings of (new) replacement bulbs and the location where each bulb was installed.

2.2.2 Time-use diary data

As with the first trial period, time-use diary data was also collected by the fieldwork contractor during the second and third trial periods. This data collected consisted of a sequence of activities for each survey respondent, each with a start and finish time. The activities recorded by the survey were allocated to categories using a modified version of the Multinational Time-Use Survey coding system[4].

2.2.3 Dwelling electricity consumption data

The analysis in this report is based on the electricity consumption data collected via the internet-connected 'Loop' electricity monitoring kit (hitherto referred to as 'Loop data'). The 'Loop' data used in the analysis consists of watt-hour (Wh) readings observed at 15-minute intervals for each participating household. This data provides the measure of electricity consumed by individual households within the treatment and control groups during the trial periods. Before analysis, the Loop electricity consumption data was processed and summarised over a number of time periods and intervals: for example, producing hourly and weekly mean consumption values for each household. Data cleaning was also conducted to ensure that faulty installations of the Loop kits and erroneous consumption values were not included in the analysis.[5]

2.3 Third party data

The analysis of trial interventions also utilised additional data from the following sources:

2.3.1 Weather data

Met Office weather data was used in the analysis to provide an estimation of household heating loads. The hourly data used was collected at Middle Wallop between the dates 2016-09-30 and 2017-03-31 and was downloaded from the Met Office Weather Observations Website[6]. The hourly weather data was pre-processed prior to use to create daily and weekly average temperatures, and to calculate heating degree-days.[7]

4 See https://www.timeuse.org/mtus for more details.

5 Further details of the data cleaning procedures can be found in the technical annex: Rushby and Harper (2018), *SAVE Loop Energy Saver Data Cleaning and Preprocessing*. SAVE Project Report, University of Southampton.

6 For more information refer to the Met Office website: http://wow.metoffice.gov.uk/

7 For details of weather data processing see Anderson and Rushby (2017), *Process Met Office WOW data for the SAVE study region*, SAVE Project Report, University of Southampton.

2.3.2 Simulated daylight data

This report uses sun-path simulation data produced by the Transient System Simulation Tool (TRNSYS) software[8] to estimate local sunrise and sunset times. The simulation used Southampton as the location.

2.4 Methods

2.4.1 Experimental design

Given the randomised control trial (RCT) design of the SAVE trials, intervention effects have been analysed by comparing the difference between control and intervention groups. Given the successful randomisation and allocation of participants to treatment and control groups, the assumption is that prior to treatment, the groups would be equal in terms of both the outcome variable and household characteristics. Any difference in consumption between the control and intervention groups is therefore assumed to be a result of the intervention alone.[9] It is assumed that all households in the study experienced the same environmental conditions during the trial weeks and therefore there is no need to correct for any differences in environmental conditions.[10] This means the results should be replicable and scalable to the wider population. Using an RCT approach limits biases that may be present in the trial groups by comparing results to a similar control group, instead of past behaviour of the treatment group.

The analysis in this report (along with previous analysis presented in SDRC 2.2) indicates that the treatment groups show small but consistent differences in consumption to that of the control group. For this reason, the analysis also employs the difference-in-differences statistical technique for analysis (see Section 2.5.2 for more information).

Due to the design of the study, it is not necessary to control for potential confounding characteristics of the households in each treatment group. However, a selection of household attributes is included in the analysis to examine characteristics that are associated with the variability in treatment effect.

2.4.2 Assumptions and limitations

2.4.2.1 Experimental design and analysis

As with any experimental study, a number of limitations apply to the findings of the trial analysis. General limitations apply to the analysis of the interventions arising from both sampling and statistical analysis. In summary, limitations of this study are related to the following:

- Recruitment of trial participants: the analysis assumes the sample was randomly assigned to treatment groups and therefore the groups are representative of the sampled population with respect to both the mix of household socio-demographic and electricity consumption characteristics (see SDRC 2.2[11]).

- Statistical power: the achieved sample size and variability of household electricity consumption limit the size of the effect that can be robustly detected (see Anderson and Rushby, 2018[12]). In general, the smaller the treatment effect, the larger the sample size required to observed that effect with confidence.

- Experimental conditions: it is assumed that all households experienced the same environmental conditions during the trial negating the need to correct for any differences despite local variation in environmental conditions (such as weather).

- Analytical assumptions: for example, parallel trend assumption of the difference-in-differences technique may not hold (see Statistical models, Section 2.5).

2.4.3 Metric of measurement

The metric of measurement used in the analysis of intervention impacts was mean 15-minute consumption summarised across various time-periods appropriate to the analysis conducted. As the distribution of household consumption was observed to be skewed, a log transformation was applied for statistical modelling. The outcome variable of the models reported is therefore log-mean 15-minute consumption.

8 TRNSYS is a graphical software tool used to simulate the behaviour of transient systems such as energy, or in this case, sun-path. The SAVE project used the TRNSYS software to model sunrise and sunset times to estimate daylight hours in Southampton during the trials. More information available here: http://www.trnsys.com/

9 Frederiks, E.R., Stenner, K., Hobman, E.V., Fischle, M., 2016. Evaluating energy behavior change programs using randomized controlled trials: Best practice guidelines for policymakers. Energy Research & Social Science 22, 147–164. https://doi.org/10.1016/j.erss.2016.08.020

10 Weather data has however been used in the analysis of the electricity consumption data in order to contextualise the observed trends.

11 Available at https://www.ssen.co.uk/save/

12 Anderson, B., Rushby, T., 2018. We Got The Power: Statistical Significance, Power, Study Design and Decision Making with A Worked Example. University of Southampton, Southampton, UK.

2.4.4 Analysis approaches

The evaluation of interventions tested within the SAVE trials involved using a number of analytical and statistical methods. A combination of methods was tailored to each intervention according to the nature of the hypothesised treatment effects. In order to examine the impact of each intervention, the trial analysis was generally configured using two approaches: 'short-term' and 'longitudinal'.

2.4.4.1 Short-term

This form of analysis was directed toward those interventions that aimed at encouraging short-term reductions in consumption over a number of hours or days during a targeted period. The events targeted varying lengths of time (one to five days) and periods of the day (4pm to 8pm and 5pm to 7pm) and therefore required a flexible and high-resolution analysis approach to detect any changes in consumption associated with the interventions. This approach was also used to examine in more detail the timing of any load reduction and/or shifting - between hours of the day and days of the week - during longer-term interventions.

2.4.4.2 Longitudinal

Longitudinal (week-by-week) analysis was used to provide a higher-level analysis of the change in consumption over a longer timescale to address two questions. The first, to quantify the time-varying treatment effect delivered by interventions such as the LED lighting upgrades. The second, to examine the effect over time, and persistence, of the data informed and shorter-term interventions. This analysis generally involved using weekly summary data, i.e. the mean 15-minute consumption of households averaged by week. Analysis examined the weekly summary data for changes in consumption measured during the whole day (all-hours) and during the targeted peak period of 4pm to 8pm only (peak-hours). Some interventions required separate measurements to be constructed for weekdays and weekends: for example, in trial period three (TP3) the dynamic pricing rebate intervention targeted only peak-hours on weekdays, therefore the main measurement of consumption tested was the mean consumption recorded for weekdays only.

2.5 Statistical models

For the analysis contained in this report, two statistical techniques are used to investigate the change in consumption attributable to the interventions tested in the second and third trial periods:

'Treatment-only' models: single-variable linear regression modelling to investigate the differences in mean consumption between the treatment and control group;

'Difference-in-differences' (DiD) models: to investigate the change in the differences in mean consumption between treatment group and the control group, and the relationship of these differences to household characteristics.

As noted above, statistical models were run on the consumption data summarised across a number of different temporal scales according to the hypothesised treatment effects.

2.5.1 Treatment only model

To examine and compare the differences in consumption between treatment and control groups, linear regression models were run using the treatment group as independent variable, the equation is as follows:

$$log(y_i)=\alpha+\beta_1\ TreatmentGroup+\epsilon_i$$

Where y_i is mean 15-minute consumption (Wh), α is the intercept (mean control group consumption), β_1 is the coefficient for the treatment group (estimate of the difference between treatment and control) and ϵi is the random error term.

Interpretation of the model results is provided by exponentiating the intercept (α) and coefficient β_1:

exp(α) gives the geometric mean[13] of the control group (intercept) in Wh;

exp(β_1) gives the ratio of the geometric means: treatment group over control group, this is the measurement of group differences reported in the model results.

13 Not to be confused with the arithmetic mean.

The treatment only models were run to examine the differences between the treatment and control groups at a number of temporal scales:

Weekly: to understand how the treatment effect varies across longer timescale, for example how the effect from LED installation varies with the reduction in daylight availability during winter;

Hourly: to understand how the treatment effect varied by hour of the day and/or day of the week, for example according to active occupancy.

2.5.2 Difference in difference model

Difference-in-difference is a commonly used statistical technique used to compare two groups that have been shown to be unequal in terms of the variable of interest (outcome or dependent variable) prior to the intervention; in this case, electricity consumption (log(mean Wh)). The technique relies upon the assumption that although the treatment and control groups are not equal, the trend of the dependent variable over time is the same for both groups (i.e. the parallel trend assumption).

For simple difference-in-difference models, dummy variables were used for time (where Time = 0 for the measurement prior to treatment, and Time = 1 for the measurement after treatment[14]) and for Treatment (Treated = 0 for the control group, Treated = 1 for the treatment group), giving the following equation for the model:

$$log(y_i)=\alpha+\beta_1\ Time+\beta_2\ Treated+\gamma_1\ (Time\times Treated)+\epsilon_i$$

Where:

y_i = mean 15-minute consumption in Watt-hours (Wh)

α = intercept (mean control group consumption at Time = 0, t0)

β_1= coefficient for difference in mean t0 to t1 (trend estimate)

β_2= coefficient for treatment group (group difference estimate)

γ_1= coefficient for treatment effect (difference-in-difference estimate)

ϵ_i = random error

An estimate of the trend in the control group (the difference mean from the reference week to the week under consideration) is given by β_1. The estimate of the difference between the consumption of the control group and the expected consumption in the treatment group is given by β_2. Finally, γ_1 is the difference-in-differences estimate: the difference between the (unobserved) expected and observed consumption in the treatment group, i.e. the treatment effect. These coefficients are illustrated in Figure 1 below.

Figure 1: Illustration of the 'difference-in-difference' linear regression model coefficients

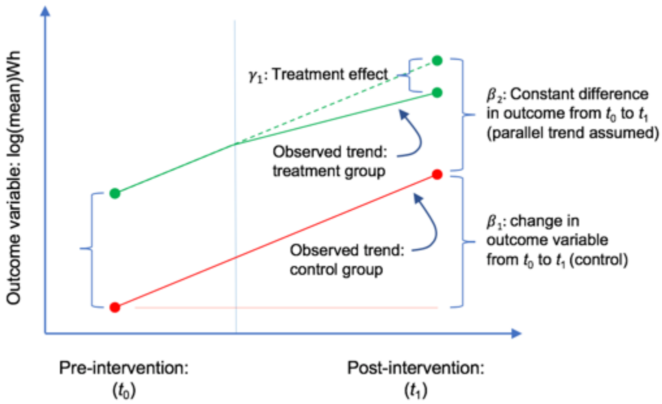

Interpretation of the difference-in-difference coefficient is the ratio of the expected log-mean consumption of the treatment group (given by $\alpha+\beta_1+\beta_2$) and the log-mean measured consumption of the treatment group (given by $\alpha+\beta_1+\beta_2+\gamma_1$), both at Time = 1. For ease of interpretation, the results are also converted to give estimated treatment effects expressed as Watt-hours per hour (Wh/h).

To estimate the effect on the outcome of another independent variable, the following equation applies:

$$log(y_i)=\alpha+\beta_1\ Time+\beta_2\ Treated+\beta_3\ Group+\gamma_1\ (Time\times Treated)+\gamma_2\ (Time\times Group)+\gamma_3\ (Treated\times Group)+\delta_1\ (Time\times Treated\times Group)+\epsilon_i$$

Estimates of the treatment effects observed in households belonging to subgroups of the independent variable in the model are as follows:

For the contrast category, the treatment effect is given by γ_1;

The interaction effect (the estimate of the difference-in-difference-in-difference (DDD) coefficient) is given by δ_1.

14 Separate regression models were run for each of the groups receiving treatment, in each case the treatment group was indicated with the dummy variable Treated = 1. In each analysis, multiple DiD models were run using a common pre-treatment baseline (reference) measurement (Time = 0, household mean consumption prior to treatment). Models were run for each period following the start of the intervention (in each case, the dummy variable Time = 1). To test the parallel trend assumption and assess the impact on the estimated treatment effect of the variation in consumption between the trial groups prior to the intervention, the regression models were also run using a number of consecutive reference weeks.

For other categories of the grouping variable, the treatment effect is $\gamma_1 + \delta_1$;

The influence of a number of additional household characteristics were modelled by the analysis, including the interaction with customer type. This was conducted to examine how the estimated treatment effect, and receptiveness to each measure, varied across different household types.

Note: generally, the linear regression models consider the whole of the treatment group, despite not all of the households in the group receiving treatment. **This analysis therefore gives an estimate of the treatment effect, given the sample population and uptake rate as achieved in each trial.**

2.5.3 Statistical power and confidence intervals

The sample size for the SAVE trials was evaluated using commonly accepted values for statistical power of 0.8.[15] Confidence levels (p-values) of model results are reported where significant and, unless noted otherwise, confidence intervals shown on charts are at the 90% confidence level.

2.5.4 Customer characteristics

As noted above, the evaluation of the SAVE trials was primarily oriented to detecting and quantifying the impact of the measures tested on reducing household demand during peak hours (16:00 to 20:00).

In order to provide customer demand profiles that represent some of the variation in electricity demand across customers with different socio-demographic characteristics, a number of 'Customer Types' were developed (described in Section 3). To examine whether the characteristics used to define the customer types also capture the variation in treatment effects, these characteristics (along with a number of others) were added to the statistical models as interaction terms (see Section 2.5.2 above).

2.6 Trial Impact Analysis

2.6.1 Overview

The interventions tested within the SAVE project were scheduled for three trial periods during 2017 and 2018:

Trial Period 1 (TP1): from 1st January to 31st March 2017;

Trial Period 2 (TP2): from 1st October 2017 to 31st March 2018;

Trial Period 3 (TP3): from 1st October to 31st December 2018.

During the trials each of the participating households were allocated to one of four trial groups: a control group (TG1), and three treatment groups (TG2, TG3 and TG4). Figure 2 shows an overview of the SAVE trials indicating the timing of the scheduled trial periods. The coloured blocks for each treatment group indicate the periods during which evaluation of trial impact was conducted.

Figure 2: SAVE trial schedule showing trial periods and treatment groups

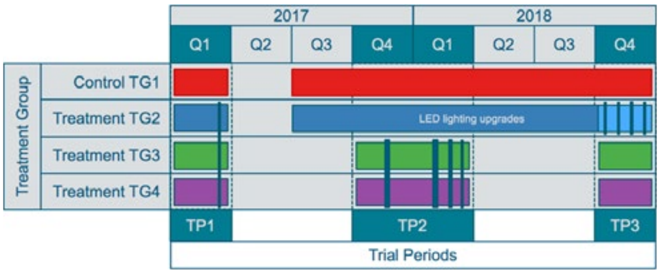

The details of the interventions implemented in the first trial period (TP1) and evaluation of the associated trial impacts are contained in SDRC 2.2. The analysis contained in this report relates to the remainder of the trials implemented following TP1. The sections that follow provide a summary of the evaluation of the interventions tested in the second and third trial periods (TP2 and TP3).

15 Statistical power indicates the probability of a Type II Error (false negative). This should not be confused with confidence interval, which indicates the probability of a Type I Error (false positive).

From July 2017, treatment group 2 (TG2) received roll-out of LED lighting upgrades, installed by fieldwork contractor operatives. In addition, this treatment group was exposed to further data-informed prompts to reduce demand on specific dates during the third trial period (TP3). This period is indicated with a lighter shade in the figure above, with target demand reduction periods ('events') indicated with vertical bars. Evaluation of the impact of the roll-out of LED upgrades is contained in Section 2.6.2 and combined LED upgrades and data informed treatment in Section 2.6.3.

Treatment groups 3 and 4 were exposed to different interventions during the final two trial periods (TP2 and TP3). During the second trial period (TP2) the following approaches were tested in groups 3 and 4:

Trial Group 3 (TG3)- Data informed engagement and price signals

Trial Group 4 (TG4)- Data informed engagement

As the intervention treatments were designed to be comparable, evaluation of these groups was conducted in parallel. A summary is provided in Section 2.6.4 below. During the third trial period, a different approach was applied to the two groups, again designed to be directly comparable. The treatments applied to each group were as follows[16]:

Trial Group 3 (TG3)- Banded price incentives (opt-in enrolment method)

Trial Group 4 (TG4)- Banded price incentives (opt-out enrolment method)

The banded price incentives were also referred to as a dynamic pricing rebate, and are referred to as such in the summary of the evaluation provided in Section 2.6.5.

A summary table containing the results from all interventions trialled during the second and third trial periods is provided in Section 2.6.6.

2.6.2 LED lighting upgrades

In the second trial period, households in trial group two (TG2) received up to ten free LED lightbulbs, installed by fieldwork contractor BMG. A pilot using 100 households was conducted in August 2017 with a full roll-out to the remaining households beginning in October and completing in January 2018. A few remaining households - newly recruited households to mitigate sample attrition - were installed in February 2018. Figure 3 shows the cumulative number of installations for the group against time. The total number of installations completed was 882.

Figure 3: Cumulative LED lightbulb installations within treatment group (number of households)

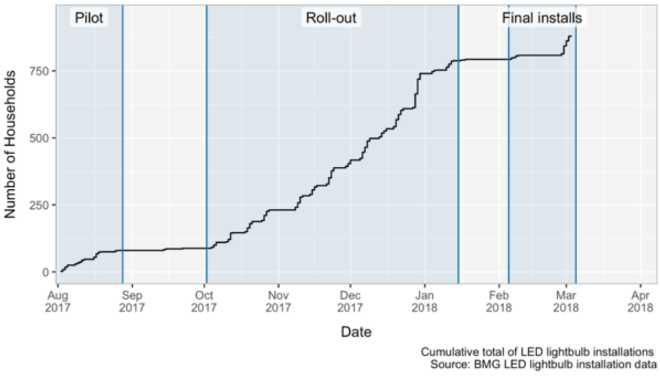

Full details of the approach to the intervention, along with installation outcomes can be found in SDRC 8.3: *LED Trial Report*.[17] The analysis initially examined data from the first winter period following installation (year 1). Due to the installation period running through winter 2017/18, the evaluation was subsequently extended to January 2019 (year 2). During this period the number of households in both treatment and control groups was subject to significant attrition from participants withdrawing from the project and data loss from the Loop electricity monitoring system. Not all of the households receiving LED upgrades therefore have corresponding electricity consumption data. The number of households with electricity consumption data in treatment and control groups declined during the analysis period by approximately 200 in the treatment group and 300 in the control group. This decline had an impact on the statistical power of the analysis.

The consumption of the LED treatment group was compared with that of the control group using the techniques set out in Section 2.4. The analysis of the impact of the LED trial was primarily oriented around a longitudinal approach, supplemented by a more granular approach to examine hour-of-day and day-of-week patterns to observe precise changes in consumption.

