Department
of Energy &
Climate Change

Appendix D – Frequently Asked Questions

Electricity Demand Reduction pilot

M&V manual

Version 0.1 01 August 2014**Error! Bookmark not defined.**

URN 14D/279

Any enquiries regarding this publication should be sent to us at edr-project@decc.gsi.gov.uk

Contents

Frequently Asked Questions

Hints and Tips for Calculating Baselines

How do I know if I have enough readings?

The three Template Calculation spreadsheets contain a cell labelled:

“Accuracy requirement met?”

If the answer is ‘No’, you may find that adding more readings changes the answer to ‘Yes’. If that is not the case then it might be true that the readings you are taking are have too large a variance to be able to quantify the savings easily and using regression analysis is more appropriate. This means you will have to identify and source independent variable data as described in the baseline section of the M&V manual.

Where do I get data from?

For data on common independent variables, there are several publicly available sources. For example, www.degreedays.net can be used to provide degree days data, and http://www.timeanddate.com/worldclock/sunrise.html can be used to provide data on the time between sunrise and sunset for different regions. There are also services that can be bought for similar or more detailed data.

What if my baseline is not accurate enough?

A first step you might take is to see whether there are any ‘outliers’. Firstly, it is useful to consider whether there have been any periods of atypical usage that are covered by any of your data points, such as bank holidays, power outages, building work, etc.

You might find it helpful to plot the data, so that outliers might be identifiable by visual inspection.

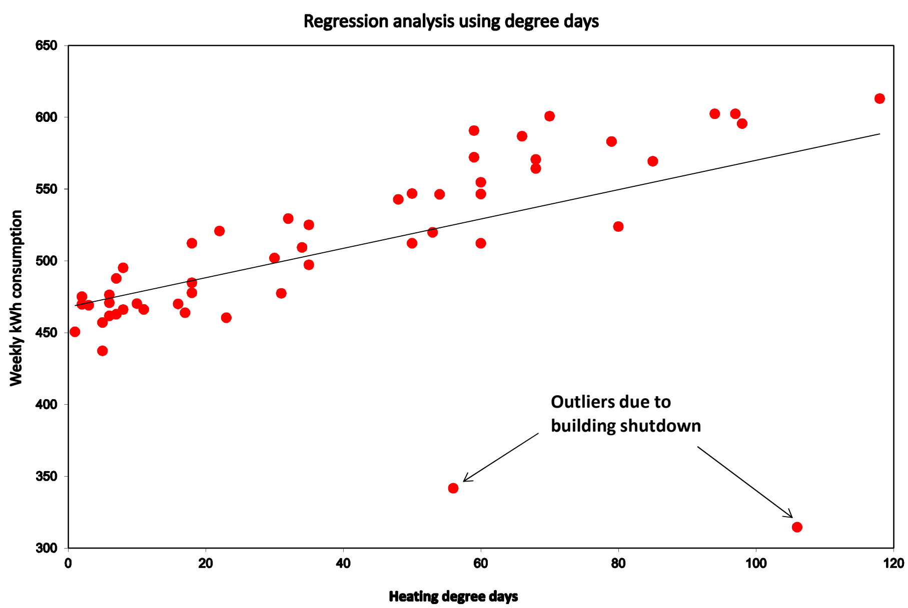

Such outliers will be apparent because they are very different from the trend shown by the rest of the data, for example they are much higher or (more commonly) lower than other data points for similar independent variables. Such outliers can be removed if there is justification for doing so, e.g. the building was not occupied due to public holidays or refurbishment.

What are peak relevant savings?

As outlined in the participant handbook, a key aim of the pilot is to test whether EDR could participate in the Capacity Market in the future. In order to participate in the Capacity Market, EDR would have to deliver reductions in electricity demand when demand on the electricity system is likely to be highest. The pilot therefore aims to explore the extent to which it is possible to incentivise demand reduction that is relevant to the winter peak period. This is defined as November-February (inclusive), business days, 4pm-8pm. To be eligible to participate in the pilot your project will need to provide at least 100kW of savings averaged over the peak period.

