ACIL ALLEN CONSULTING

REPORT TO
AUSTRALIAN ENERGY MARKET OPERATOR

25 JULY 2014

CONNECTION POINT DEMAND FORECASTING

FINAL REPORT OF INDEPENDENT ADVISER

ACIL ALLEN CONSULTING PTY LTD
ABN 68 102 652 148

LEVEL FIFTEEN
127 CREEK STREET
BRISBANE QLD 4000
AUSTRALIA
T+61 7 3009 8700
F+61 7 3009 8799

LEVEL TWO
33 AINSLIE PLACE
CANBERRA ACT 2600
AUSTRALIA
T+61 2 6103 8200
F+61 2 6103 8233

LEVEL NINE
60 COLLINS STREET
MELBOURNE VIC 3000
AUSTRALIA
T+61 3 8650 6000
F+61 3 9654 6363

LEVEL ONE
50 PITT STREET
SYDNEY NSW 2000
AUSTRALIA
T+61 2 8272 5100
F+61 2 9247 2455

SUITE C2 CENTA BUILDING
118 RAILWAY STREET
WEST PERTH WA 6005
AUSTRALIA
T+61 8 9449 9600
F+61 8 9322 3955

ACILALLEN.COM.AU

Contents

List of figures

List of tables

1 Introduction

In 2013 the Australian Energy Market Operator (AEMO) engaged ACIL Allen Consulting (ACIL Allen) to develop a methodology that AEMO could use to produce forecasts of maximum electricity demand at the transmission connection point (CP)[1] level in the National Electricity Market (NEM). That methodology is summarised in ACIL Allen's report to AEMO of 26 June 2013 (ACIL Allen's methodology report).

Since the release of the report, AEMO has used ACIL Allen's methodology to prepare forecasts of maximum demand at all connection points in New South Wales and Tasmania.[2] As it prepared the forecasts, AEMO obtained ongoing advice and assistance in two ways:

1. ACIL Allen acted as an independent adviser to AEMO. In practice this amounted to two broad tasks:
   a) ACIL Allen advised AEMO on a series of practical issues associated with implementing ACIL Allen's methodology and validating the results
   b) later in the process, assisted AEMO in responding to issues raised in the 'red flag reviews'.
2. Frontier Economics acted as an independent reviewer of AEMO's work. It conducted 'red flag reviews' in which it raised concerns relating to the forecasts that AEMO was developing. It also conducted a full review of the forecasts, though ACIL Allen only commented on the red flag reviews.

On completion of the forecasting process, AEMO provided ACIL Allen with a summary of the way that the connection point forecasts had been prepared. AEMO asked ACIL Allen to confirm that ACIL Allen's methodology described in that summary was consistent with ACIL Allen's methodology described in ACIL Allen's report of June 2013.

The purpose of this report is to provide that confirmation. It is structured as follows:

- chapter 2 provides an overview of ACIL Allen's methodology
- chapter 3 compares AEMO's implementation of ACIL Allen's forecasting methodology with the original methodology as it was prepared. That chapter concludes that AEMO has succeeded in applying ACIL Allen's methodology, albeit with a few reasonable modifications.
- chapter 4 summarise the issues raised by Frontier Economics in its 'red flag' reviews relating to the weather normalisation procedure and ACIL Allen's advice to AEMO regarding those issues.

---

1 That is, the point where the transmission network meets the distribution network. In some jurisdictions this is referred to as a terminal station.

2 That is, the point at which the electricity transmission network connects to the electricity distribution network. In some cases these are referred to as terminal stations.

2 Overview of ACIL Allen methodology

This section of the report provides an overview of ACIL Allen's methodology, which is depicted graphically in Figure 1. Each of the seven steps is addressed in turn.

Figure 1 **Forecasting methodology**

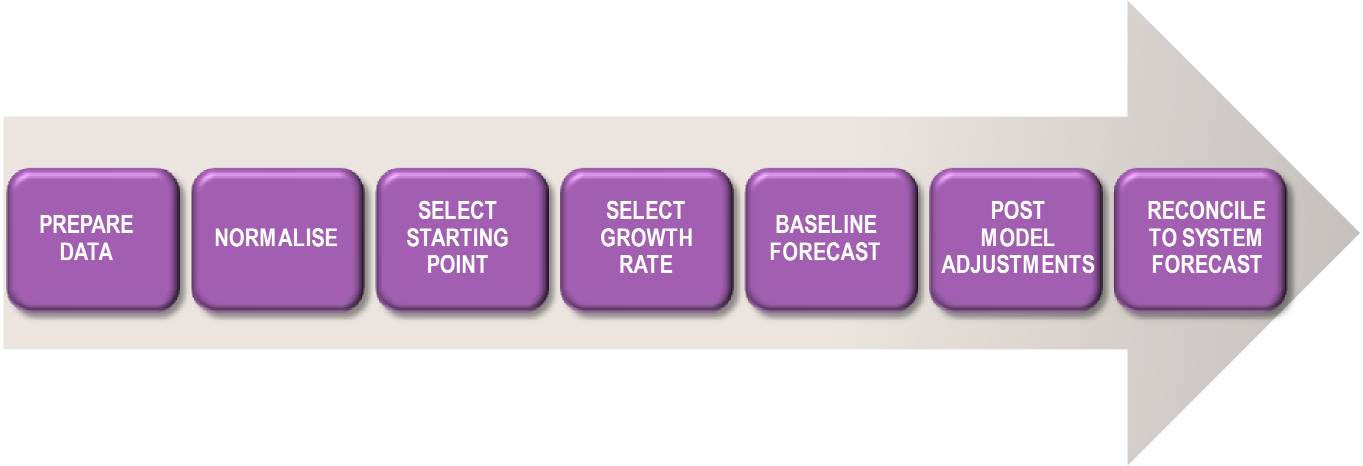

2.1 Prepare data

2.1.1 Demand data

The first step in ACIL Allen's methodology is to collect the necessary data and manage it appropriately. Three datasets are required.