16 A full explanation of this change and DNO application is provided in SDRC 8.4 and 8.7, available online at https://www.ssen.co.uk/save/
17 Available online at https://www.ssen.co.uk/save/

2.6.2.1 Group Comparison
First, a comparison of weekly mean 15-minute consumption is provided for all-hours (i.e. considering consumption occurring during any hour of the day) and peak-hours only (only considering consumption occurring between 16:00 and 20:00 hours). LED lightbulbs were installed within a wide time period during the study (shown above in Figure 3 above). In the following figures, households which have and have not had LED bulbs installed are shown as a distinct group and indicated by colour. The distributions of household mean consumption within each group (treatment and control) were tested to determine if consistent differences exist using 'treatment only' regression modelling. The results confirm earlier analysis showing small differences in average consumption between the control and LED treatment groups prior to the intervention. Estimation of the treatment effect therefore required the use of 'difference-in-differences' statistical models.

2.6.2.2 Weekly consumption trends, year 1
The figures in this section show the weekly mean of the mean electricity demand per household (averaged over peak hours), comparing the control group against the treatment group overall (blue line), and for those household where LEDs have been installed (green). Figure 4 shows the weekly mean 15-minute electricity demand in peak hours for the period July 2017 to June 2018.

Figure 4: Weekly mean 15-minute consumption (Wh) by group, peak hours only: July 2017 to June 2018

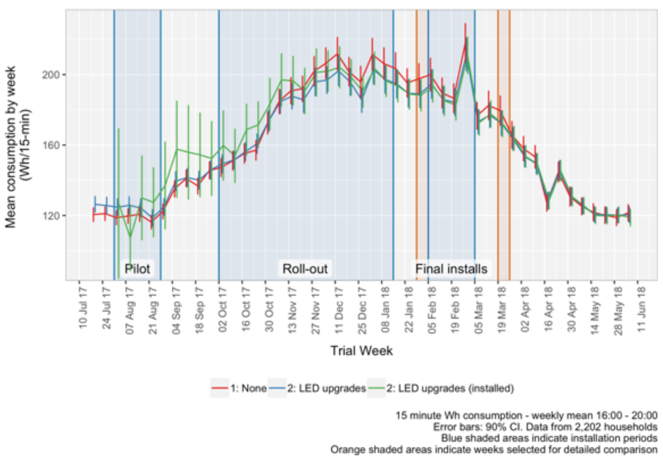

The wide confidence intervals around the 'LED upgrades (installed)' group from July until November 2017 indicate the (initially) small number of households that had received LED upgrades. The figure also clearly shows the increased consumption during the winter months. Although more difficult to observe, it can be seen that the treatment group does reduce consumption relative to the control group moving from marginally above during August to October 2017, to below the control group during November and remaining lower through to March. The orange shaded areas indicate the two weeks selected for more granular analysis.[18] In Figure 5 below, the blue shaded bars highlight weeks during trial period three (TP3) - running from October to December 2018 - where households in the LED intervention group were exposed to additional data-informed behaviour-change treatment. The treatment consisted of postal, online and text messaging asking householders to reduce energy demand during periods ranging from a number of hours to a whole week (peak-hours only). These periods are examined in more detail in Section 2.6.3.

Figure 5: Weekly mean 15-minute consumption (Wh) by group, peak hours only: May 2018 to January 2019

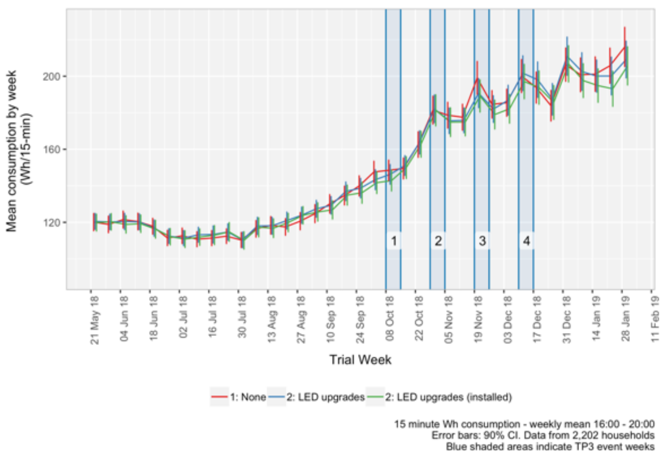

Figure 5 clearly shows mean consumption during peak hours increasing again into the autumn and winter months with colder temperatures and darker evenings increasing demand for heat and artificial lighting. Differences in consumption between treatment and control groups are observed during the winter period. Figure 6 shows the peak hours consumption of the LED treatment group (and subgroups) relative to the control group. Points below the horizontal red line indicate mean (of household mean) consumption in the treatment group below that of the mean of the control group. This chart shows more clearly the movement of the treatment group compared to the control group. Note that from December 2017, the number of participants in the 'LED not installed' group (purple line) is small—so small that sample effects begin to contribute to the high variability of consumption shown.

18 The two weeks selected for detailed analysis were the weeks commencing 29th January 2018 (following the main roll-out of LEDs), and 19th March 2018 (following the final installations). This analysis is presented in SDRC 8.3: LED Trial Report available online at https://www.ssen.co.uk/save/.

Figure 6: Weekly mean peak hours consumption of treatment group (and subgroups) relative to the control group

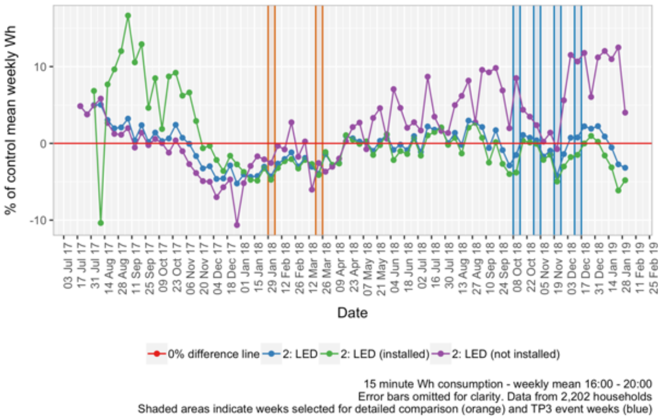

In Figure 6, it can be seen that during July 2017, the period before the intervention, the consumption of the control group is slightly lower than the treatment group (by chance). The difference between groups is small (less than 5%) and variable. As the confidence intervals shown in Figure 5 overlap, it is unlikely that these differences are statistically significant. This was confirmed by running 'treatment only' models as described in Section 2.5.1. The differences observed in the results were consistent with the installation of LED upgrades, the majority of which took place between October and January. The re-convergence of consumption in the treatment group with the control group in March and April shows that the treatment effect (reduced electricity consumption) within peak hours is seasonal and affected by the reduced daylight hours in the winter months. The results of the 'treatment only' models also show that while statistically significant differences in consumption between the groups were not found, small but consistent differences between the groups were present before the roll-out of LED upgrades. Further analysis, using a differences-in-differences approach was therefore conducted to account for these differences and to estimate the treatment effect.

Two sets of difference-in-differences models were run: using consumption from all hours of the day, and peak hours only. Each set of models estimated the difference-in-differences using a number of reference weeks to test the influence of between-group variability on the estimated treatment effect. For ease of interpretation, Figure 7 and Figure 8 show the results of the models expressed as change in consumption, Watt-hours per hour (Wh/h). The lines represent the (geometric) mean change in the weekly mean consumption by treatment group converted to Watt-hours per hour (Wh/h). The figures show the results gained for each contrast week as separate grey lines (with 90 percent confidence intervals overlaid). The mean of the treatment effect estimates (calculated using all contrast weeks) is shown in blue. Using data for all hours of the day, Figure 7 shows that the maximum observed change relative to the control group occurred during the week commencing 1st January. During this maximal week, the mean change in the treatment group (relative to the control) was a reduction of 31 Watts per household (90% confidence interval = 2 to -61 Watts). Equivalent to 733 Watt-hours per household per day and 5.1 kWh per household per week.

Figure 7: Change in hourly mean 15-minute household consumption, converted to constant power equivalent in Watts *(all hours)*

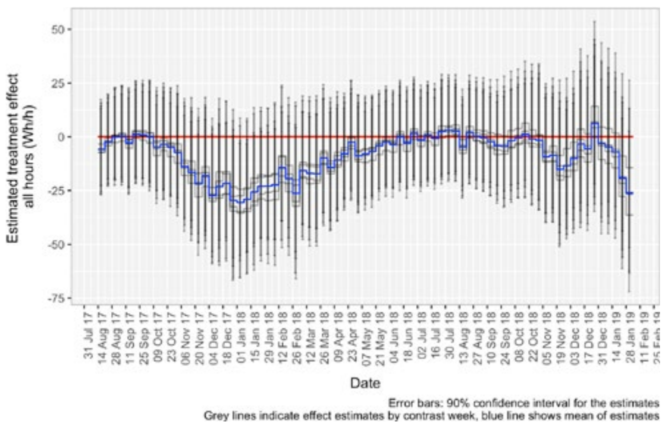

Figure 8 shows that during the targeted peak hours (16:00 to 20:00), the maximum observed change (relative to the control group) occurs during the weeks commencing 25th December 2017 to 15th January 2018.

Figure 8: Change in hourly mean 15-minute household consumption, converted to constant power equivalent in Watts (*peak hours*)

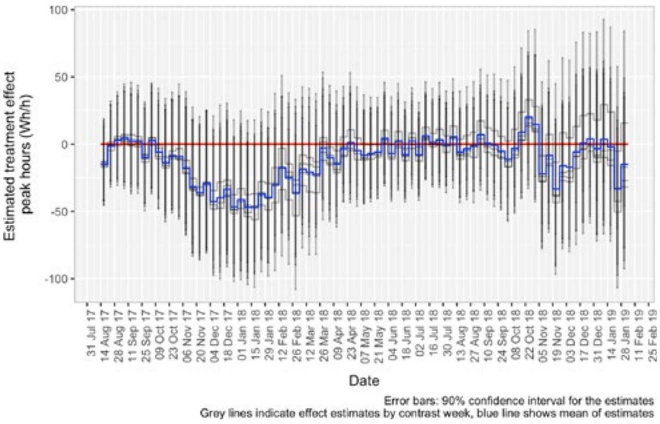

During the maximal week, the mean change in consumption in the treatment group over the peak hours (4 to 8 pm) was equivalent to a reduction of 47 Watts per household (90% confidence interval between 8 and -97 Watts). This is equivalent to reduced consumption over the 4-hour peak period of 186 Watt-hours per household per day.[19]

2.6.2.3 Household characteristics and treatment effect

One of the functions of the Customer Model is to understand customer receptiveness to measures and thus represent the variability in treatment effect from interventions across different household types.

To investigate these relationships a number of DiD regression models were run using household characteristics as interaction terms. The variables currently used to define 'customer types'[20], were modelled to examine the degree of variation in treatment effect captured. None of the interaction terms modelled for the LED intervention were found to be statistically significant and the results below should be treated as indicative only. In summary, when controlling for household size, dwelling size and heating fuel, the key findings from the models relevant to the customer types are:

- Household size
  - the greatest treatment effect was observed in households with one occupant[21]
- Bedrooms
  - the greatest treatment effect was observed in the largest homes (5+ bedrooms)
  - the treatment effect increases with size of dwellings (no. of bedrooms)
- Heat source
  - greater effect was observed in electrically-heated households and households with 'other' primary heat source than in gas-heated households

Other household characteristics

To examine the variation in treatment effect with other household characteristics, a selection of other variables were modelled as interactions. For detailed analysis results refer to SDRC 8.3 *LED Trial Report*.[22]

2.6.3 LED lighting upgrade and data-informed interventions

This section provides a summary of the analysis of the impact of the data-informed event interventions trialled during the third SAVE project trial period (TP3): October 2018 to December 2018. During this period, householders in intervention group 2 were prompted - through a variety of printed (via post) and electronic (via email, text and online) materials - to reduce their electricity consumption during peak hours (between 16:00 and 20:00). As detailed above, a large proportion of the participating households within this intervention group had previously also received LED lighting upgrades, installed between July 2017 and March 2018. The targeted periods for reduced consumption during the TP2 events were as follows:

- Event 1: peak hours during Wednesday 10th October 2018;
- Event 2: peak hours each day for weekdays during the week commencing 29th October 2018;
- Event 3: peak hours each day for weekdays during the week commencing 19th November 2018;
- Event 4: peak hours during Thursday 13th December 2018.

---

19 Due to the length of time taken to roll-out of the LED upgrades, the evaluation of the intervention impact was extended to the end of the monitoring period (January 2019). In the period 1st October to 31st December 2018, this trial group was exposed to a series of additional data-informed interventions, and also suffered significant attrition in sample size, reducing from approximately 800 to 550 households in the treatment group. These factors are likely to have contributed to the increased variability in estimated treatment effect during this period and the wider confidence intervals surrounding the estimates.

20 The variables used to define customer types are: household size (number of people), dwelling size (number of bedrooms) and primary heating source (fuel).

21 While this may seem counter-intuitive, the analysis could not disentangle the relationship between the physical intervention and resulting energy behaviours. Uncertainty in the observed effect also means that firm conclusions about the interactions cannot be drawn.

22 Available online at https://www.ssen.co.uk/save/

The target given to trial participants during the third trial period (TP3) was to reduce electricity consumption during the network peak, between 16:00 and 20:00 hours (4 to 8pm). As with the analysis of the intervention impacts for the first and second trial periods, the analysis examined household consumption during three time periods:

- Pre-peak: the four hours prior to the peak period (12:00 to 16:00);
- Peak: the four hours of the peak period (16:00 to 20:00);
- Post peak: the four hours following the peak period (20:00 to 00:00).

Due to the length of these periods, and the possibility of changes in consumption occurring within smaller time intervals or outside these times, analysis was conducted using consumption summarised to hourly periods.

2.6.3.1.1 Data summary

For each event, three weeks of consumption data was used: the week preceding the event, the week of the event and the week after. This allows any consumption shifted away from the event week (i.e. to hours and days before and after the event period) to be measured. The 15-minute consumption data was summarised before the analysis was performed, taking the weekly mean for each household for each hour of the day. Separate values were calculated for weekdays and weekends. Where required, weekdays were further divided into days during an active intervention, and 'normal' weekdays (when no intervention was active).

2.6.3.1.2 Event 1

Event 1: peak hours during Wednesday 10th October 2018. The analysis used three weeks of data from 1st October 2018 to 21st October 2018. Figure 9 shows a comparison of the group mean 15-minute consumption by hour of the day, day type (i.e. weekday and weekend) and by week for the period, curtailed to show only the period 12:00 (noon) to 24:00 (midnight). The trial group receiving treatment is compared to the control (Group 1). From inspection, the mean consumption profiles for each group are very closely matched with only small differences apparent between the treatment and control groups observed.

Figure 9: Hourly mean of 15-minute household consumption (Wh) by intervention group: event 1, pre-peak, peak and post-peak periods

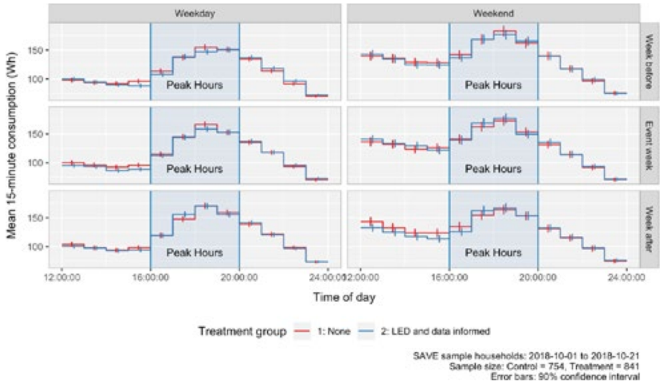

As with the longitudinal analysis, 'treatment only' models were used to test for differences between the treatment and control group. Two timescales were used as changes in energy consuming activities potentially occur at different intervals determined by the temporal patterning of routines in households i.e. some households may perform activities on certain days of the week (weekly schedule) and some are more flexible from day-to-day. For the week-to-week comparison, differences in consumption were tested on three consecutive Wednesdays (the week before, the event week and the week following). For the day-to-day comparison, differences were tested for the day prior to the event (Tuesday), the day targeted in the event (Wednesday) and the day following the event (Thursday). Only marginal differences were observed for each time period, with no evidence of any significant shifting to non-trial weekdays or weekends.

Difference-in-difference models were also run for the same time periods with the reference (pre-intervention) time period, t_0, set as the week or day prior to the event week, or day.

For short events, it is assumed that consumption activities avoided during the targeted time are more likely to be shifted to the hours or day immediately before and/or after the event day. Therefore, for Event 1, the difference-in-differences analysis was performed for the day-to-day comparison.

Figure 10 shows that, compared to the day prior to the event, the measured consumption in the treatment group during the event day is marginally higher than predicted for all time periods (averaged for all-hours, pre-peak, peak and post-peak).

Figure 10: Estimated day-to-day treatment effect as ratio of observed to expected 15-minute consumption in the treatment group by hour of day: Event 1

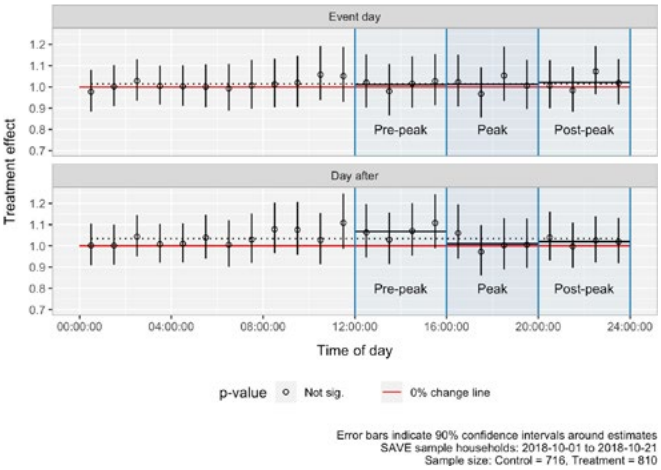

Models were also run to test for a week-to-week treatment effect, where weekly schedules might dominate consumption patterns. The observed marginal increase in consumption during peak hours on the event day of 2.1 percent is not consistent with the experiment hypothesis. Using a significance level of 90 percent, the results are not significant.

2.6.3.1.3 Event 2

Event 2 targeted peak-hours during the all weekdays of the week commencing 29th October. 15-minute consumption data for the period 22nd October to 11th November 2018 was used for the analysis of Event 2. Only the week-to-week models were run to estimate the treatment effects for this event period. The results show only a marginal reduction (2.2%) in consumption during weekday peak hours in the event week[23]. A greater reduction of 6.2% was observed during peak-hours in the week following the event. Similarly, the average reductions observed across all hours were greater during the week following the event. This can be explained as an effect due to the LED lighting upgrades: British Summertime ended during the weekend of the event week, with the clocks moving back one hour and reducing available daylight.[24]

2.6.3.1.4 Event 3

Event 3: peak hours each day for weekdays during the week commencing 19th November 2018. 15-minute consumption data for the period 12th November to 2nd December 2018 was used for the analysis of Event 3. Inspection revealed that the consumption in the treatment group is slightly lower than the control group during the middle of the day and into the evening peak period during the event week. Clear differences in consumption between the treatment and control groups during the peak hours of the event week, specifically the hours 18:00 to 19:00.

The treatment only models confirmed the observations made from the consumption data, with the average consumption during weekday peak-hours approximately 5% lower in the treatment group although the differences between groups were generally shown not to be significant at the 90% confidence level. The results for difference-in-differences show an average reduction of 2.9 percent during the peak-hours and an average reduction in consumption across all hours of 1.6 percent. No evidence of shifting consumption to outside of peak hours was found, as reductions were also observed during the pre-peak and post-peak periods of the event week.

2.6.3.1.5 Event 4

Event 4 targeted peak hours during one day only: Thursday 13th December 2018. 15-minute consumption data for the period 3rd December to 23rd December 2018 was used for the analysis of the event. For the week-to-week comparison of Event 4, differences in consumption were tested on three consecutive Thursdays (the week before, the event week and the week following). For the day-to-day comparison, differences were tested for the day prior to the event (Wednesday), the day targeted in the event (Thursday) and the day following the event (Friday).