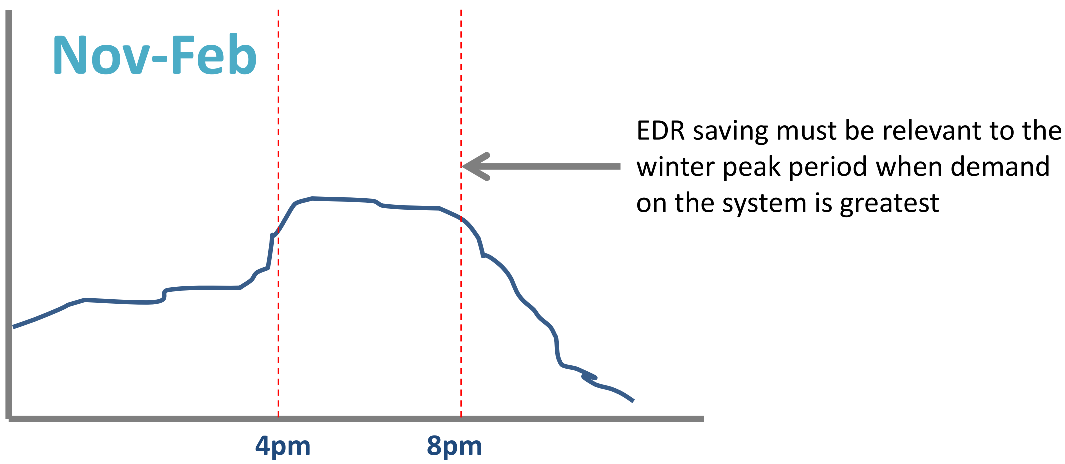

Many measures being bid in will also deliver savings outside this winter peak period. However, for the purposes of calculating and bidding in savings for the pilot it is only peak-relevant savings that are eligible.

For your project, it is likely you have already calculated or have a good idea of the electricity savings over time (kWh) from your project. You can use this information to work out if your project meets the payback requirements to participate in the pilot (see the 'Site payback' section of the manual for more detail).

To participate in the pilot you also need to calculate the kW (or capacity) value of peak relevant savings that your project will deliver. This capacity value will need to take account of the extent to which measures included in your bid are relevant to the peak period. For example, a measure saving 8 kW, but only operational for half of the peak period would be credited with an average capacity value of 4 kW.

Why is a whole year of data suggested for baselining seasonal variations when we're only interested in peak savings in November to February?

It is important to try to capture all seasonal conditions when defining a baseline model, otherwise the results are less rigorous. Technically, the use of only data from November – February would make it more likely that "extrapolation" would be needed in order to quantify savings. It is that extrapolation that reduces the confidence we can have in the results. By taking a whole year, we minimise the need for extrapolation and ensure that the results are more reliable than they would otherwise be.

Why is it acceptable to use data which covers periods outside the peak demand hours, such degree days, which cover a whole day?

Although the peak demand hours are of interest, there is not usually any independent variable data that is captured at high enough frequency to isolate those hours. If that is the case, then metrics like degree days do give some indication of the conditions at the site, and therefore are likely to improve the rigour of the calculation compared to using no independent variables.

What are Degree Days?

Degree days can be used to normalise for the effect of seasonal variations on a building's energy use. They can be used to identify trends in consumption, but also provide a basis for measuring significant changes in energy efficiency with a greater level of accuracy than a simple year-on-year comparison.

In numerical terms, degree days represent the cumulative difference between external temperature and a 'baseline' temperature over a period of time. For example, if a week experiences particularly cold weather, there will be a large difference between actual temperatures and the baseline and a relatively high degree days total.

The following chart illustrates how the degree days total accumulates – the shaded area between the baseline temperature, in this case 15.5℃, and the actual temperature is the degree day value for the period. Note that when the actual temperature is above the baseline then the degree days are taken as zero.

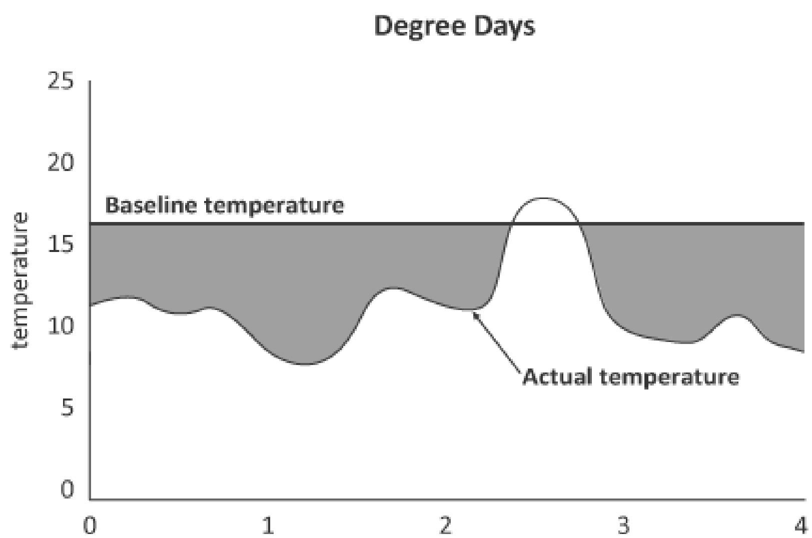

In the previous example, the degree days are calculated when actual temperatures are below the baseline temperature - this type of calculation means that they are heating degree days and would be appropriately used to develop a baseline for equipment used for space heating, or if heating makes up a significant proportion of the overall building load.