Over the three consecutive Thursdays, the mean consumption of the treatment group increased relative to that of the control group. There was no evidence from the models that the event produced a reduction in consumption in the treatment group, although increased consumption was observed in the post-peak hours of the event day - an average increase of approximately 8 percent - consistent with the hypothesis that energy-using activities are being deferred until after the peak hours. Further detailed analysis from the evaluation of the event-based interventions trialed in TP3 is provided in SDRC 8.4 and 8.7.

23 calculated as an average of the estimated hourly difference-in-difference coefficients.

24 As the reference weeks used in the analysis of the subsequent events occur after the clock change, the clock change has no impact.

2.6.3.2 Summary

In this trial period and within this treatment group, it is difficult to disentangle the interactions of the existing LED intervention with the data-informed 'event' interventions. The intervention took place during a period moving into winter with darker evenings where a reduction in consumption relative to the control group would be expected even in the absence of any additional intervention, due to the pre-existing LED upgrades in this group. There is little evidence in the short-term analysis to support the hypothesis that the data-informed interventions produced significant additional reductions during peak hours. The largest estimated treatment effect was observed during Event 3, a 2.9% reduction. It should be noted that this result was not statistically significant at the 90% confidence level and the contribution of each individual intervention (LED upgrades and data-informed events) is unknown. Taking these results in the wider context and comparing to those obtained in the data-informed engagement during trial period 2 (reported in Section 2.6.4 below), suggests that some pre-engagement work may be useful in mobilising customer response to event-based interventions.

2.6.4 Data informed and price signals

This section provides analysis of the impact of the interventions trialled within trial groups 3 and 4 (TG3 and TG4) during the second SAVE project trial period (TP2): October 2017 to March 2018. The treatments applied to these groups were as follows:

TG3: data informed and price signals

TG4: data informed only

During the trial, householders in intervention groups 3 and 4 were prompted - through a variety of printed (via post) and electronic (via email and online) materials - to reduce their electricity consumption during peak hours (between 16:00 and 20:00). In addition, intervention group 3 received price signalling (i.e. vouchers/cash for meeting specified percentage reduction in peak hours electricity consumption).[25]

The targeted periods for reduced consumption during the TP2 challenges were as follows:

Challenge 1: peak hours each weekday during the week commencing 20th November 2017;

Challenge 2: peak hours each weekday during the week commencing 29th January 2017;

Challenge 3: peak hours during Tuesday 6th and Wednesday 7th March 2018;

Challenge 4: peak hours (17:00 to 19:00 only) during Tuesday 20th March 2018.

To understand the impacts of the interventions tested during this trial period, a longitudinal analysis of weekly mean consumption during peak hours was undertaken alongside a more detailed analysis at the hourly timescale for each of the challenges.

2.6.4.1.1 Longitudinal analysis

As with the analysis of the LED treatment group provided above, initial analysis used weekly summaries of the 15-minute electricity consumption data to evaluate the impact of the intervention across the full extent of the trial period. Figure 11 shows the mean of household mean peak-hours consumption for each treatment group relative to the control group, i.e. the percentage difference in mean consumption of each treatment group from the control group. Also shown in the figure is the weekly heating degree-days for the trial period (black line), to indicate changes in heating requirements. It is observed that prior to the start of trial period 2 (pre-October 2017), mean consumption in treatment group 4 is up to approximately 6% lower than the control (purple line), clearly showing the requirement to use difference-in-difference models to account for the pre-treatment asymmetry between groups. Inspection reveals no obvious correlation between the observed differences between treatment groups and heating demand.

25 For full details of the intervention schedule and materials, refer to SDRC 8.4 & 8.7 Data Informed Engagement and Price Signals Report, available online at https://www.ssen.co.uk/save/.

Figure 11: Weekly mean 15-minute peak-hours consumption of treatment groups relative to control group, data informed and price signals interventions, trial period 2

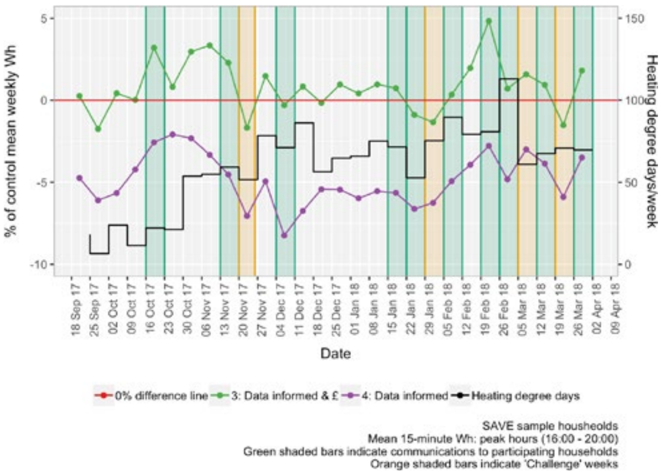

Figure 12: Estimated treatment effects by intervention group as mean change in consumption in peak-hours, trial period 2

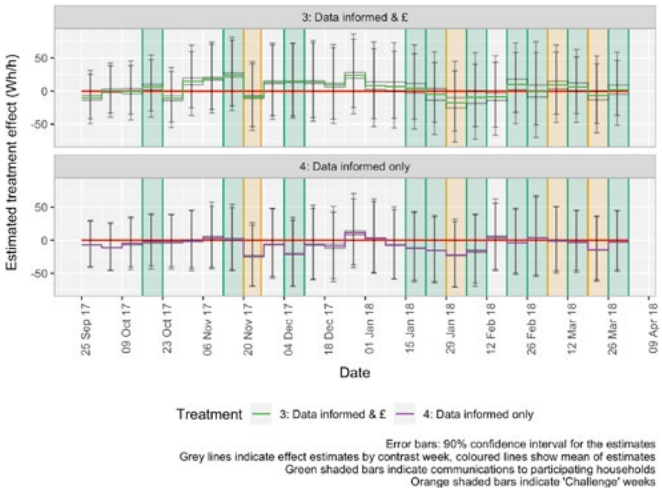

A set of models were tested to estimate the differences between the expected and observed consumption in the treatment groups. The dependent variable was log-mean 15-minute consumption measured during weekday peak-hours only (the hours targeted by the interventions tariff). To account for variation in estimated treatment effects due to non-parallel trend in treatment and control, models were run using multiple contrast weeks.

Figure 12 shows the results from the regression models for the 'Engagement + £' group (treatment group 3) and 'Engagement only' group (treatment group 4) in the top and bottom panels respectively. The estimated treatment effects (difference-in-difference estimates) are plotted against the vertical axis and show the mean difference between expected and observed consumption in the treatment group during weekday peak-hours (16:00 - 20:00) by calendar week. The error bars represent the 90% confidence intervals around the estimated values and are overlaid.

From Figure 12 the following observations are made regarding the responses of the treatment groups to the intervention using the longitudinal analysis:

The maximum estimated treatment effect in group 3 (data-informed and price signals) occurred during the 2nd Challenge week, with a mean load reduction of -18 Wh/h (90% CI -69 to 38 Wh/h);

The maximum estimated treatment effect in group 4 (data-informed only) occurred during the 1st Challenge week, with a mean load reduction of -24 Wh/h (90% CI -69 to 25 Wh/h);

Treatment group 3 appear to increase consumption in the 3 weeks prior to Event 1, in contrast to group 4 where there is very little change;

The cumulative effect is that impact during the 1st event week is very small for group 3;

Following the 1st event, consumption in treatment group 3 increases consumption to above the expected level, while consumption in group 4 is below the expected level, indicating some persistence in the treatment effect within this group only;

Both treatment groups show consistent reduction week-on-week through January, with a maximum effect observed during the week of the 2nd event.

The size of the confidence intervals around the effect estimates should be noted and none of the weekly results shown above are significant at the 90% confidence level. The effect sizes estimated by the longitudinal analysis provide an analysis of the treatment effect using a reference point at a time prior to the start of the trial period. The estimated effects therefore give a cumulative effect for any week during the trial. If a trial group increased consumption relative to the control group in week one, then decreased consumption by the same amount in week two, if all else remains equal the estimated (cumulative) treatment effect for week two will be zero. Using a week-to-week comparison (and a reference point the week prior, as in the 'short term' analysis below) however, would show the treatment effect as a reduction in consumption. The two reference points (and associated estimates) are useful in understanding both the longer- and shorter-term responses the intervention methods trialed and impact on the network.

2.6.4.1.2 Short term analysis

More detailed analysis of response to the challenges was modelled using hourly summary data. The models measured week-to-week and day-to-day changes in consumption relative to the control group using the difference-in-differences technique. These models do not capture the longer-term effects, for example the increased consumption in treatment group 3 in the weeks prior to Challenge 1. The analysis does however provide further information about the timing of the response to the interventions: i.e. in which hour of the day the maximum load reductions occur, and the timing of any load shifting. Figure 13 presents the results of shorter-term analysis of changes in consumption from the week prior, to the week of Challenge 1. It shows that the maximum load reduction in both treatment groups occurs between 5 and 6pm and 'spill-over' of load reduction to the hours either side of the peak period is observed in treatment group 3 but not group 4. Marginal increases in consumption immediately after the peak period in group 4 indicate that shifting of peak loads may have occurred, although the magnitude of the effect indicates that the reduction in peak hours is primarily load reduction for the targeted weekdays. As noted above, while this analysis quantifies the shift in consumption from the week prior to the week of the Challenge, it presents a false picture of the treatment effect for the data informed and price signal group (treatment group 3) as it does not account for the increased consumption over the weeks before the Challenge (see Figure 12).

Figure 13: Challenge 1 estimated treatment effects by intervention group as mean change in hourly consumption

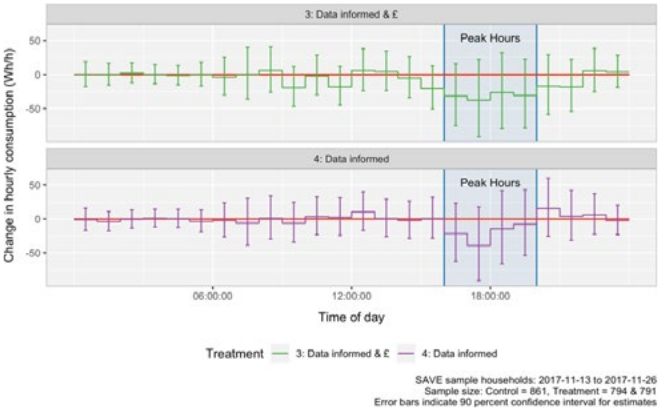

The hourly granularity of the short-term impact analysis also captured the treatment effects observed during the reduced hours targeted during Challenge 4 that were obscured in the longitudinal analysis.[26] In this challenge, only the central hours of the peak were targeted: 17:00 to 19:00. The hourly results show that a much greater reduction was observed between 17:00 to 19:00 hours for group 3 (with increased consumption in the first and fourth hour of the peak period). A small increase in consumption was also observed outside of the targeted hours. The average reduction over these two hours was approximately 7 percent (35 Wh/h). For group 4 the average reduction was approximately 3 percent (16 Wh/h). In the week-to-week results, the estimation of the mean treatment effect across peak hours is smaller for both treatment groups and is under 1 percent. As with the other Challenges, the results were not statistically significant with high levels of uncertainty around the treatment effect estimates.

2.6.4.1.3 Household characteristics

As with the other interventions, models were run to examine the relationship of household characteristics to the magnitude of treatment effect. The model results show that while differences were found in the observed treatment effects between groups, no statistically significant interactions were observed for households grouped by the variables tested. The variability in consumption within groups provides wide confidence intervals around the estimated interaction coefficients.

26 Similar analysis was conducted for Events 2 and 3 and is provided in SDRC 8.4/8.7 Data Informed Engagement and Price Signals Report, available online at https://www.ssen.co.uk/save/.

Analysis conducted using appliance ownership found that when controlling for household size, dwelling size and primary heat source, that households owning high power-rating appliances (including EVs, dishwashers, tumble dryers, and plug-in heaters) have higher electricity consumption during peak hours. The results were statistically significant for tumble dryers and dishwashers at the 99% confidence level.

In terms of modelling the interaction of household characteristics with *treatment effect* within the Customer Model, for the treatments examined in this analysis:

- household size and heat source provide more useful predictors of the level of response than dwelling size (measured using number of bedrooms);
- other household characteristics (such as income and level of qualifications) and appliance ownership (dishwashers, tumble dryers and plug-in heaters) may also interact with treatment effects, however the analysis conducted for this report does not provide evidence with an acceptable level of confidence.

2.6.4.1.4 Summary

The findings from trial period 2 highlight several phenomena that are relevant to the wider roll-out of demand response interventions and network planning:

1. Where customers are incentivised to reduce consumption relative to a 'baseline' measured from historical consumption, caution is required. Households may increase consumption during the baseline measurement period risking higher network loads during peak times and also cancelling-out any demand reduction during the targeted period.[27]
2. Load reductions from data-informed only treatments outperformed those in the group where financial incentives were offered. This confirms findings from other trials that financial incentives do not consistently provide demand reductions.

2.5.6 Banded price incentives (dynamic pricing rebate)

During the third trial period (TP3), trial groups 3 and 4 formed two groups exposed to a trial of dynamic pricing rebate tariffs. The two treatment groups experienced differential recruitment conditions as follows:

TG3: dynamic pricing rebate, opt-in enrolment. By default, households in this treatment group received no treatment and were invited to opt-in to participate;

TG4: dynamic pricing rebate, opt-out enrolment. By default, households in this treatment group were enrolled to receive treatment by default, and invited to opt-out if they did not want to participate.

Both groups were offered a financial incentive to reduce consumption to below a pre-defined threshold during weekday peak hours (16:00 to 20:00).[28] The rebate incentive level was set at 10 pence per hour for the first six weeks of the trial period (capped at £20 per household), and increased to 30 pence per hour for the remainder of the period (with cap for maximum incentive raised to £50 per household). The maximum rebate each household could earn was therefore £2 per week for the initial 'low' incentive period, increasing to £6 per week. As the rebate intervention ran for a period of three months, a longitudinal approach to evaluating the impact was used, with the treatment effect measured as the average change in 15-minute consumption during weekday peak-hours.

Household electricity consumption (Loop) data for September through to the end of December 2018 was used to evaluate this intervention and the 15-minute consumption data was summarised to provide weekly mean demand during weekday peak-hours for each household.

The number of households within each treatment group declined across the trial period due to attrition of the sample. Tables showing the sample size for each group are shown in Appendix A.1.1.

Comparing the mean consumption of those participating and not-participating within each treatment sub-group revealed evidence of self-selection: households joining the trial by opting-in (participating households in treatment group 3) were observed to have lower consumption, on average, than those not opting-in (see Appendix A.1.2 for more details). The implications of this finding are relevant for a network operator rolling out incentives using similar recruitment methods: using an opt-in enrolment method may not achieve large participation rates, particularly in areas dominated by high-demand customers.

27 See Gyamfi et al., (2013), Residential peak electricity demand response—Highlights of some behavioural issues. Renewable and Sustainable Energy Reviews 25, 71–77. https://doi.org/10.1016/j.rser.2013.04.006

28 Low, medium or high threshold values were informed by prior measurement of each household's consumption or characteristics where no consumption data was available.

In Figure 14, mean peak-hours consumption for each treatment group is presented relative to the control group and alongside the weekly heating degree-days for the trial period. It is observed that prior to the start of the trial period, mean consumption in treatment group 3 (opt-in) is up to 2.5% higher than the control group and mean consumption in treatment group 4 (opt-out) is up to 5% lower than the control group. This Figure clearly shows the requirement to use difference-in-difference models to account for the pre-treatment asymmetry between groups. Inspection reveals no obvious correlation between the observed differences between treatment groups and heating demand.

Figure 14: Weekly mean 15-minute peak-hours consumption of treatment groups relative to control group, dynamic pricing rebate intervention

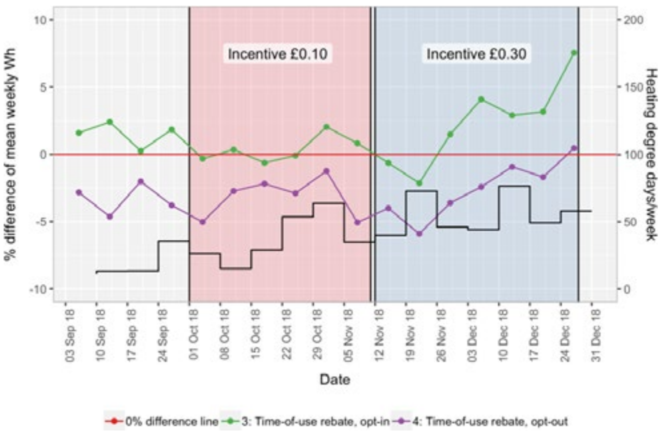

To estimate the treatment effects, two sets of difference-in-differences models were used. The first set used the entire sample of households within each treatment group to estimate the effect while embedding the enrolment rate, i.e. accounting for the differential rates for each enrolment method.[29] These results provide estimates of the expected effect of rolling out a similar scheme (under the same conditions and within a similar population). The second set of models used only those households within each group that were participating.[30] These results provide an understanding of how the participants responses (treatment effect) vary with the enrolment method used. For the first half of the trial period (with low incentive rate - £0.10/hour), the estimated mean treatment effect in the 'opt-in' group is generally larger than that of the 'opt-out' group. The treatment effect is also observed to be more consistent in the 'opt-in' group compared to the 'opt-out' group, particularly during the low incentive period. This finding indicates that participants opting-in to a dynamic pricing rebate scheme may offer greater reduction that those opting-out (particularly where low incentive rates are offered). However, as the participation rate within the 'opt-in' group was much lower, the estimated effects are higher in the 'opt-out' group when accounting for enrolment rate. In contrast, the treatment effect in the 'opt-out' group increases markedly in the final week of the low incentive period (week commencing 5th November) and a peak during the week commencing 19th November, before reducing gradually to the end of the trial period. This finding indicates that the postal communication of the increased incentive rate caused an increase in engagement the week before the rate was raised (week commencing 5th November) and translated into an increase in treatment effect, most pronounced in the opt-out group.

The estimated treatment effects are shown in Figure 15. The treatment effects are presented as Watt-hour reductions and calculated for each **entire treatment group** (including both participating and opted-out households), the **participation rate is therefore embedded**.

Figure 15: Estimated treatment effects of dynamic pricing rebate as mean change in consumption in peak-hours, participation rate embedded

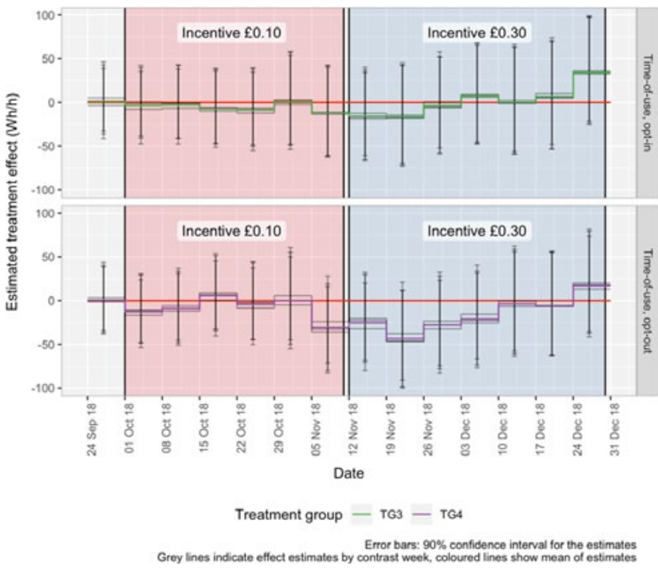

The maximum estimated load reduction for both treatment groups was observed during the week commencing 19th November 2018 – the second week of the high incentive period – with mean effect sizes as follows:

Treatment group 3 maximum -17 Wh/h (90% CI, -71 to 43);

Treatment group 4 maximum -44 Wh/h (90% CI, -97 to 15).

29 The whole group sample sizes in these models were: Control = 709, Opt-in = 662, Opt-out = 657 households.

30 The sample sizes for participating households only were: Control = 709, Opt-in = 297, Opt-out = 633 households.

Separate models were created to examine only those households within each trial group that participated (i.e. excluding households not opting-in or choosing to opt-out), Figure 16 shows the resulting treatment effect estimates and provides a comparison of the recruitment methods. The figure shows that excluding non-participating households increases the estimated treatment effect within the opt-in group (group 3)but not the opt-out group (group 4). This is due to a much greater number of non-participating, and non-engaged households in the former group. The figure also shows that a more consistent treatment effect is observed in the opt-in group, that may be attributed to the self-selection within this group and indicating that these households were more engaged throughout the duration of the trial. However, the maximum observed effect is lower in the opt-in group that the opt-out group for the higher incentive period, indicating that participating households are not primarily motivated by the financial incentive.

Figure 16: Estimated treatment effects of dynamic pricing rebate as mean change in consumption in peak-hours, *participating households only (participation rate excluded)*

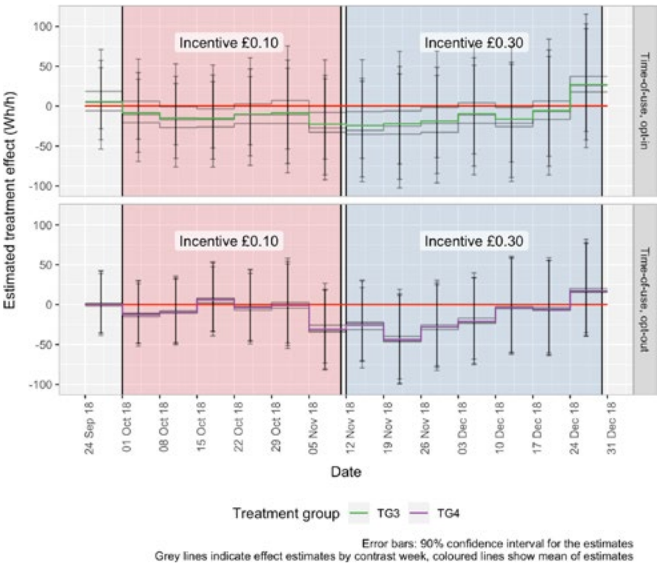

Again, the maximum estimated load reduction occurred in the higher incentive period and for each treatment group was as follows:

Treatment group 3 maximum -24 Wh/h (90% CI, -83 to 41), observed week commencing 12th November 2018;

Treatment group 4 maximum -44 Wh/h (90% CI, -97 to 15), observed week commencing 19th November 2018.

2.6.5.1 Household characteristics

Additional models were created to examine the interaction of the customer type variables with treatment effect. The models used data for only participating households within each treatment group. Controlling for household size, dwelling size and primary heating fuel, the interaction models show the following effects (on average) by sub-group:

- Household size (people) - one and two-person households show the strongest observed response, followed by households with four or more occupants. Three-person households exhibit the smallest response to the intervention;
- Dwelling size (bedrooms) - for both treatment groups, households occupying larger dwellings (3 and 4+ beds) show the greatest average estimated treatment effect. For the opt-in group households, those with 2 and fewer bedrooms saw a smaller effect on average but this interaction was less prevalent in the opt-out group;
- Primary heating fuel - in both treatment groups, the strongest response was observed in households heated by 'other' fuels (not gas or electricity), followed by households primarily heated electrically. On average, the smallest effects were observed in gas-heated households.

Models were also created to examine the interaction of other household characteristics (such as employment status, presence of children, ethnicity, tenure) with the observed treatment effects. A larger effect was estimated in households without children, although this interaction was less prevalent in the 'opt-out' treatment group and a significant interaction was observed with ethnicity. Due to the small samples within sub-groups, these results are indicative only and should not be used to infer the level of treatment effect expected within different groups of consumers. Very large confidence intervals surround many of the estimated effects when examining sub-groups.[31] However, the SAVE trial results indicate that the levels of demand response may vary across different groups of households and that extending the customer model typology with these characteristics would capture more variability in demand profiles and the differential response to interventions aimed at reducing peak load.

31 For more detailed results refer to SDRC 8.4 & 8.7 Data Informed Engagement and Price Signals Report, available online at https://www.ssen.co.uk/save/.

2.6.6 Summary

In order to meet project objectives outlined at the head of this section, the evaluation of the impact of trial interventions was conducted using a number of statistical techniques and along a variety of time intervals. In summary:

Analysis of the direct measurement of household energy consumption has provided for the detection of any 'subconscious' behaviour change that occurred as a response to treatments;

The analysis was conducted along a number of time scales to distinguish between novelty effects and longer-term change - from longitudinal analysis of average weekly treatment effects, down to more granular day-to-day analysis of average effects at hourly intervals;

The ability to detect statistically significant effects was provided by the use of robust statistical methods to test key measurements, with the results presented with appropriate confidence intervals;

The statistical modelling of groups defined by household characteristics was used to examine which customers were most receptive to intervention treatments.

While the size of the recruited sample and treatment groups was designed for the hypothesized effect sizes, attrition of the sample during the trials and the small size of observed treatment effects mean that for many of the interventions, statistically significant results have not been obtained. Small samples of households within sub-groups have also limited the ability of the analysis to identify groups with stronger and weaker response to the interventions with an acceptable level of confidence.

Through implementation of a large-scale randomized control trial, the SAVE trials have provided estimates of the scale of peak-hours demand reduction that can be expected from a range of DNO-led interventions. The LED trial provided empirical data supporting a consistent mean demand reduction during the peak hours from a roll out of installed LED lighting upgrades. In addition to the theoretical demand reduction, the intervention **provided evidence of the actual load reduction obtained from the upgraded lighting when *in-use***. The effect was seasonal with greatest load reduction observed during mid-winter and minimum daylight availability. At the maximum, the estimated average observed load reduction was equivalent to 47 W per household (90% CI, -96 to 7 W) across the peak hours.

Table 2 below contains a summary of the estimated treatment effects for trial group 2 (TG2) during the second and third trial periods: the *LED* and *LED* plus data informed interventions. Values in bold show the average estimated percentage difference from the expected consumption in the treatment group during peak hours. The Watt-hour values provide estimates of the maximum effect size and the time period within which the observation was made. The table shows that the LED intervention provided the largest peak-hours load reduction in this treatment group.

Table 2: Summary of observed treatment effects, treatment group 2

Trial period, group, intervention		Element	Effect size (peak hours)		Max. impact (Wh) (90% CI) and period observed
			%	Wh/h	
TP2	TG2: LED lighting upgrades	Max average weekly treatment effect (all hours)	**-7.0**	-31	-31 (-61 to 2) w/c 1st Jan 2018
		Max average weekly treatment effect (peak hours)	**-7.0**	-47	-47 (-96 to 7) w/c 15th Jan 2018
TP3	TG2: LED and data informed		**+2.1***	+8*	-23 (-70 to 30) 6-7pm
			-2.2*	-12*	-20 (-78 to 44) 5-6pm
			-2.9*	-16*	-22 (-74 to 35) 7-8pm
			-0.2*	-1*	-11 (-62 to 48) 4-5pm

Notes:
* Value calculated as mean of hourly effect estimates across targeted peak hours
# Challenge 4 targeted the central peak hours only (17:00 to 19:00)

Table 3 below contains a summary of the estimated treatment effects for trial groups 3 and 4 during the second trial period: the data informed and price signal treatments. Again, values in bold show the average estimated percentage difference from the expected consumption in the treatment group during peak hours. The table highlights the difference in treatment effects estimated by the longitudinal and short-term methods.

Table 3: Summary of observed treatment effects, data informed and price signal interventions, trial period 2

Element	TG3: Data informed + £			TG4: Data informed only		
	%	Wh/h	Max. impact (90% CI) and period observed	%	Wh/h	Max. impact (90% CI) and period observed
Longitudinal (max. avg. weekly effect)	**-2.8**	-18	-18 (-69 to 38) Challenge 2	**-4.2**	-24	-24 (-69 to 25) Challenge 1
Short-term						
Challenge 1:	**-5.5***	**-32***	-38 (-91 to 22) 5-6pm	**-3.8***	**-21***	-39 (-91 to 18) 5-6pm
Challenge 2:	**-0.8***	**-4***	-9 (-70 to 58) 6-7pm	**-1.3***	**-7***	-26 (-76 to 30) 7-8pm
Challenge 3:	**+3.0***	**+13***	-20 (-83 to 51) 6-7pm	**+2.4***	**+10***	-8 (-69 to 60) 6-7pm
Challenge 4:#	**-7.0***	**-35***	-40 (-90 to 18) 5-6pm	**-3.0***	**-16***	-19 (-77 to 46) 6-7pm

* Value calculated as mean of hourly effect estimates across targeted peak hours
# Challenge 4 targeted the central peak hours only (17:00 to 19:00)

Finally, while the sample sizes of the treatment groups were not of a sufficient size to detect statistically significant treatment effects, the dynamic pricing trials provided observations consistent with participating households responding to price signals. Table 4 shows the estimated treatment effects for treatment groups 3 and 4 during the third trial period: the dynamic pricing rebate intervention. Comparing across the results tables it can be seen that the estimated effects were greatest for the LED and dynamic pricing rebate interventions, with a similar percentage effect size of approximately 7 percent.

Table 4: Summary of observed treatment effects, dynamic pricing rebate intervention, trial period 3

Incentive	TG3: Opt-in recruitment			TG4: Opt-out recruitment		
	%	Wh/h (90% CI)	Timing of max. effect*	%	Wh/h (90% CI)	Timing of max. effect
Low incentive	**-2.1**	-12 (-62 to 42)	w/c 05-11-2018	**-5.5**	-31 (-78 to 21)	w/c 05-11-2018
High incentive	**-2.6**	-17 (-71 to 43)	w/c 19-11-2018	**-7.1**	-44 (-97 to 15)	w/c 19-11-2018

* Timing of maximum effect refers to Wh effects. In percentage terms, the maximum value observed during the low incentive period for the opt-in group was the week commencing 12th November 2018.

3

CUSTOMER TYPES

3.1 Overview

This section provides details of the development of the Customer Model carried out following the initial implementation (detailed in SDRC2.2), with a focus on 'Customer Types', created to represent the variability in peak-hours electricity demand across households grouped by a number of characteristics. The Customer Model provides electricity demand profiles for each Customer Type under both 'baseline' and 'intervention' profiles for a number of scenarios.[32]

Project objectives evidenced in this section:

Customer model development: a model which will allow the output of appropriate demand profiles under baseline and intervention conditions to allow network simulation and scenario-based planning exercises.

Following the first period of trials, the development of the initial Customer Model to meet these requirements using project data was detailed in *SDRC 2.2 SAVE Updated Customer Model* (Rushby et al., 2017). The report presented the Customer Model capabilities, with two outputs:

- the calculation of observed (raw) or aggregated household level consumption profiles created directly from the observed consumption and survey data (Option 1);
- the estimation of small-area level (Output Area) consumption profiles created using spatial micro-simulation to combine observed household level consumption profiles with UK Census data (see SDRC2.2) (Option 2).

Early discussions of the requirements suggested that the latter small area level consumption profiles would be required as input into the Network Model. This would enable appropriate adjustments to the observed profiles to take account of the different socio-demographic composition of the areas on the low-voltage network to be modelled. Using these profiles, the Network Investment Tool would be able to account for geographical variation in demand profiles and intervention responses that would be expected due to differing social composition, alongside expected geographical variation in network re-reinforcement costs.

However, the emerging input requirements of the Network Model set out in SDRC 7.2[33] indicated that the required alignment of the Customer Model and Network Model should be provided through the development of a number of 'customer types' (based on option 1 above) to provide average consumption profiles for each defined type with no geographical adjustment. The 'customer types' would consist of a range of household types which would attempt to capture some of the diversity in consumption across households but would not vary between small geographical areas as only one profile per customer type is provided for the whole of the study area. Inevitably this will reduce the 'true' heterogeneity that can be modelled. However, as the mix of customer types within the population varies by small area, applying the correct mix of customer types to any area - and thus any network asset situated within that area - provides an associated variation in customer demand profiles. Therefore, through the provision of 'customer type' profiles, aligned to household characteristics available for Census Output Areas, the innovative used of spatial variation in customer load profiles is retained within the Network Investment Tool.

In implementing these 'non-spatial' profiles, the selection of the appropriate set of 'customer type' demand profiles to be used in the modelling of a given network topology was also removed from the scope of the Customer Model. This was because the selection of the correct profiles for each individual household is dependent on customer connectivity and allocation to the Network Model considered to be the role of the network planner. This functionality is now provided by the Census Interface module and the process of applying 'customer type' demand profiles to the Network Model is detailed in SDRC 8.5 and 8.6.

32 Within the SAVE project, the agreed outputs to be provided by the Customer Model were mean half-hourly demand (and standard deviation) for each Customer Type (see Section 3.3.1).

33 SDRC 7.2: Project SAVE Network Modelling Tool, EA Technology Limited, available at https://www.ssen.co.uk/save/

The remainder of this section describes the empirical development of the customer types and the process used to generate demand profiles from the data collected by the SAVE project. It is organised into the following sub-sections:

- defining customer types: includes the methodological approach used and practical constraints around the 'customer types'; and analysis of the consumption data collected from the sample of households participating in the SAVE project to recommend a classification of 'customer types' that sufficiently capture differences in household consumption profiles;
- generation of baseline and intervention demand profiles: process used to aggregate household level consumption data in order to generate demand profiles (and examples of output);
- generation of intervention impact profiles: details the process used to generate profiles that represent the change in demand under intervention conditions;
- summary and limitations.

3.2 Defining Customer Types

3.2.1 Methodological approach

The Customer Model spatial micro-simulation (SMS) process described in SDRC 2.2 used a set of variables ('constraints') in the SAVE household survey which match exactly to those available in Census small area tables and which are reasonable predictors of the outcome of interest; evening peak (16:00 - 20:00) electricity consumption. The selection of these variables used a regression approach and version 1 of the Customer Model used nine constraints in the weighting process within the SMS. Clearly these nine could have been a potential basis for a customer typology since they provide an adequate representation of the variation in electricity consumption across households (good regression model fit). However, the large number of possible combinations of these nine household attributes (and the resulting small counts of SAVE sample households) means that they are not suitable for use as the basis for a customer typology.[34]

The development of a more practical set of 'customer types' is described in the sections that follow, and is performed with a view to balancing the following constraints:

- model fit: variables (and sub-categories[35]) selected to provide maximum predictive power of peak hours consumption (kWh);
- number of households per type: the number of households of each 'customer type' that can be drawn from the SAVE sample should be maximised;
- number of 'customer types': the number of types should be minimised to avoid over-complex implementation and constraints within the Network Model[36].

In practice the last two constraints act in the same direction since a smaller number of types will in general have larger numbers of households in each type. The selection of variables used to define 'customer types' is presented in the following section.

3.2.2 Selection of 'customer type' variables

Previous modelling of the SAVE data (reported in SDRC 2.2) produced a ranked list of variables associated with household evening peak consumption. The top three ranked variables were (in order of best predictor): household size (number of persons), number of rooms and main heat source.[37] These variables provided the top candidates used for aggregation of 'customer type' consumption profiles and therefore their ability to predict the outcome variable needed to be assessed using linear regression modelling. As with previous analysis, the outcome variable used was mean consumption (in kWh per household) during the evening peak hours (16:00 - 20:00) for January 2017. Measurement was taken over all non-zero half-hourly kWh values for that period with no account taken of weekdays versus weekends, nor of different half-hours within the peak period. Half-hourly consumption was calculated for each household from the sum of pairs of 15-minute observations, providing kWh consumed per household per half hour for the entire period.

Due to the requirement for the use of a normally distributed dependent variable within the statistical modelling, the (summarised) mean household consumption values have been log-transformed.[38]

34 The number of combinations of household characteristics would have numbered many hundreds. Refer SDRC 2.2 for more details.

35 An additional constraint in the selection of sub-categories for each candidate variable was exerted by the respective variables within the Census data. Alignment was required between the two surveys, and the SAVE household survey data generally utilised a greater number of sub-categories than the Census data (at the required geographical scale).

36 The Network Model is limited to importing fifty demand profiles.

37 See SDRC 2.2 SAVE Updated Customer Model available at https://www.ssen.co.uk/save/.

38 This removes the positive skew in the distribution of consumption values.

To examine the candidate variables in detail, multiple linear regression models were run with multiple variations of coding of sub-categories of each of the variables for combinations of two and three variables. Selection of the most appropriate definition of the 'customer types' was conducted using an iterative process, and involved both reviewing regression model results (model fit, significant differences between categories etc) and visual inspection of the resulting demand profiles. For details, of the original categories and regression model results see Appendix A.2. In the following section the selected best fit combination results are presented.

3.2.2.1 Two-variable definition

Due to the small number of non-gas customers within the SAVE sample, the definition of 'customer types' was initially proposed using two variables: household size and dwelling size (using a proxy measurement of number of rooms). This model was found to explain approximately 27 percent of the variation in peak hours electricity consumption (1st column of results shown in Table 18, Appendix)[39]. As two alternative proxies for dwelling size are available within the Census output area level statistics: number of (all habitable) rooms and number of bedrooms, therefore both were tested. As, using number of bedrooms in place of number of (all) rooms, resulted in only a small reduction in model fit (2nd column of results in Table 18, Appendix A.2), this measure was chosen due to common usage.

The SAVE sample does not include households for every combination of the original survey responses to the two 'customer type' variables. This is illustrated by Table 5 which shows the number of households in the January 2017 data for each combination of household size and number of rooms. It is immediately clear that many combinations have very few or zero households (NA indicates that data is missing from household surveys). There are also some unexpected cases (e.g. 5 people in 1 room). This problem is further exacerbated as the sample size is reduced for the final 'customer type' demand profiles, as profiles are provided for each trial group (reducing the sample of households to approximately a quarter for each).

Table 5: Household counts for 'customer type' combinations, household size (rows) x number of rooms (columns), original coding, all households

	Number of rooms									
Household size	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9+**	**NA**
1	0	4	9	63	84	104	82	64	76	19
2	4	1	9	33	128	178	251	186	418	35
3	0	1	4	5	49	79	125	81	174	30
4	0	0	3	3	17	64	97	94	248	36
5	1	0	2	1	6	23	38	31	83	10
6	0	0	1	0	6	9	8	8	24	4
7	0	0	0	0	1	2	4	1	8	1
8+	0	0	0	1	1	1	2	1	4	0
NA	0	0	0	0	1	1	1	0	0	628

39 Refer to the 'adjusted R2' in the model results table. The academic literature contains a wide range of studies into the relationships between socio-economic, dwelling and appliance characteristics and household electricity demand. Models using only socio-economic characteristics typically report the prediction of similar amounts of variation, although the majority of studies examine consumption over longer time-scales. For a review see *Jones et al., 2015. The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings*. Renewable and Sustainable Energy Reviews 43, 901–917. https://doi.org/10.1016/j.rser.2014.11.084

To ensure that the sample was adequately distributed with sufficient samples in each combination, the number of candidate 'customer types' was therefore reduced by collapsing the survey responses into four categories for each of the variables chosen (household size and number of bedrooms) resulting in 16 (4 x 4) customer types and as a consequence increasing the sample size in many of the 'customer type' categories (Table 6). Collapsing the categories was carried out using an iterative approach to optimise model fit and variance captured across categories.[40]

Collapsing the categories for household size and dwelling size resulted in a marginal reduction of the model fit (adjusted R2 reduced from 0.27 to 0.25) and the regression model results contained in Table 7 show that the average peak-hours consumption of households within each sub-category is significantly different to the reference category (1 person, 0-2 bedroom household). Table 6 shows the resulting cell counts using collapsed categories of the number of bedrooms measure for all households in the SAVE sample. The model fit is approximately the same when using number of bedrooms as a proxy for dwelling size, in place of number of rooms, (adjusted R^2 = 0.25).