Degree days can also be used for equipment used for cooling, or for at the whole building level where cooling load is substantial, for example in air-conditioned offices or supermarkets with a large refrigeration load. In this case, we are more interested in calculating how much higher external temperatures have been above a baseline, which will give us a total for cooling degree days. This calculation can be carried out similarly to heating degree days for a given baseline temperature.

It should be noted that base temperatures are building specific, although there are values that are commonly used, e.g. 15.5 for general use with most buildings and 18.5 for hospitals. Publically available data for degree days is usually calculated for different base temperatures in half degree intervals. You may find that using degree days with a base temperature different from 15.5 gives you a more accurate baseline, which can be assessed via the standard error in the regression analysis template. The standard error will be lowest for the degree days for a base temperature that is closest to the actual base temperature of your building.

What if I want to aggregate savings across lots of sites?

Those who are aggregating savings across multiple sites can use a sampling approach to measure the demand reduction of their project. Where there are multiple implementations of the same EDRM, a suitably sized sample can be taken to measure the demand reduction and subsequently scaled up to represent savings across a number of sites. This is acceptable provided that the project meets the relevant payback criteria. Both technologies and time of use must be sufficiently similar across the different sites to enable a sample to act as a reliable and reasonable proxy to enable savings to be calculated across all of the sites.

What do I do in cases where technology is not definitively off or on e.g. production ramps up and down?

There are some technologies where the judgement about whether equipment is definitively on or off is not necessarily straightforward. For example, the lighting load in a building may begin to phase down beyond 6pm as people leave an office building. Or a motor or pump may ramp up or ramp down over time.

How loads that cycle on or off are taken into account will depend on your chosen measurement method. For some you will have to estimate the number of hours that equipment is on for. In instances where there are multiple installations and the exact time of use varies across the installations within a site to be regarded as ‘on’ at least 60% of the technology in question e.g. lighting should be operational to be regarded as on.

Questions on the Statistical Terms used in the Manual

What is the “Standard Error”?

There are two common uses of the phrase “Standard Error” which mean different things. These are:

The Standard Error of the Estimate; and

The Standard Error of the Mean

In this context, we are interested in The Standard Error of the Estimate which can be stated formally as:

$$SE_{\hat{Y}} = SQRT(\Sigma(\hat{Y}_i - Y_i)^2/(n - p - 1))$$

Where:

Y_i are the observed values (in our case, this will usually be energy or power data).

$\hat{Y}_i$ are the predicted values of Y from the regression model n is the number of observations

p is the number of independent variables in the regression equation

What is a p-value?

A p-value is technically a probability. In the context of energy savings, they are most often used to determine whether there is a strong relationship between energy consumption/power drawn and a particular variable.

The important thing to remember is that the more relevant the variable is, the smaller its p-value will be.

Technically, what is happening is, a statistical hypothesis test is being carried out, such that the null hypothesis is that there is no relationship between the variable and energy consumption/power drawn. P-values represent the probability of the null hypothesis actually being true in the context of the available data, so small values indicate that the null hypothesis (no relationship) is false, i.e. there is actually a relationship between the variable and energy consumption/power drawn, therefore that variable should be taken into account.

A standard threshold for p-values, which is often said to indicate “Statistical significance”, is 0.05. So if the p-value associated with a particular variable is below 0.05, that variable is sometimes said to be significant at the 95% level of significance, indicating that the variable in question should be taken into account.

How does the Sample Size calculator work?

Where a sample of data is to be taken, the sample size calculator will output an appropriate sample size given four pieces of information:

- The size of the population (so if you’ve got 400 light fittings from which you want to draw a sample, set the population size to 400).
- The cv (coefficient of variation), this is given by the standard deviation divided by the average of the sample. When we haven’t started sampling yet, 0.5 is usually used.
- The confidence level. The recommended confidence level is 90%.
- The range of precision. The recommended precision level is 10%

Note, the easiest way to understand what is meant by confidence and precision is to consider what, ultimately, they will allow us to say about the sample:

To specify 90% confidence and 10% precision is to specify that the sample should be big enough that there will be 90% probability of the population mean (average) being within the range ±10% of the sample mean.

Once some you have some values, the cv in the sample size calculator can be updated to give a better and better idea of what the sample size should be.

A fuller explanation of the process the sample size calculator goes through is given below.