Table 6: Household counts for 'customer type' combinations, household size (rows) x number of bedrooms (columns), collapsed coding, all households

	Number of bedrooms				
Household size	**0-2**	**3**	**4**	**5+**	**NA**
1	232	185	52	16	20
2	321	516	280	87	39
3-5	133	598	365	132	77
6+	3	35	23	21	5
NA	0	2	0	1	628

Table 7 Regression results: household and number of bedrooms (recoded)

	Dependent variable: log(mean)Wh
People 2	0.496*** (0.435, 0.558)
People 3-5	0.757*** (0.694, 0.820)
People 6+	0.940*** (0.804, 1.076)
Bedrooms 3	0.143*** (0.088, 0.197)
Bedrooms 4	0.327*** (0.265, 0.389)
Bedrooms 5+	0.359*** (0.274, 0.444)
Constant	-1.891*** (-1.948, -1.833)
Observations	2,985
R2	0.249
Adjusted R2	0.248
Residual Std. Error	0.569 (df = 2978)
F Statistic	164.899*** (df = 6; 2978)
Note:	*p<0.1 **p<0.05 ***p<0.01

Using these 'customer type' definitions, the resulting demand profiles are illustrated in Figure 17 using the baseline (pre-trials) study period data for all households from January 2017.

Figure 17: Mean 30-minute Wh observations by customer type using data for January 2017, faceted by number of bedrooms, colours indicate number of people

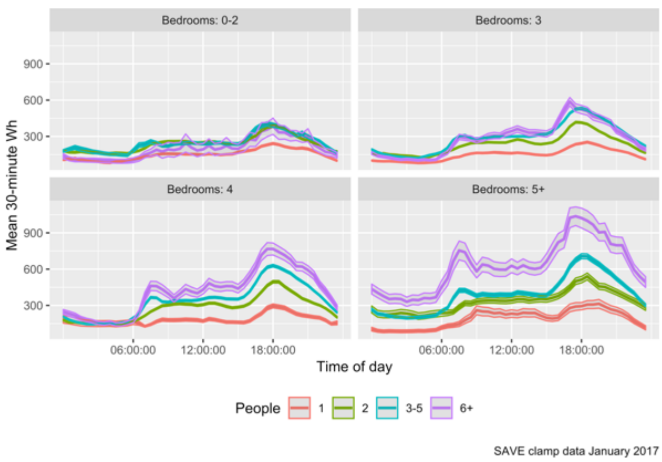

40 An additional constraint to the final definition of the 'customer types' was the requirement to align with categories within the Census data. This was required for operation of the Census Interface module of the Network Investment Tool (see SDRC 8.5 and 8.6).

It can be seen in Figure 17 that profile shapes are broadly similar across customer types. Profiles for larger households generally exhibit a larger and more defined morning peak, as well as a larger peak during the peak hours of 4-8 pm Higher over-night consumption is observed in larger dwellings with more than one occupant. Profiles for the more common customer types are 'smoother', with less variation from one 30-minute time interval to the next, due to the larger sample size within these groups. The small samples of less common customer types give a more irregular profile. The wider confidence intervals around the profiles for the less common customer types reflect the smaller sample sizes for these combinations.

3.2.2.2 Incorporating primary heat source

As noted above, previous analysis has shown that peak-hours electricity demand and profile shape varied according to primary heating fuel (Figure 18) and was found to be the third-ranked variable for predicting peak-hours consumption. The role of electric heating is seen as an important factor in network forecasting and therefore during development the 'customer types' were also disaggregated by fuel type. This allows the customer types to capture the variation in peak-hour load *and* profile shapes, and allows better representation of demand profiles for households using electric and other non-gas heating within network simulations.

Figure 18: Mean 30-minute Wh observations by primary heat source, 5 categories (left) 3 categories (right)

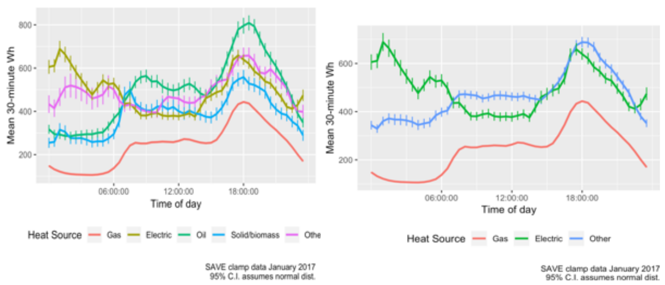

Adding the primary heat source variable to the original coding of household size and number of bedrooms improved the linear regression model fit to capture 31 percent of the variation in peak-hours demand, however due to constraints on the number of 'customer type' demand profiles to be imported to the Network Model and the number of non-gas-heated households within the SAVE sample, the number of non-gas categories needed to be limited. Again, households were grouped into fewer categories than contained in the original responses to the household survey, with all electrically-heated households in one category (containing night-storage and other electrical heating) and other non-gas in another category (containing oil, solid-fuel, biomass and other). This recoding reduced model fit by approximately two percentage points to 29 percent (see 3rd results column of Table 19, Appendix).

At this stage, the 'customer types' definitions were aligned to match the categories within the Census data. This involved recoding responses to the SAVE household survey for household size, number of bedrooms and primary heating fuel. The recoding did not have a significant impact on the fit of the model (see regression results provided in Table 20, Appendix).

To further maximise sample size, non-gas-heated customer types were limited to disaggregating households within fuel-type groups by number of bedrooms only (as opposed number of bedrooms and household size). As a result the number of variables used for the definition of customer types varies across fuel types.[41]

Finally, to maintain sample size, the customer type categories for gas-heated households containing the smallest samples were merged. The final 'customer types' are defined using different variables according to the primary heat source (fuel) of the sample households and are as follows:

- Gas: disaggregated by household size and no. of bedrooms (14 profiles);
- Electric: disaggregated by no. of bedrooms (4 profiles);
- Other: disaggregated by no. of bedrooms (3 profiles).

41 Disaggregating 'other' non-gas-heated households using household size rather than dwelling size provided a better model, with an adjusted R2 value of 0.185. Dwelling size was chosen for consistency in the 'customer type' definition. Final model fit was assessed separately for each fuel, providing adjusted R2 values (for within primary heating group) as follows: 0.30 for gas-heated households, 0.10 for electrically-heated households, and 0.10 for other non-gas-heated households.

The final customer types, and sample sizes using data from January 2017, are shown in Table 8:

Table 8: Final customer types represented in SAVE sample, January 2017

		Number of bedrooms			
Heat source	**Household size**	**0-1**	**2**	**3**	**4+**
Gas	1	55	122	173	59
Gas	2	31	238	466	337
Gas	3	80		247	144
Gas	4+	37		337	357
Electric	All	38	40	49	17
Other	All	40		56	60

3.3 Generation of demand profiles

The next step is to use these 'customer type' demand profiles to improve the representation of electricity consumption across different households in local network topologies. Since the profiles were defined using information readily available from the Census (see Section 3.2.2), it follows that the profiles can be allocated to local areas using UK Census small area household counts. This then allows network planners to determine appropriate demand profiles to be selected for modelling according to the mix of households in any small area[42] and so appropriately represent the 'likely' heterogeneity in a given location.

In addition to providing representative baseline demand profiles that can be allocated to local areas, a second objective of the SAVE project was to build and demonstrate a model that could undertake simulations of interventions trialled within the project and to allow the results to be extrapolated to the general customer population. This capability is provided by the 'SAVE customer type demand profile generator' which produces electricity demand profiles for **each** customer type under a number of scenarios. Each scenario reflects the average loads that would be expected to occur for households in each Customer Type under the following conditions:

Baseline: household electricity demand under the 'control' condition (i.e. no intervention);

Treatment: household electricity demand under a number of 'treatment' conditions tested during the SAVE project (e.g. interventions such as event/challenge days, or dynamic pricing rebate incentive);

The resulting demand profiles provide the network planner (within the SAVE project, the Network Model) with the appropriate household demand to compare loading on the network under control and intervention conditions.[43]

3.3.1 Demand profile statistical definition

For the purposes of the SAVE project, the following statistics (metrics) are used for the 'Customer Type' consumption profiles which are passed from the Customer Model to the Network Model. Each 'Customer Type' profile under each scenario consists of:

- Mean half-hourly consumption (kWh) for each of the 48 'half-hours in the relevant summarised period: this provides the *p* value for Debut[44];
- Standard deviation (kWh) for each of the 48 'half-hours in the relevant summarised period: this provides basis for the *q* value for Debut.

These metrics are calculated for each 'Customer Type' and treatment group, and for each defined time period. As the previous discussions will have made clear, this means that some of the profiles are based on small numbers of households (see Table 9, Table 10 and Table 11 below) and the potential for small-sample effects such as extreme values or unrepresentative profiles will be high in these cases.

3.3.2 Profile generation methods

'Customer type' demand profiles are generated for each scenario using two methods:

- Households with gas as the primary heat source: demand profiles are generated by summarising the consumption across households in each group (Customer Type) to provide mean and standard deviation (SD);
- Households primarily heated with electricity or other non-gas fuels: due to the small numbers of relevant households (see Table 8), synthetic demand profiles are generated by scaling the mean demand profile and standard deviation of all households (grouped by primary heat source).

42 Modelling scale can vary but SAVE applied the method using Output Area level statistics from Census data (available at Nomis, https://www.nomisweb.co.uk).
43 For details of how the Customer Type demand profiles are selected and applied, see SDRC 8.5/8.6.
44 Debut is the load flow application underpinning the Network Model (see SDRC 7.3).

For both methods, the consumption data is initially aggregated to 30-minute time intervals using summed pairs of the 15-minute data for each household. Then, for scenarios with longer-term impacts, where intervention effects were evaluated as weekly averages (for example interventions such as upgraded LED lighting intervention and dynamic pricing rebates), the demand profiles were summarised across a number of days to create average weekday demand for each 30-minute time-period.[45] When creating a profile that aggregates across a time period, observations from all days in the period are used. For example, for a customer type with 100 households in the sample:

- A demand profile for a single day or date will use 100 observations for each half-hour metric;
- A demand profile aggregating over five days will use 100 x 5 = 500 observations for each half-hour metric.

3.3.2.1 Profiles for gas-heated households

For gas-heated 'customer types', the demand profiles are generated by calculating the mean and standard deviation of all (30-minute) consumption observations within each 'customer type' for each half-hourly time-slice in the relevant period. Example profiles for gas-heated households are shown in Figure 19.

Figure 19: Example customer type demand profiles, gas-heated households, Winter

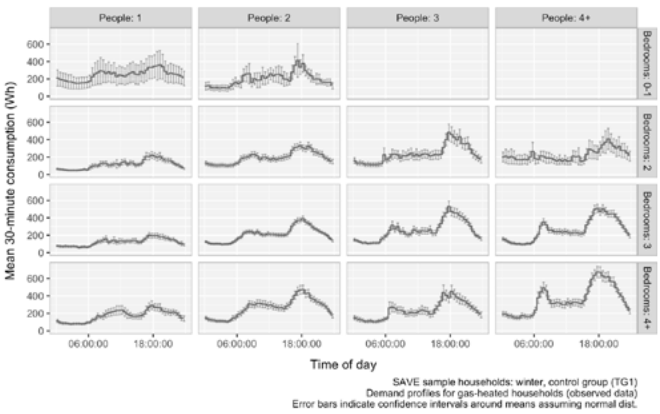

The sample size for the gas-heated customer type demand profiles illustrated above are shown in Table 9[46]. The collection of electricity consumption data was problematic for some households due to communication problems. This affected sample size available to drawn upon. The large variation in sample size by Customer Types is also noted, reflecting the distribution of the sample across groups and identifying rare types. To mitigate the small samples in some customer types, the 0-1 bedroom and 2-bedroom categories were combined for 3 and 4-person households. The profiles are improved by combining these categories, however the small sample effects are still visible in the wider confidence intervals and as greater variability of consumption across consecutive half-hours in some customer types illustrated in Figure 19 (e.g. 2P-1B and 4P-2B).

Table 9: Sample size (number of households) for customer type demand profiles, Control group (TG1), Winter

Customer Type	Minimum sample	Maximum sample
Gas_1P_0-1B	16	16
Gas_1P_2B	31	31
Gas_1P_3B	38	40
Gas_1P_4+B	21	21
Gas_2P_0-1B	8	8
Gas_2P_2B	48	48
Gas_2P_3B	118	121
Gas_2P_4+B	59	59
Gas_3P_0-2B	21	23
Gas_3P_3B	60	60
Gas_3P_4+B	31	33
Gas_4+P_0-2B	9	11
Gas_4+P_3B	78	82
Gas_4+P_4+B	73	76

3.3.2.2 Synthetic profiles for non-gas-heated households

As noted above, the SAVE sample contains insufficient numbers of households heated primarily with fuels other than gas for their profiles to be considered representative within each customer type category (see Table 10 and Table 11 below). Due to the small sample size for these households, the observed profiles were more likely to be affected by unrepresentative households, therefore synthetic consumption profiles were constructed. Standard deviation values for the synthetic profiles were calculated using the same approach.

45 In these scenarios, profiles were generated for three 'day-types': weekdays, Saturdays and Sundays, although only profiles generated for weekdays involved summarising data across multiple days.

46 In any period, the sample of households contributing data may vary due to a number of factors such as sample attrition, participant withdrawal or communication issues resulting in differing sample sizes across time periods. Minimum and maximum sample sizes are therefore shown in the table.

Such profiles were created to allow appropriate representation of these customers within the Network Model. The method used is summarised below:

- Aggregate non-gas households into *electric all* and *other all* heat-source categories;
- Calculate the mean consumption profile (using 30-minute total consumption) to produce mean Wh per half-hour for each heat-source category;
- Calculate the mean daily total consumption for each heat-source category across all dwelling sizes;
- Calculate the mean daily total consumption for each by heat-source category for each dwelling size (number of bedrooms);
- Scale the mean consumption half-hourly Wh (2) according to the ratio of daily total consumption for each dwelling size to all dwellings (i.e. (4)/(3)).

Clearly the use of synthetic profiles constructed in this way results in an even greater loss of the heterogeneity across households by household size and by primary heat source categories (i.e. electric storage heaters with electric other, or oil with solid fuels and others), as shown in Figure 20. However, the figure shows that the resulting profiles exhibit more consistency, i.e. smoother profiles with less 'noise'.

Examples of the synthetic profiles are shown below. Figure 20 shows example synthetic profiles generated for electrically-heated households by dwelling size (number of bedrooms), compared to synthetic profiles for gas-heated households (shown in black). The (mean) observed profiles are also shown (as lighter lines) for each dwelling size as a comparison to the synthetic profiles, showing the extent to which the synthetic profiles differ from the observed consumption.

Figure 20: Example customer type demand profiles, electrically-heated households, weekdays January 2018

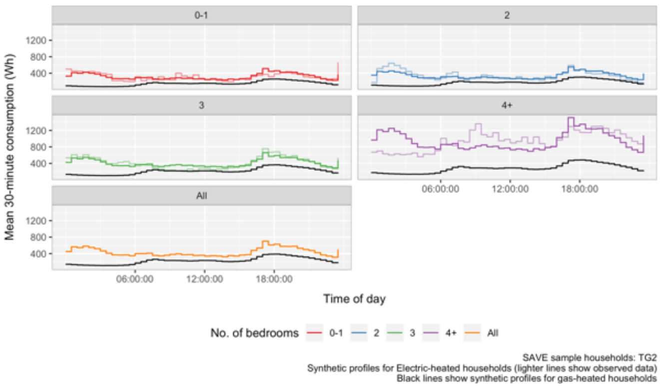

While the synthetic profiles for households occupying dwellings of up to 3 bedrooms match the observed data well (although under-estimating the overnight peak for 2-bedroom households), the synthetic profile for larger households (4+ bedrooms) does not fit the observed data for this sample. The data shows that the mean demand profiles for electrically-heated households in larger dwellings are different in shape than those in smaller dwellings. Comparing these profiles to those obtained using the whole sample confirms that this is also true when drawing on a larger sample, and cannot be attributed to small sample effects. While the synthetic profiles are a good fit for smaller dwellings, for larger dwellings that are less likely to be heated primarily with night-storage heaters, the synthetic profiles over-estimate load during the overnight hours (and underestimate load during the day and late evening). This should be considered by planners when assessing network assets with a high proportion of large, electrically heated dwellings. In the majority of instances however, this should not be an issue as the profiles match quite well between 4pm and 8pm, peak time.

The mean total daily consumption values used for scaling are shown in Table 10, and reveal the small sample size underpinning the profiles for electrically-heated households. Further anomalies within the synthetic profiles were observed and varied across treatment groups. For example, the scaling factors calculated for electrically heated households were observed to be higher for the 3-bedroom households than for 4+ bedroom households for the control group. This was likely caused by a small sample of unrepresentative households.

Table 10: Mean total daily consumption (Wh) by dwelling size, electrically-heated households, trial period 2, LED treatment group

Bedrooms	Day-type	total Wh	Households	Scaling value
0-1	Weekday	15,301	8	0.73
2	Weekday	16,428	12	0.78
3	Weekday	19,693	18	0.94
4+	Weekday	45,090	5	2.15
All	Weekday	20,946	43	1.00

Figure 21 shows the synthetic profiles generated for other non-gas-heated households by dwelling size, again compared to synthetic profiles for gas-heated households (shown in black) and observed data (lighter coloured lines). There are differences between the synthetic profiles and the observed profiles for the households occupying dwellings with 2 or fewer bedrooms and underestimation of the afternoon/evening peak in 3-bedroom households. Table 11 shows how these smaller dwellings have the smallest samples of all, and the need to synthetic profiles to manage suspected small sample effect in 'other' heated households.

Figure 21: Example customer type demand profiles, other non-gas-heated households, weekdays January 2018

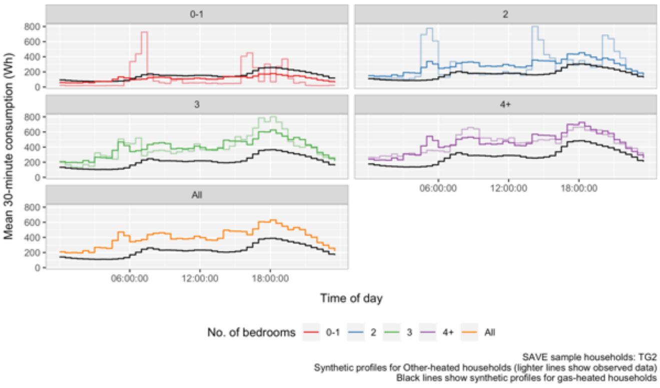

For 'other non-gas' households, the profiles for the observed data for 0-1-bedroom and 2-bedroom households indicate that secondary electrical heating loads are present and provide the peak loads for these customers. These customer types provide a small proportion of the households within the other non-gas-heated households category and so the shape of the synthetic profiles are dominated by the more numerous households with 3 or more bedrooms. As a result, the synthetic profiles for other non-gas-heated household types are clearly underestimating the peaks in consumption that exist in the observed data. This is (at least in part) due to the synthetic profiles being 'smoothed' due to the 'diversity' effect of the larger sample which has implications for network planning. In an area with very little diversity i.e many households of a single type (specifically smaller dwellings), the magnitude of the peak load on the network may be under-estimated and the timing of the peak may not be accurate.

With a small sample of households contributing data for these customer type categories, the 'true' representation of mean demand profiles among these households is not known. The true heterogeneity across households by number of bedrooms is also difficult to assess and thus the difference between the synthetic and 'real' demand profiles – and the associated impact on the accuracy of simulated loads cannot be quantified accurately. Network modelling results based on these customer types should therefore be viewed and interpreted with caution. As an alternative, planners may choose to use existing WinDebut profiles for Economy 7 for simulations.

The small sample size underpinning the profiles is shown in Table 11.

Table 11: Mean total daily consumption (Wh) by dwelling size, other non-gas-heated households, trial period 2, LED treatment group

Bedrooms	Day-type	total Wh	Households	Scaling value
0-1	Weekday	5,356	1	0.28
2	Weekday	13,630	7	0.72
3	Weekday	18,804	12	0.99
4+	Weekday	21,822	18	1.15
All	Weekday	18,931	38	1.00

3.3.3 Comparing customer profiles

The SAVE customer type profiles offer a number of improvements over existing input load profiles. To illustrate, the SAVE profiles for gas-heated households have been compared to demand profiles from the Energy Networks Association (ENA) ER P5 profiles[47]. The results are shown in Figure 22 below.[48] The SAVE profiles (black profile in the figure) for 1-person households with 1-bedroom and 2-person households with 3 or fewer bedrooms are comparable with the 'low income' P5 profiles (blue in figure). The profiles for one-person households generated by the SAVE data are generally lower than all of the P5 profiles. SAVE demand profiles created for 3-person households are generally comparable to the 'medium income' profiles from P5 (green). Only the profile generated for 4(+)-person and 4(+)-bedroom households provides higher consumption values than the 'high income' P5 profile (red).

47 Energy Networks Association, *Engineering Recommendations P5: Design methods for LV underground networks for new housing developments*, Issue 6. 2017

48 To enable the comparison, the customer type profiles were converted from half-hourly consumption in Watt-hours (Wh/30-min) to constant power equivalent in kilo-Watts (kW).

Figure 22: Comparison of household load profiles: SAVE gas-heated customer types and P5 profiles

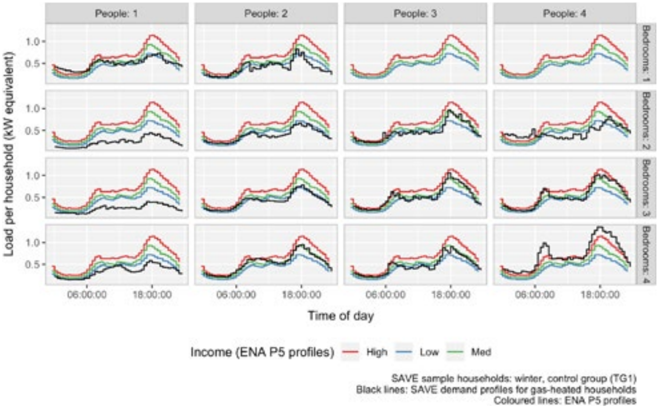

Figure 23 shows the same P5 profile data against the SAVE synthetic non-gas-heated customer type profiles. In the left panel, it is clearly shown that the P5 profiles provide significantly smaller loads for households living in dwellings with three or more bedrooms compared to the SAVE electrically-heated baseline demand profiles. In contrast the P5 profiles provide higher loads for smaller dwellings (one or fewer bedrooms). The SAVE profiles provide overnight loads that are more representative of households with electric storage heating (e.g. Economy 7) and show that the P5 profiles are clearly underestimating overnight demand for electrically heated households. Comparing the P5 profiles with the SAVE other-non-gas baseline profiles (Figure 23, right panel), the SAVE profiles provide a significantly more granular insight and hence more variability in demand.

Figure 23: Comparison of household load profiles: SAVE non-gas-heated customer types and ENA P5

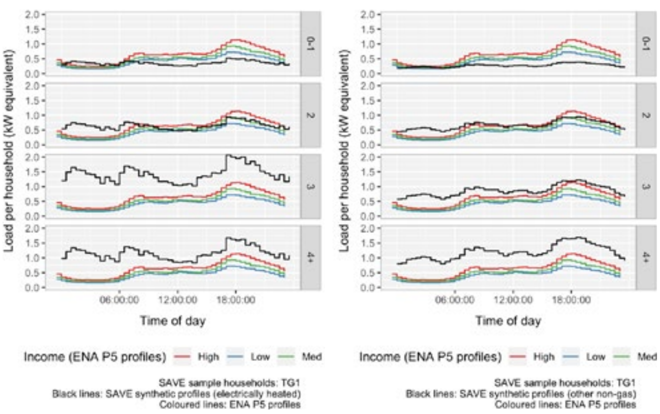

3.4 Intervention impact profiles

Initially, the impacts of SAVE interventions were to be embedded within the customer type demand profiles. This method is detailed in SDRC 2.2 and was implemented by providing a set of 'baseline' demand profiles for each customer type under control conditions, plus one set of 'intervention' demand profiles under each of the intervention conditions, following the method outlined above (Section 3.3). Using this method, the treatment effects would be represented by the difference between 'baseline' (control) and 'intervention' (treatment) profiles. By providing the two sets of demand profiles as input to the Network Model, simulation of the two conditions is possible.

A number of developments ruled out this original approach. First, the original approach required equivalency of treatment and control groups prior to the intervention which was shown to be invalid during the evaluation of trial impacts (see Section 2.6 above). The observed (pre-treatment) differences between the control and treatment groups invalidated the assumption of equivalency and, depending on the direction of the pre-existing difference, may have acted to mask the effect of interventions, or reveal effects where none were present.

Second, the generation of synthetic profiles remove any heterogeneity of treatment effect for non-gas groups. Further, any differences between the synthetic profiles of each trial group should be treated with caution as they may be an artefact of the small samples and profile generation method and not derived from the treatments.

In response to these issues, the Customer Type model implemented an alternative approach to providing intervention profiles using the difference-in-differences method to estimate treatment effects for each customer type (and for each intervention), and applying the resulting 'impact profile' to the baseline demand profile. The intervention impact profiles were estimated as the difference-in-differences of household consumption pre- and post-intervention (or reference week, t_0 and test week, t_1). The intervention impact profiles were generated using Coordinated Universal Time (UTC) timestamps to align with the demand profile generation and requirements of the Network Model. This differs from the evaluation presented in Section 2.6, which was conducted using Greenwich Mean Time (GMT) and British Summer Time (BST) to account for daylight saving and to capture effects at the time householders experience the intervention. While discrepancies will therefore exist between the impact estimates generated by the two methods when daylight saving is in force (from the last weekend in March through to the last weekend in October), this does not affect network simulation.[49] Figure 24 shows an example of the resulting intervention impact profiles for the LED upgrade treatment for gas-heated customer types.

Figure 24: Example intervention impact profiles for gas-heated customer types, LED intervention

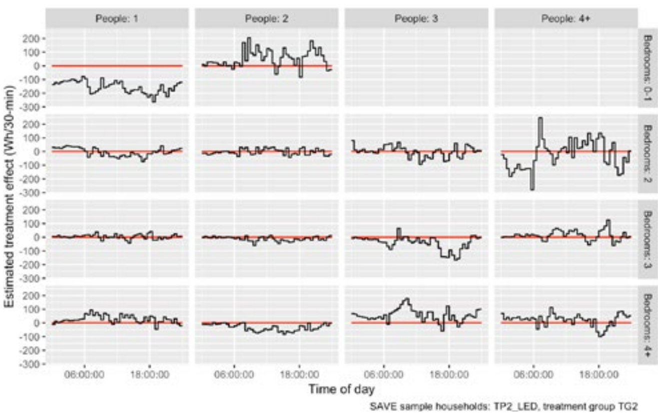

Figure 24 shows the highly variable impact profiles gained across the customer types and clearly indicates the small sample effects in the less common categories. For example, the estimated treatment effects for 1-person, 0-1-bedroom households show a maximum of approximately 400 Watts (200 Wh/30-min) reduction in load. This reduction is very large and unlikely to be solely attributable to treatment effect. Similarly, there are large increases in demand shown for some customer types that are unlikely to be treatment effects, rather artefacts of the small sample and natural variability in consumption not related to the intervention.

To address the small sample problems, and for the purposes of generating more reliable impact profiles to apply in network simulation, the sample was disaggregated into fewer customer types. Gas-heated households were disaggregated only by dwelling size (number of bedrooms), providing four intervention impact profiles. Treatment group profiles were constructed by applying the impact profiles to the customer type 'baseline' demand profiles. The baseline profiles are modified with exactly the same intervention impact values for each dwelling size, and as the baseline profiles for the full set of customer types are used, the resulting intervention profiles are unique for each.

Figure 25 provides an example of the construction of customer type demand profiles under treatment conditions for the LED upgraded lighting scenario (only gas-heated households are shown). In the figure, the black lines represent the impact profile for the intervention (i.e. the treatment effect), grey lines show the baseline (control group) demand profiles, and the blue lines show the resulting intervention profile (sum of *baseline* and *impact*) which provides the demand profile under intervention conditions for each customer type. Note that due to the collapsed customer type categories, no profiles are shown for 3 and 4+ person households with fewer than two bedrooms. The figure shows that the same intervention impact profile has been applied to all customer types of each dwelling size (bedrooms) category (i.e. all customer types in each row). The intervention impact profiles consistently show a reduction in consumption during the evening peak hours with some variation by number of bedrooms, consistent with the observed treatment effects from the evaluation of the LED intervention.[50]

49 The calculation of the difference-in-differences also used the arithmetic mean in order to align with the provision of mean and standard deviation as inputs to the Network Model. It should be noted that this is in contrast to the measurement provided by the trial impact analysis which estimated the geometric mean of treatment effects. Slight discrepancies between the estimates of the two methods will exist.

50 Network models also require standard deviation as input. As the impact profiles model only the change in consumption, statistically accurate values standard deviation of the resulting consumption values cannot be provided. Differences in standard deviation across the treatment groups (i.e. difference from control group) were adopted and therefore also include any pre-existing, between-group differences in variance not related to the intervention. An example of the standard deviation intervention impact profiles is provided in Appendix A.4.

Figure 25: Example of intervention group load profiles constructed using baseline and intervention impact, LED upgrades treatment (weekdays)

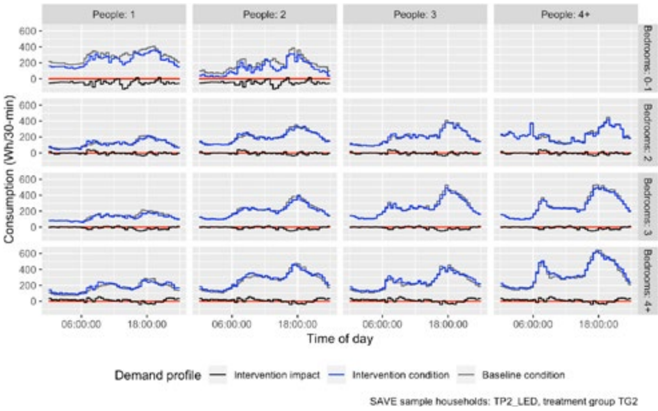

Large treatment effects, with high variability between half-hours and a distinct reduction in the afternoon, are observed for the households with 0-1 bedroom.

As shown in Table 8 above, only a small proportion of the SAVE sample households were primarily heated with fuels other than gas: approximately 5 percent each for electric and other non-gas heating fuels. As with the generation of demand profiles under baseline conditions, the small sample of these households also caused limitations in generating intervention impact profiles. To provide representation of the impact of interventions for these customer types, a single profile was generated for electrically-heated households and a single profile for other non-gas households. The profiles for change in mean consumption are shown in Figure 26 and show that the estimated effects are much larger, exhibit far greater noise (high variability between consecutive half-hour periods) and have much wider confidence intervals than those produced for the gas-heated customer types. Large changes in load are indicated during overnight periods which do not align with the observed treatment effects in the wider sample. These effects are more likely an artefact of (uncontrolled) pre-existing differences in consumption in households of this type between the treatment and control groups or the background variability in household consumption over time ('noise'). These intervention impact profiles are therefore considered not to be representative of expected treatment effects.

Figure 26: Example intervention impact profiles for non-gas heated households showing change in mean consumption, LED upgrades treatment group

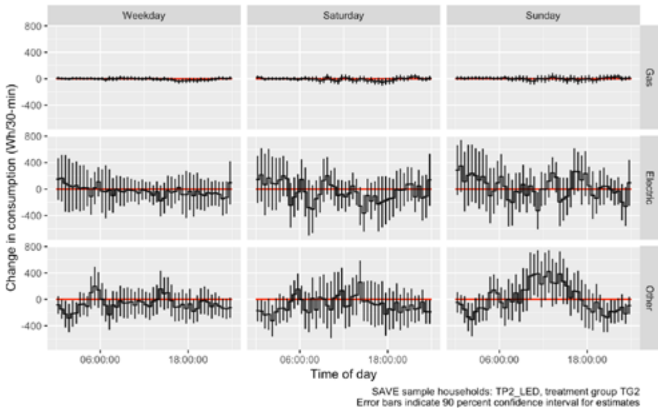

While the small sample sizes in the less common customer types result in the generation of unrepresentative impact profiles, the method shown for constructing intervention profiles allows the substitution of more representative impact profiles from other more common customer types, or the use of impact profiles generated from larger groups of households. The applicability of impact profiles will however vary from case to case. For example, impact profiles for gas-heated customers may be applicable to other customer types for the LED intervention as the size of the treatment effect is less likely to be affected by primary heating fuel than other interventions. In contrast, for dynamic pricing and event-based interventions, primary heating source may play a larger role in how households respond. Non-gas households potentially have more scope to shift demand from high consuming appliances, including primary and secondary heating loads, than households with gas as their primary heat source.[51] It follows that the use of the impact profiles generated from gas-heated customer types in this instance may be conservative, however the use of the impact profiles generated for non-gas households is not recommended for the reasons already stated.

51 This finding supported by analysis of events run through community coaching trials. See SDRC 8.8: Community Energy Coaching Trial – Final Reporting available at https://www.ssen.co.uk/save/.

3.5 Summary and recommendations

A customer typology has been developed using the three household characteristics available in both the SAVE household surveys and the Census small area statistics: household size, number of bedrooms and primary heat source. Out of a selection of candidate variables, these characteristics were found to best predict the variability in household consumption during peak hours. In the supporting analysis the customer types as defined were found to explain approximately 29 percent of this variability.[52] Through this development, **the Customer Model has been shown to provide a much more diverse range of demand profiles than those in current guidance**, and can be adopted by network planners to better represent the variety of customers in the population for network modelling and simulations.

The functionality of the Customer Model framework has also been demonstrated, showing that through the implementation of the customer typology and the generation of baseline demand and impact profiles, the treatment effects attributable to SAVE interventions can be passed as outputs from the Customer Model to network modelling applications. This **allows network planners to simulate loads under both baseline and intervention conditions** in order to evaluate the effectiveness of the SAVE interventions on specific networks, each with a specific mix of customer types.

The creation of intervention *impact* profiles also allows the estimated treatment effect from an intervention to be applied to any baseline demand profile outside of the time period during which that intervention was trialled or to different customer types (a different population). For example, a week-long event-based intervention may be implemented during a week in December (coinciding with the period best representing the typical winter peak load), not November as trialled. This capability **allows a network planner flexibility in the simulation scenarios that can be run** and can be tailored to the constraints for any specific network. Caution should be exercised as impact profiles applied in this way may not be representative of the expected treatment effects.[53]

Whilst the representative SAVE sample was designed to be of sufficient size to be able to detect statistically significant treatment effects during the intervention trials, when broken down into the customer types, the ability to provide representative demand and impact profiles was reduced. This resulted in the following important limitations for the Customer Model outputs:

- quantity of households within some customer types may provide unrepresentative baseline demand profiles;
- synthetic profiles suppress the heterogeneity of *baseline* demand profiles for non-gas households and result in unrepresentative profiles for some sub-groups;
- a combination of high background variation in demand, unrepresentative households (or unusual behaviour), and small samples act to generate estimated impact profiles for some customer types that cannot be attributed to SAVE interventions.

Through the implementation of synthetic demand profile creation and the generation of intervention impact profiles, the project has demonstrated methods with which to mitigate some of the limitations with respect to the uncommon customer types. However, for the few networks that dominated by the less common customer types, the outputs from the Customer Model are less robust. The collection of further data to compliment the SAVE sample for these uncommon household types - particularly households using non-gas fuels as primary heating – is therefore recommended.

Through the customer typology, **the Customer Model provides a set of household demand profiles generated from a high-quality, representative sample of households**. The model also provides intervention profiles generated from robust, industry-leading trials into domestic demand response. The Customer Model provides a platform which can incorporate future developments such as supplementary datasets and results from similar high-quality energy efficiency and demand response trials.

52 This does not imply that aggregating household level kWh data to mean(kWh) and sd(kWh) for each Customer Type will adequately represent the distribution/variance in household evening peak electricity consumption of the resulting groups. This is important as layers of aggregation (means of means and standard deviation calculations on means of means) may suppress variation with the result that the output profiles cease to represent the actual variance experienced by the network. To make this clear, the results of calculating the mean and standard deviation for the kWh consumption values at different levels of aggregation were examined. The results are presented in Appendix A.3.

53 In a new context, the experimental controls provided by the trial design are no longer present. During SAVE, environmental and other variables are held constant across the treatment and control groups during trials negating the need to control for such variables within the statistical models. Where different conditions are present, the treatment effects may vary from those estimated.

4

FUTURE DEVELOPMENT AND WIDER APPLICABILITY

The Customer Model developed under the SAVE project provides an inherently replicable and scalable framework, with a number of routes available for further development and the wider applicability of the model. In brief, these routes are as follows:

- Refinement of customer type profiles: integrating new sources of input data to further improve demand and intervention impact profiles;
- Integration of low-carbon technologies: providing 'impact profiles' for LCTs such as air and ground-source heat-pumps, electric vehicles and solar photovoltaic panels for customer types;
- Refinement to the customer typology: development of the customer model to respond to challenges linking data;
- Application of the SAVE methodology to wider UK context: requirements for DNOs to adopt the Customer Model framework within their operations.

These future development pathways are expanded upon in the following sections.

4.1 Refinement of customer type profiles

The SAVE customer type demand and intervention profiles are underpinned by an industry-leading, high-quality large and representative sample of households in southern England. This sample was designed to be large enough to detect statistically robust results for the energy efficiency and demand response interventions trialled during the SAVE project, however the emerging requirement and use of the sample in the development of a customer typology revealed limitations of the sample to robustly represent the rarer customer types. The relative rarity of households not using gas as primary heat source, contrasts with the potential impact of such households on modelling the LV network: they exhibit a wider range of peak loads and demand profiles with significantly different shapes. These households are therefore of significant interest to network operators (although in relatively few network areas currently, numbers of such households are likely to increase with the decarbonisation of heat). Due to the small sample of these households within the SAVE sample, the profiles provided by the Customer Model have larger uncertainties. There is an opportunity to provide more robust and representative profiles for these customer types.

Providing a dataset with a greater sample of these households would improve the representation of non-gas customer types within the customer model and allow the preservation of heterogeneity within these groups. This is currently suppressed by the implementation of synthetic demand profile generation. The SAVE trials have also identified that non-gas-heated households may also show greater potential for demand response, an indication which requires further examination in trials with larger populations of these groups. If confirmed by further studies, the potential impact of SAVE interventions under the electrification of heat could increase.