A process for calculating an initial sample size is given in the IPMVP as:

$$n_0 = (z^2 * cv^2)/e^2$$

Where:

- n0 is the initial estimate of the required sample size before sampling begins

- cv is the coefficient of variation, defined as the standard deviation of the readings divided by the mean. Until the actual mean and standard deviation of the population can be estimated from actual samples, 0.5 may be used as an initial estimate for cv.
- e is the desired level of precision. Precision refers to the error bound around the true estimate (i.e., ±x% range around the estimate). Confidence refers to the probability that the estimate will fall in the range of precision.
- z is the standard normal distribution value for the desired confidence level. For example, z is 1.96 for a 95% confidence level (1.64 for 90%, 1.28 for 80%, and 0.67 for 50% confidence).

For example, for 90% confidence with 10% precision, and a cv of 0.5, the initial estimate of required sample size (n0) is:

$$n_0 = (1.64^2 * 0.5^2)/0.1^2 = 67$$

Adjust initial sample size estimate for small populations. The necessary sample size can be reduced if the entire population being sampled is no more than 20 times the size of the sample. For the initial sample size example, above, (n0 = 67), if the population (N) from which it is being sampled is only 200, the population is only 3 times the size of the sample. Therefore the “Finite Population Adjustment” can be applied. This adjustment reduces the sample size (n) as follows:

$$n = n_0 N / (n_0 + N)$$

Applying this finite population adjustment to the above example reduces the sample size (n) required to meet the 90%/±10% criterion to 50.

The sample size calculator allows the user to enter the confidence, precision, coefficient of variation and population, and delivers the recommended sample size as per the calculations described above. Resulting calculations are rounded up to the nearest whole number.

kW vs kWh Explained

The usual ‘currency’ in which energy savings are measured is the kWh, that is, electricity use over time. For example, a motor running at 2 kW for an hour will use 2 kWh of electricity.

- A kW quantifies the rate at which energy is used
- A kWh quantifies the total energy consumption, i.e. power x time

Take the following example:

A 100-watt light bulb on for 10 hours consumes 1000 watt-hours (or 1kWh)

Power x Time = Energy Consumption

x = 1000 watt-hours (1 kWh)

100 watts 10 hours

For the entirety of the time the light bulb is on, it demands 100 watts (or 0.1 kW)

Consider then ten 100-watt light bulb burning for only 1 hour. They will also consume 1000 watt-hours (1kWh).

Power x Time = Energy Consumption

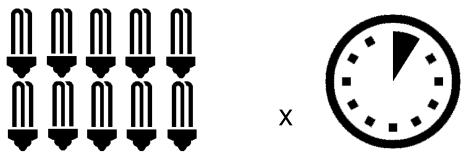

x = 1000 watt-hours (1 kWh)

1000 watts 1 hour

However, the demand whilst the light bulbs are on is 1000 watts (or 1 kW).

Calculating demand from kWh

For some of the measurement approaches, the data you will use will be in kWh, for example if it is taken from utility meter data or from submeters set to record kWh. These kWh values will need to be converted into an average kW demand value over peak hours. Appended to this manual are template spreadsheets that will help you to do this for the measurement methods described in the later sections of the manual (specifically in section 4.3).

For example, say we want to calculate the peak demand for a whole building for a given week in November and the total kWh value recorded by utility half hourly data for 4 – 8 pm Monday to Friday is 3000 kWh.

The average kWh demand = 3000 ÷ 5 (no. of days) ÷ 4 (no. of hours)

= 150 kW

Note that this doesn’t provide instantaneous demand, which may be available with some types of metering, but provides a proxy in the form of an average value for peak hours.

Deemed savings

What if the technology I wish to replace is not on the deemed savings list?

Please use one of the alternate M&V approaches detailed in this manual. Only technologies that are on the list can be deemed in the pilot.

Where can I find more information on the calculations and assumptions used in the deemed savings methods?

Please refer to the supporting Technical data/calculation sheets for each specific technology.

What evidence do I need to submit with my application?

Please refer to the Deemed Savings Manual (Appendix B).

Am I allowed to use the deemed savings method for projects across multiple sites?

Yes – but where the original equipment or the replacement equipment are different across the sites, or where the operational characteristics are different (e.g. different operating hours) then the equipment for each site should be grouped, and each group should have a separate line on the spreadsheet calculator.

Can I combine a number of technology types as part of a deemed project?

Yes – unless otherwise stated in the exclusions section set out in Appendix B: Deemed Savings Manual.

Department of Energy & Climate Change
3 Whitehall Place
London SW1A 2AW
www.gov.uk/decc
URN 14D/279