Smart meters provide a large potential source of customer demand data that is compatible with the Customer Model. The SAVE Customer Model and outputs are underpinned by the high quality of the SAVE dataset, specifically the linking of household electricity consumption data to detailed household surveys of socio-demographic characteristics. Great attention has been paid to recruiting the SAVE sample to ensure that households are representative of the wider population and in evaluating the recruitment outcomes. To retain the high quality of the data upon which the customer type demand profiles are based, any additional data used to compliment the SAVE dataset should be of a similarly high standard and be subject to the same robust quality checks. To allow implementation of smart meter data as an input to the Customer Model, such data should be paired with, at a minimum, the household characteristics required to identify the customer type, along with other socio-demographic data to ensure that the sample of households can be evaluated to determine the extent to which they reflect the wider population within each customer type.[54]

4.2 Refinement of customer typology

The SAVE dataset contains 2 years of consumption data which can be further exploited using advanced analytical techniques. While outside of the scope of the statistical modelling undertaken within the SAVE project, further modelling (for example time-series analysis and models controlling for environmental variables) may improve confidence levels surrounding the predicted baseline loads for the defined customer types by utilising a greater range of data from the full dataset.

54 For more details around recruitment and evaluation of the SAVE sample see SDRC 2.2: SAVE Updated Customer Model available online at https://www.ssen.co.uk/save/

With additional data for rarer household types and/or advanced analytics, the relationships between household characteristics and peak demand (and profile shape) could be further developed to include new variables. Greater disaggregation of households using additional characteristics could further improve the diversity captured by the customer types and allocation of demand profiles.

4.3 Integration of low carbon technologies

LCTs present considerable risks to DNOs as potential significant growth in household demand during network peak-hours. The potential impact of LCTs is such that they alone may shift the localised network peak-hour period (e.g. overnight EV charging). A similar approach to that applied to household demand profiles could be applied to individual technologies such as electric vehicles and heat-pumps, allowing network planners to select corresponding loads to apply when modelling future load growth.

Again, the quality of the input data – including linked household socio-demographic characteristics - is essential in maintaining the quality of the resulting output demand profiles for each LCT technology. Careful consideration should be given to how linked data is collected and provided to ensure that the associated demand (or generation) profiles can be evaluated in terms of how representative the sample is of the wider population. Analysis should also be conducted to determine whether variation in the profiles for each technology are associated with the customer typology developed within the SAVE project, or whether each LCT requires a specific customer typology to capture diversity across groups of customers. The method to establish such technology specific typologies would follow the method set out in development of the customer types, as follows:

- Selection of customer type variables: identify household characteristics associated with the LCT appliance load/ generation profile (peak hours load and/or characteristics of profile shape);
- Generation of baseline demand profiles using customer typology;
- Generation of impact profiles under a range of intervention conditions, for example managed charging for EVs and direct control of heat-pumps.

4.4 Applying SAVE methodology to wider UK context

The SAVE project was designed to test the effectiveness of a range of demand response interventions to reduce and/ or shift peak-hours demand on winter weekdays. Further, the project sought to develop a generalisable customer model to assess a range of intervention scenarios through the generation of demand profiles under baseline and intervention scenarios. The Customer Model, through the customer typology, baseline demand and intervention impact profiles, has delivered these objectives.

The applicability of the Customer Model to a wider geographical area is dependent upon how well the Solent region can be assumed to match the wider UK: the sample was designed to be generalisable to the wider customer population within the study region. In pursuit of this objective the SAVE household sample was designed and to be representative of households in the Solent area of southern England and evaluated on this requirement.[55]

While the Customer Model outputs can be adopted and applied to network simulation throughout the UK, a detailed comparison of the SAVE demand profiles with consumption profiles from other regional samples has not been conducted. This is necessary as figures for household electricity consumption indicate regional differences across England and the UK[56]. Further work is required to assess the applicability of the SAVE demand profiles for predicting demand of customers in other geographical regions, and to give confidence to other DNOs wanting to apply the SAVE Customer Modelling framework to their regions. Ideally, the following actions are recommended to ensure the wider applicability of the SAVE Customer Model to another region within the UK:

- comprehensive comparison of SAVE customer type profiles to consumption profiles (and linked household socio-demographics) from the new area to assess any significant differences in demand;
- where the SAVE consumption profiles are representative, replacement Census data from the new area can be used to allocate customer type profiles for simulation;
- where the SAVE consumption profiles are not representative of the new area, a similar trial could be required in the new area (or access to smart meter data with linked household socio-demographics), with the Census data also replaced.

55 For detailed description of the sampling methodology and recruitment outcomes, refer to SDRC 2.2: *SAVE Updated Customer Model* available online at https://www.ssen.co.uk/save/

56 See average domestic electricity consumption figures by region in Sub-national Electricity and Gas Consumption Statistics: Regional and Local Authority, BEIS (2018).

Top-down analysis of feeder or substation demand monitoring could be used in place of large household datasets to understand the margin of error between simulated and observed aggregate loads. The margin of error could then be compared across different networks (and regions) to assess the accuracy of the customer type demand profiles in different geographical regions.

Analysis of aggregate load data could also be used to model the relationship between aggregate loads and socio-demographic data for each area. These approaches could be used to verify the wider applicability of the SAVE customer type demand profiles. For this approach, the SAVE community coaching trials revealed the importance of the accurate mapping of customer connections and consumption data for other non-domestic loads.[57]

57 See SDRC 8.8: *Community Energy Coaching Trial – Final Reporting* available at https://www.ssen.co.uk/save/.

5

CONCLUSION

Through the implementation of a best practice trial design, the SAVE project has delivered a high-quality sample, representative of SSEN's wider customer base in Solent region.[58] As such, the household electrical consumption and linked household data provides an industry leading dataset to underpin the generation of customer demand profiles for network planning and investment tools. Furthermore, a range of energy efficiency and behavioural interventions has been carried out using a randomised control trial, considered to be the best practice method for experimental work.[59] Through the use of the high-quality dataset associated with the trial, the SAVE Customer Model has provided the following:

First, an evaluation of the impact of the interventions tested under the SAVE project was conducted, providing estimates of the range of likely impacts of each on household electricity demand. In addition, the evaluation of the trial provided an increased understanding of the differential effects and persistence of each intervention, along with indications of how the response might vary across different households. For the upgraded LED lighting treatment, the maximum average estimated treatment effect was a 7 percent reduction in consumption during peak hours, equivalent to 47 Watts per household (90% CI, -96 to 7 W). The maximum effect was observed during the week commencing 15th January 2018. These effects along with those observed during event-based interventions with this group are shown in Table 12.

Table 12: Summary of observed treatment effects, treatment group 2

Trial period, group, intervention		Element	Effect size (peak hours)		Max. impact (Wh) (90% CI) and period observed
			%	Wh/h	
TP2	TG2: LED lighting upgrades	Max average weekly treatment effect (all hours)	**-7.0**	-31	-31 (-61 to 2) w/c 1st Jan 2018
		Max average weekly treatment effect (peak hours)	**-7.0**	-47	-47 (-96 to 7) w/c 15th Jan 2018
TP3	TG2: LED and data informed	Event 1: 1 day 10-10-18 (text message)	**+2.1***	+8*	-23 (-70 to 30) 6-7pm
		Event 2: weekdays w/c 29-10-18 (email + Loop)	**-2.2***	-12*	-20 (-78 to 44) 5-6pm
		Event 3: weekdays w/c 19-11-18 (postcard)	**-2.9***	-16*	-22 (-74 to 35) 7-8pm
		Event 4: 1 day 13-12-18 (text message)#	**-0.2***	-1*	-11 (-62 to 48) 4-5pm

* Value calculated as mean of hourly effect estimates across peak hours
# Results shown for Event 4 are the day-to-day estimates.

For the data-informed and price incentives (TG3) and data-informed only (TG4) treatment groups, the results of the trial evaluation are shown in Table 13 and Table 14 below.

In trial period 2 the maximum treatment effects, averaged over the peak hours (16:00 to 20:00), were as follows:

- Data-informed + £ (TG3): 7 percent reduction (35W) observed during Challenge 4
- Data-informed only (TG4): 3.8 percent reduction (21W) observed during Challenge 1

58 Evaluation of the recruitment outcomes was provided in SDRC 2.2 available online at https://www.ssen.co.uk/save/.
59 Frederiks et al., 2016. Evaluating energy behavior change programs using randomized controlled trials: Best practice guidelines for policymakers. Energy Research & Social Science 22, 147–164. https://doi.org/10.1016/j.erss.2016.08.020

Table 13: Summary of observed treatment effects, data informed and price signal interventions, trial period 2

Element	TG3: Data informed + £			TG4: Data informed only		
	%	Wh/h	Max. impact (90% CI) and period observed	%	Wh/h	Max. impact (90% CI) and period observed
Longitudinal (max. avg. weekly effect)	**-2.8**	-18	-18 (-69 to 38) Challenge 2	**-4.2**	-24	-24 (-69 to 25) Challenge 1
Short-term						
Challenge 1:	**-5.5***	**-32***	-38 (-91 to 22) 5-6pm	**-3.8***	**-21***	-39 (-91 to 18) 5-6pm
Challenge 2:	**-0.8***	**-4***	-9 (-70 to 58) 6-7pm	**-1.3***	**-7***	-26 (-76 to 30) 7-8pm
Challenge 3:	**+3.0***	**+13***	-20 (-83 to 51) 6-7pm	**+2.4***	**+10***	-8 (-69 to 60) 6-7pm
Challenge 4: #	**-7.0***	**-35***	-40 (-90 to 18) 5-6pm	**-3.0***	**-16***	-19 (-77 to 46) 6-7pm

Notes:
* Value calculated as mean of hourly effect estimates across targeted peak hours
# Challenge 4 targeted the central peak hours only (17:00 to 19:00)

In the dynamic pricing trial held during trial period 3, the maximum treatment effects were observed during the high incentive period. The reductions, averaged over the peak hours (16:00 to 20:00), were as follows:

- Opt-in recruitment (TG3): a 2.6 percent reduction (-17W, 90% CI -71 to 43);
- Opt-out recruitment (TG4): a 7.1 percent reduction (-44W, 90% CI -97 to 15).

Table 14: Summary of observed treatment effects, dynamic pricing rebate intervention, trial period 3

Incentive	TG3: Opt-in recruitment			TG4: Opt-out recruitment		
	%	Wh/h (90% CI)	Timing of max. effect*	%	Wh/h (90% CI)	Timing of max. effect
Low incentive	**-2.1**	-12 (-62 to 42)	w/c 05-11-2018	**-5.5**	-31 (-78 to 21)	w/c 05-11-2018
High incentive	**-2.6**	-17 (-71 to 43)	w/c 19-11-2018	**-7.1**	-44 (-97 to 15)	w/c 19-11-2018

* Timing of maximum effect refers to Wh effects. In percentage terms, the maximum value observed during the low incentive period for the opt-in group was the week commencing 12th November 2018.

Second, a customer typology has been developed to represent the variability in electricity demand during peak hours - and daily demand profiles - across different households. Out of a selection of candidate variables, the three characteristics found to best predict the variability in household consumption during peak hours were used to define the customer types: number of people, number of bedrooms and primary heat source. An iterative process was used to balance the requirements of the customer type definitions: to represent as much variability in peak-hours demand and profile shape as possible, while maintaining a sample size sufficient to generate representative profiles for each type.[60] This typology provides a framework for greater disaggregation of the loads expected for different customers than that currently used by network planners, and consequently allows better representation of different household types in modelling and simulations. Table 15 shows the customer types created for the SAVE customer model.

Table 15: Final customer type categories represented in the SAVE Customer Model

<table>
<tr><th></th><th></th><th colspan="4">Number of bedrooms</th></tr>
<tr><th>Heat source</th><th>Number of people</th><th>0-1</th><th>2</th><th>3</th><th>4+</th></tr>
<tr><td>Gas</td><td>1</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td></tr>
<tr><td>Gas</td><td>2</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td></tr>
<tr><td>Gas</td><td>3</td><td colspan="2">✓</td><td>✓</td><td>✓</td></tr>
<tr><td>Gas</td><td>4+</td><td colspan="2">✓</td><td>✓</td><td>✓</td></tr>
<tr><td>Electric</td><td>All</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td></tr>
<tr><td>Other</td><td>All</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td></tr>
</table>

Third, household electricity demand profiles have been produced for a number of time periods (scenarios) for each customer type using the representative SAVE sample. While only a number of scenarios have been created, the dataset and method supports the creation of profiles for any week (or range of individual days) between January 2017 and December 2018. In addition, intervention impact profiles have been generated for a number of the interventions tested during the trials, providing estimated treatment effects for the SAVE interventions. Combining these outputs allows network planners to simulate loads on low-voltage network assets under both 'baseline' and 'intervention' conditions. By using the customer typology in conjunction with publicly available Census Output Area statistics, the Customer Model provides an innovative tool with which network planners can allocate customer load profiles to specific network geographies, simulate customer loads, and assess the feasibility of a number of demand-side interventions in any specific location. The process is visualised in Figure 27 below. From left to right, the figure provides an example of using Census statistics to allocate customer demand and intervention profiles to an element of the network. The Customer Model creates the load profiles for each customer type for a specified scenario: in this case an intervention using LED lighting upgrades (shown within the green box).[61] The appropriate quantities of each customer type – and their associated profiles - are selected according to the Census statistics to be applied within the network simulation. In this example only six customer profile types exist in the target network and are selected (within the orange-shaded box).

Figure 27: Allocating customer load profiles using Census statistics

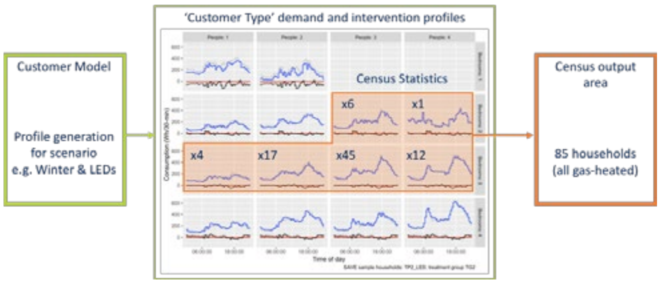

60 In order to maintain the sample size for less common customer types, some categories of the were combined, for example households with 3 or 4 persons in dwellings with 0-1 bedrooms were combined with those in 2-bedroom dwellings. For more details of the development of the customer typology refer to SDRC 2.3: Customer Model Final Report, available at https://www.ssen.co.uk/save/.

61 Profiles for only gas-heated households are shown in the example for clarity.

Fourth, the Customer Typology method provides a replicable and scalable process. This allows other DNOs to use the process to generate customer types and load profiles specific to the mix of households within their own geographical region. SAVE has provided the methodology and process, along with the data requirements for replicating the Customer Model more widely across other UK DNOs. The Customer Model provides an expandable framework, allowing suitable, high-quality demand data from other sources (with the appropriate linked household characteristics) to provide input to the model in order to expand and refine both the customer typology and output load profiles.

6

APPENDICES

6.1 Dynamic pricing rebate tariff analysis

6.1.1 Sample size

The number of households within each sub-group are shown in Table 16. The household counts include only those with electricity consumption (Loop) data.

Table 16: Number of households in treatment sub-groups, dynamic pricing rebate trial, Week 1

	Control	Dynamic pricing	Non-participant	Total
1: None	709	0	0	709
3: Dynamic pricing rebate, opt-in	0	297	365	662
4: Dynamic pricing rebate, opt-out	0	633	24	657
Total	709	930	389	2028

Note that there is some attrition in the sample over the course of the trial period. Table 17 shows the household counts for the final full week of the trial period (Week 13).

Table 17: Number of households in treatment sub-groups, dynamic pricing rebate trial, Week 13

	Control	Dynamic pricing	Non-participant	Total
1: None	692	0	0	692
3: Dynamic pricing rebate, opt-in	0	292	371	663
4: Dynamic pricing rebate, opt-out	0	631	20	651
Total	692	923	391	2006

6.1.2 Consumption in treatment subgroups

Figure 28 shows the sub-groups within trial group 3 (opt-in). This treatment group used an opt-in recruitment method: members of the group did not participate as default, shown as 'Non-participant (Out)', and were required to opt-in to take part and receive rebate incentives, shown as 'TOU-rebate (In)'. The consumption observed in the sub-groups show that households choosing to take part (opting-in) have, on average, lower consumption than the group average (and control group).

Figure 28: Mean 15-minute electricity demand by intervention sub-group, TG3 (Opt-in), peak hours

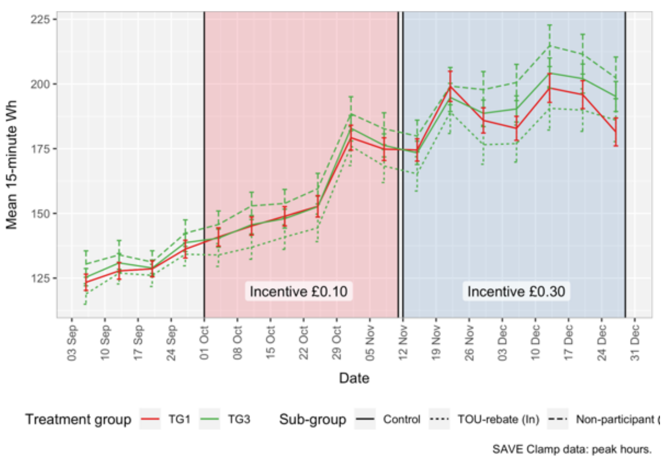

Figure 29 shows the sub-groups within trial group 4 (opt-out). This treatment group used an opt-out recruitment method: members of the group were enrolled to participate as default, shown as 'TOU-rebate(In)', but could opt-out if they did not want to take part to receive rebate incentives, shown as 'Non-participant (Out)'. In this treatment group, the mean consumption of the participating households is shown to be very close to that of the overall group average and marginally lower than that of the control group. The mean consumption of those households opting-out is more variable than the mean of the treatment group overall and control group means.

Figure 29: Mean 15-minute electricity demand by intervention sub-group, TG4 (Opt-out), peak hours

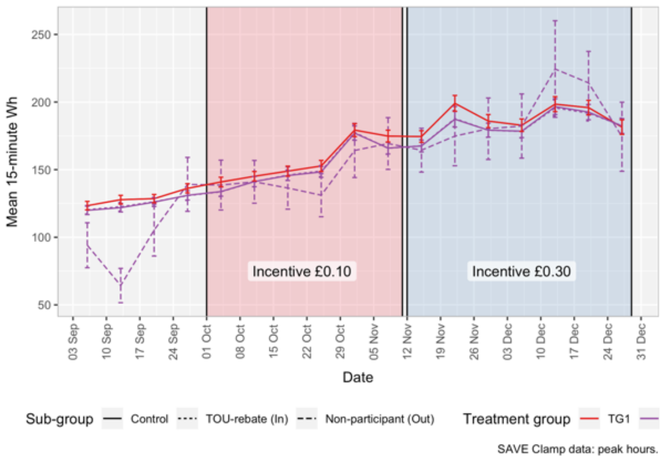

6.2 Customer type regression model results

Table 18: Regression results, original coding of household and dwelling size

	Dependent variable:	
	People x Rooms (1)	**People x Bedrooms (2)**
People 2	0.488*** (0.426, 0.549)	0.504*** (0.442, 0.566)
People 3	0.681*** (0.609, 0.753)	0.683*** (0.609, 0.756)
People 4	0.787*** (0.714, 0.860)	0.793*** (0.718, 0.867)
People 5	0.929*** (0.832, 1.027)	0.904*** (0.803, 1.005)
People 6	1.029*** (0.873, 1.186)	0.980*** (0.820, 1.140)
People 7	0.953*** (0.672, 1.234)	0.902*** (0.617, 1.186)
People 8+	0.847*** (0.494, 1.199)	0.810*** (0.454, 1.167)
Rooms 2	-0.554 (-1.223, 0.114)	
Rooms 3	-0.589** (-1.125, -0.054)	
Rooms 4	-0.457* (-0.963, 0.048)	
Rooms 5	-0.553** (-1.050, -0.055)	
Rooms 6	-0.431* (-0.926, 0.065)	
Rooms 7	-0.388 (-0.884, 0.107)	
Rooms 8	-0.271 (-0.767, 0.225)	
Rooms 9+	-0.147 (-0.641, 0.348)	
Bedrooms 1		-0.859 (-1.975, 0.258)
Bedrooms 2		-0.901 (-2.014, 0.211)
Bedrooms 3		-0.755 (-1.868, 0.357)
Bedrooms 4		-0.584 (-1.696, 0.529)
Bedrooms 5		-0.500 (-1.616, 0.615)
Bedrooms 6		-0.434 (-1.563, 0.695)
Bedrooms 7		-0.693 (-1.825, 0.440)
Bedrooms 8		-0.802 (-1.959, 0.355)
Bedrooms 9		-0.590 (-1.736, 0.556)
Bedrooms 10		-0.730 (-1.931, 0.471)
Bedrooms 11		-0.899 (-2.087, 0.289)
Bedrooms 12		-0.339 (-1.556, 0.879)
Bedrooms 13		-1.387* (-2.959, 0.186)
Constant	-1.405*** (-1.901, -0.908)	-0.995* (-2.109, 0.118)
Observations	2,991	2,985
R^2	0.271	0.260
Adjusted R^2	0.267	0.255
Residual Std. Error	0.562 (df = 2975)	0.567 (df = 2964)
F Statistic	73.782*** (df = 15; 2975)	51.946*** (df = 20; 2964)
Note:		*p <0.1** p<0.05 ***p<0.01

Table 19: Regression results, three-variable model comparison

	Dependent variable:		
	Heat Source: 6 categories	**Heat Source: 5 categories**	**Heat Source: 3 categories**
People 2	0.512*** (0.452, 0.572)	0.512*** (0.451, 0.572)	0.506*** (0.446, 0.566)
People 3-5	0.779*** (0.717, 0.841)	0.777*** (0.715, 0.839)	0.771*** (0.709, 0.833)
People 6+	0.961*** (0.828, 1.093)	0.959*** (0.826, 1.091)	0.951*** (0.818, 1.083)
Bedrooms 3	0.177*** (0.123, 0.230)	0.177*** (0.123, 0.231)	0.179*** (0.125, 0.233)
Bedrooms 4	0.365*** (0.303, 0.427)	0.369*** (0.307, 0.431)	0.371*** (0.309, 0.433)
Bedrooms 5+	0.377*** (0.293, 0.460)	0.377*** (0.294, 0.461)	0.383*** (0.299, 0.466)
Elec. storage	0.353*** (0.252, 0.454)		
Elec. other	0.772*** (0.508, 1.037)		
Oil	0.521*** (0.381, 0.661)		
Solid/Biomass	0.202** (0.026, 0.377)		
Other	0.511*** (0.361, 0.661)		
Elec. all		0.404*** (0.309, 0.500)	
Oil		0.521*** (0.381, 0.661)	
Solid/Biomass		0.201** (0.026, 0.377)	
Other		0.511*** (0.361, 0.661)	
Elec. all			0.405*** (0.309, 0.500)
Other all			0.437*** (0.346, 0.527)
Constant	-1.976*** (-2.033, -1.918)	-1.976*** (-2.033, -1.918)	-1.973*** (-2.030, -1.915)
Observations	2,969	2,969	2,969
R^2	0.291	0.289	0.287
Adjusted R^2	0.289	0.287	0.285
Residual Std. Error	0.553 (df = 2957)	0.554 (df = 2958)	0.554 (df = 2960)
F Statistic	110.410*** (df = 11; 2957)	120.289*** (df = 10; 2958)	148.817*** (df = 8; 2960)
Note:			*p<0.1 **p<0.05 ***p<0.01

Table 20: Regression results, Customer Types, final categories (Census aligned)

	Dependent variable:
	Final Customer Types
CustomerTypeElectric-2B	0.062 (-0.189, 0.313)
CustomerTypeElectric-3B	0.370*** (0.133, 0.607)
CustomerTypeElectric-4+B	0.926*** (0.607, 1.244)
CustomerTypeGas-1P-0-1B	-0.861*** (-1.092, -0.630)
CustomerTypeGas-1P-2B	-0.681*** (-0.885, -0.477)
CustomerTypeGas-1P-3B	-0.566*** (-0.763, -0.369)
CustomerTypeGas-1P-4+B	-0.459*** (-0.688, -0.231)
CustomerTypeGas-2P-0-1B	-0.257* (-0.521, 0.008)
CustomerTypeGas-2P-2B	-0.181* (-0.373, 0.011)
CustomerTypeGas-2P-3B	-0.046 (-0.232, 0.139)
CustomerTypeGas-2P-4+B	0.100 (-0.088, 0.289)
CustomerTypeGas-3P-0-2B	-0.078 (-0.294, 0.138)
CustomerTypeGas-3P-3B	0.160 (-0.032, 0.352)
CustomerTypeGas-3P-4+B	0.334*** (0.134, 0.535)
CustomerTypeGas-4+P-0-2B	0.031 (-0.226, 0.287)
CustomerTypeGas-4+P-3B	0.289*** (0.101, 0.477)
CustomerTypeGas-4+P-4+B	0.519*** (0.331, 0.706)
CustomerTypeNA	-0.102 (-0.395, 0.190)
CustomerTypeOther-0-2B	-0.008 (-0.256, 0.240)
CustomerTypeOther-3B	0.443*** (0.213, 0.673)
CustomerTypeOther-4+B	0.629*** (0.401, 0.857)
Constant	-1.240*** (-1.418, -1.061)
Observations	2,991
R^2	0.290
Adjusted R^2	0.285
Residual Std. Error	0.554 (df = 2969)
F Statistic	57.654*** (df = 21; 2969)
Note:	*p<0.1 **p<0.05 ***p<0.01

6.3 Aggregation effects

The comparison of regression models suggest that the reduced-coding of the Customer Type variables appear to capture a reasonable amount of the intra-household variation in the log(mean kWh), however:

- the models use log(kWh) not kWh (which the Network Model requires) and
- this does not imply that aggregating household level kWh data to mean(kWh) and sd(kWh) will adequately represent the distribution/variance in household evening peak electricity consumption of the resulting groups.

Layers of aggregation (means of means and standard deviation calculations on means of means) may well suppress variation to the extent that the Network Model ceases to be a realistic representation of the actual variance experienced by the network and thus under-estimate the 'headroom' required.

To make this clear, the results of calculating the mean and standard deviation for the kWh consumption values at different levels of aggregation are presented in the following sections.

6.3.1 Means of all observed 15 minutes

Table 21 shows mean kWh consumption using the original 15-minute observations (Table 17). The mean kWh across the trial (BMG) groups is very similar as are the standard deviation and thus the upper and lower confidence intervals (*C.I. Lower* and *C.I. Upper* in tables) which give 95% confidence bounds for the mean. Thus, in the case of the control group (Group 1), if the SAVE sample were repeated 100 times, in 95 of those samples the mean evening peak consumption calculated in this way would be expected to be between 0.212 and 0.213 kWh per half hour. Finally, the coefficient of variation (COV, ratio of standard deviation to mean) is greater than 1. This has no real meaning on its own but is useful for comparisons with other calculation methods or distributions (as below).

Table 21: Mean and s.d. of 15-minute Wh observations by trial group (16:00 - 20:00 period), all observations

Trial Group	Observations	Households	mean (kWh)	S.D. (kWh)	C.I. Lower	C.I. Upper	COV
BMG Group 1	411046	979	0.213	0.246	0.212	0.213	1.157
BMG Group 2	431591	1025	0.216	0.260	0.215	0.217	1.205
BMG Group 3	365296	863	0.220	0.255	0.220	0.221	1.159
BMG Group 4	383821	905	0.213	0.258	0.212	0.214	1.210

6.3.2 Means of all 'observed' half-hours

Table 22 repeats this analysis but is based on all half-hourly values for each household (which has already suppressed 15-minute level variation by summing pairs of 15 minute observations) across the period selected. This is made clear by the substantially decreased number of observations compared to the number of households. Whilst the mean values are higher (as would be expected), the standard deviation is also slightly higher and this is reflected by the (very) slightly wider 95% confidence intervals. It should also be noted that the COV is also similar. This suggests that analysis at the half-hourly level would not substantially mask the underlying variation captured by the 15 minute data.

Table 22: Mean and s.d. of half-hourly Wh by trial group (16:00 - 20:00 period), all half-hourly observations

Trial Group	Observations	Households	mean (kWh)	S.D. (kWh)	C.I. Lower	C.I. Upper	COV
BMG Group 1	231438	979	0.377	0.436	0.376	0.379	1.154
BMG Group 2	242867	1025	0.384	0.463	0.382	0.385	1.207
BMG Group 3	205546	862	0.392	0.454	0.390	0.394	1.160
BMG Group 4	216080	905	0.378	0.460	0.377	0.380	1.214

6.3.3 Means of household level means

Finally, Table 23 shows the mean half-hourly kWh consumption for the evening peak period based on the overall household mean evening peak consumption as used in the regression models. Thus, each household has only 1 observation - the mean of all half-hours in the period selected. Thus the table shows the mean of these household means.

As before the mean (of means) kWh shown in the table are relatively similar for each trial group, as are the standard deviation values. However comparing with the previous table shows that whilst the mean is similar, the standard deviation is much lower, as is the COV demonstrating that the within-household half-hourly variation has been substantially suppressed.

Table 23: Mean and s.d. of half-hourly mean Wh by trial group (16:00 - 20:00 period), using household means

Trial Group	Observations	Households	mean (kWh)	S.D. (kWh)	C.I. Lower	C.I. Upper	COV
BMG Group 1	979	979	0.371	0.255	0.355	0.387	0.686
BMG Group 2	1025	1025	0.380	0.289	0.362	0.398	0.761
BMG Group 3	862	862	0.386	0.271	0.368	0.404	0.702
BMG Group 4	905	905	0.373	0.276	0.355	0.391	0.739

Figure 30 makes the variance suppression very clear by showing the coefficient of variation (COV) for each calculation method on the same chart. There is very little difference between the COV values using the observed 15/30-minute values. However, the COV value for the 'mean of the means' is substantially lower. It seems likely therefore that using a single standard deviation value calculated over a set of household means will substantially under-represent the actual variation. This implies that any aggregation to household types needs to be a per-half-hour calculation for all half hours and not a calculation based on household level means over a group of half-hours.

Figure 30: Comparison of coefficient of variation of mean by calculation method

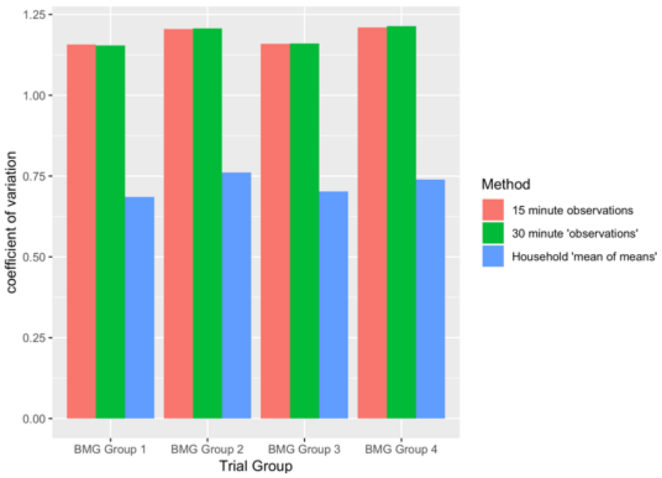

To illustrate the distributions that the Network Model needs to mimic for each customer type if it is to re-create the variance found in the original data at the 30-minute level, Figure 31 uses a violin plot to show the distribution of all *individual* 30-minute kWh observations in the period for each customer type (defined by recoded household size and number of bedrooms).[62] Note that this will already have suppressed variation at the 15-minute level and does not represent the small cell-count uncertainty highlighted above.

Figure 31: Distribution of 30-minute observations by 'customer type', all households

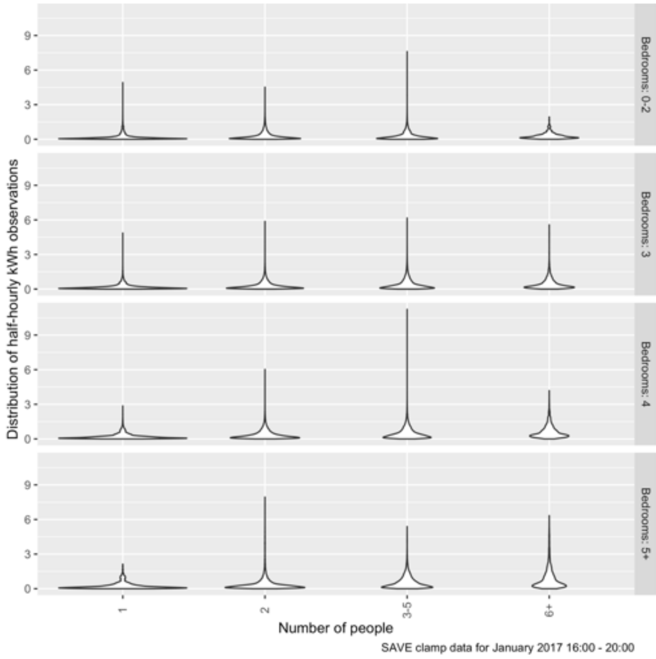

Inspection of the distributions shown in Figure 31 confirm the positive skew present in the 30-minute consumption data for all customer types. Due to the non-normal (positively skewed) distribution of half-hourly consumption observations within each customer type, caution should be exercised around the mean and standard deviation values. These measures do not best represent the distribution of the underlying data and may lead to the inaccurate representation of risk in external models.

6.3.4 Summary

The tables and charts above confirm that increasing levels of aggregation will increasingly suppress the variance in the data and this is especially true where observations are averaged over both time and households at the same time. This implies that any aggregation to household types needs to be a per-half-hour mean and standard deviation calculation for each half-hour under consideration across all households within each 'type' (as currently implemented in the customer model and detailed in Section 3.3).

62 Note that the customer type definition shown in the figure differs to the final typology as the analysis was produced prior to the re-alignment with Census data.

6.4 Intervention impact profiles

Figure 32 below illustrates the intervention impact profiles for standard deviation for gas-heated households. In the figure, the black lines show the 'impact' profile, grey profiles show the standard deviation profiles for the 'baseline' (control group), and the blue line show the resulting profile for the intervention group. The profiles are modified with exactly the same values for each dwelling size (i.e. black lines indicate impact profiles are the same in each row). However, as the baseline profiles for the full set of customer types are used, the resulting standard deviation profiles are unique for each.

Negative values are shown within the profiles for 2-person 1-bedroom and 1-person 4+ bedroom households, highlighting a limitation of this method: the generation of negative standard deviation values. These profiles will need further modification prior to use in modelling or simulation applications. In this case is would be preferable to substitute the 'baseline' standard deviation profile for these customer types. Generally, the profiles exhibit a reduction in standard deviation during times of the day when load reduction is observed – shown as negative values for the black profile in Figure 32 - as would be expected with an associated reduction in the variance of consumption values.

Figure 32: Example of intervention group standard deviation profiles constructed using baseline and intervention impact, LED upgrades treatment (weekdays)

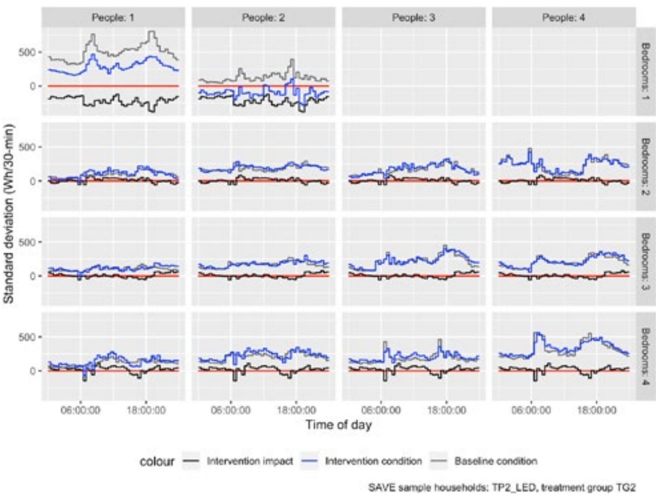

Figure 33 below illustrates the intervention impact profiles for standard deviation for non-gas-heated households. Very large values are observed and as with the generation of the values for impact on mean consumption, these result from a combination of high background variation in demand, unrepresentative households, and small samples.

Figure 33: Example intervention impact profiles for non-gas heated households showing difference in standard deviation from control group, LED upgrades treatment group

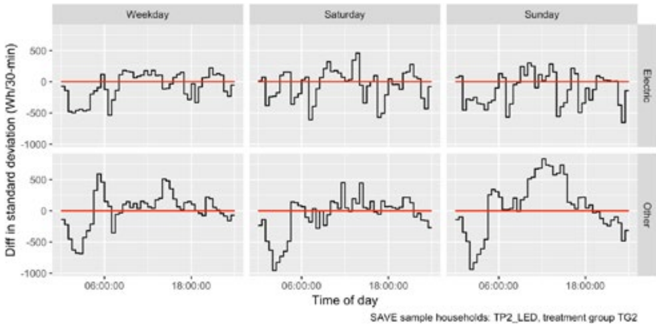

6.5 Data requirements and model build

The requirements detailed in this section provide the minimum data inputs to assemble and generate the Customer Model outputs described in this report: the customer type household demand profiles and intervention impact profiles. These inputs therefore also set out the data required to replicate the Customer Model in order to apply the SAVE Customer Type methodology to another distribution network operator area of operations and/or geographical region.

6.5.1 Household data

In order to define the customer type categories, a number of household attributes are required, including (at a minimum): household size (number of persons), dwelling size (number of bedrooms), and primary heating source (fuel). These characteristics should be supplied for all households contributing electricity consumption data. For the SAVE sample households, this information was collected and supplied by the fieldwork contractor (BMG Research) through a survey conducted with trial participants. The data file included socio-economic and demographic data for the households participating in the fieldwork. Update surveys were conducted periodically throughout the project to ensure that basic household attributes such as number of occupants were updated. A data processing script implemented within R assembles the appropriate survey data file by combining the original survey with the relevant update files according to the time period under consideration.

6.5.2 'Loop' electricity consumption data

The Customer Model requires household electricity consumption data This was provided under the SAVE project by the projects data supplier, Navetas. The 'Loop' data provided cumulative watt-hour (Wh) readings observed at 15-minute intervals for each participating household.

Prior to analysis and the generation of customer type demand profiles, the Loop data is pre-processed to remove erroneous and interpolated consumption values[63]. The 15-minute consumption data was aggregated to 30-minute consumption totals for use in generating customer type demand and intervention impact profiles.

6.5.3 Census 2011 output area data

Although Census data is not used directly within the final (non-spatial) Customer Model, the construction and definition of the Customer Types was aligned to match the categorisation within the Census statistics available at the *Output Area* (OA) scale. This allows the use of Census data to allocate appropriate 'baseline' demand and intervention impact profiles to households of each type and in the correct proportions within each geographical area in the SAVE study region. This functionality is provided by the *Census Interface*.

This is in the form of output area (OA) level tables containing aggregate household counts for a range of household and household response person attributes. The data was downloaded from Nomisweb (https://www.nomisweb.co.uk/census/2011/data_finder).

Census data was used more directly within the initial spatial microsimulation element of the SAVE Customer Model[64]. This element of work was not taken forward to the final Customer Model as the functionality was separated out and provided by the Census Interface[65].

63 For further details of 'Loop' data pre-processing, see Rushby and Harper (2018), SAVE Loop Energy Saver Data Cleaning and Pre-processing. SAVE Project Report, University of Southampton.

64 Refer to SDRC 2.2 for details of the spatial microsimulation module of the Customer Model and the use of Census data within this modelling approach.

65 Refer to SDRC 8.5 & 8.6 Pricing Model, Customer Model and Network Model for details of the implementation of the Census Interface. Available at https://www.ssen.co.uk/save/