The first is a time series of high frequency demand (15 or 30 minute interval) for each CP to be forecast. Ideally this time series should go back for at least 10 years.

These data should be well understood and should relate closely to what is being forecast. Three factors that may require adjustments to the historical data:

1. network configuration
2. block loads
3. output of embedded generation.

2.1.2 Weather data

The next dataset to collect is weather data for normalisation. Daily maximum and minimum ambient temperatures should be obtained as well as other weather variables, such as rainfall, humidity etc.

The BOM publishes weather data for many weather stations in Australia and one must be chosen and assigned to each CP for normalisation.

The choice of weather station and variable is empirical. It would be appropriate to test data from several stations to identify which is most closely correlated to demand at each CP, which may reflect microclimatic influences etc.

30 years (or more) weather data should be collected. It does not matter that this will be longer than the time series of demand data as a longer time series is required for weather normalisation.

Weather data will typically have some missing observations. These should be imputed, as discussed in ACIL Allen's methodology report.

2.1.3 Other data

A variety of other data may be useful in forecasting demand at CPs. Some of these data will not be used formally, but may assist in making the various judgements that must be made along the way. These include data relating to:

1. planned transfers between connection points (for either post model adjustments or altering network configuration)
2. changes in relationship between drivers and demand due to:
   a) changed use of demand management
   b) ongoing uptake of embedded generators, in particular solar PV systems
   c) block loads, historical and future
   d) regional economic and population data
3. metadata concerning CPs, such as:
   a) industrial/ commercial/ residential mix of customers
   b) nature of industry/ commerce.

It is important that the data used in preparing forecasts are accurate, sourced transparently and cannot be said to have been chosen selectively to the forecaster's benefit.

A good approach is to take historical data from publicly available sources such as the Australian Bureau of Statistics (ABS) and the BOM. Of course some data are only available from NSPs. Those data should be taken from the most accurate source possible according to a consistent process, such as giving preference to data from revenue meters over other sources. The sources should be accurately described.

2.2 Normalise

The objective of demand forecasting is not to forecast what *actual* electricity demand will be in any given year.

Rather, the objective is to forecast what demand *would be* under 'normal' conditions. Generally this is taken to mean demand under normalised weather conditions.

This requires that the component of demand that is weather sensitive is 'normalised' out of historical data before forecasts can be produced.

In some cases the key source of variability may be something other than weather. For example, in some areas the key source of variability in demand may be water pumping load, which may be only loosely connected to temperature. In these cases it may be more appropriate to normalise for this other factor. However, the ability to do so depends on identifying that other factor and appropriate data to measure it.[3]

The weather normalisation procedure comprises four steps:

1. prepare the dataset for normalisation

[3] In the specific case of water pumping load, it may be most appropriate to treat it as a block load by deleting it from the historical series and adding projections of it to the forecast.

2. estimate the relationship between temperature and demand at the CP
3. create a distribution of maximum demands for each CP for each year
4. identify 'normal' maximum demand from that distribution.

The procedure is performed separately for each season for which demand forecasts are required, typically winter and summer.

The appropriate dataset to use for normalisation is a subset of the demand data collected at stage 1. Generally, it should:

1. reflect only the season of interest, i.e. summer or winter
2. be one year's data unless conditions were very mild or very extreme
3. be truncated to remove:
   a) 'mild' days
   b) non-working days.

Weather may be daily maximum and minimum temperature, the average of these or another (weather related) variable. This could be examined empirically to identify which is most closely correlated with demand at the CP in question. It may change between CPs and between seasons. For example winter demand may be correlated with temperature at 6:00PM rather than daily (overnight) minimum.

As an example, consider using maximum and minimum demand to compute a linear regression of the following form:

$$MD_d = m * MAXtemp_d + n * MINtemp_d + c + e_d$$

where MD_d is maximum demand observed on day d for all days in the dataset $MAXtemp_d$ and $MINtemp_d$ are daily maximum and minimum temperature respectively *m, n* and *c* are regression parameters, and *e* is an error term.

The coefficients of the regression model and the standard error of the estimate are collected and used with *all* of the available weather data to produce estimates of what daily maximum demand *would have been* in the most recent season under *all* historical weather conditions observed.

The final step of the normalisation process is to identify the 'normal' maximum demand for the season from that distribution of simulated demands.

This is simply a matter of collecting demand at the desired percentile. For example, 50 POE demand is the 50^{th} percentile of the 3000 demands produced in the previous step. 10 POE demand is the 90^{th} percentile etc.

Any POE level can be taken from the distribution of simulated demands.

2.3 Selecting the starting point

The next step is to select the starting point for the forecasts. Conceptually, this is the weather normalised demand in the last year for which actual data are available.

Practically, two options are available and a judgement must be made.

The options are to define the starting point:

- 'off the point' taking the simulated 50 (or 10) POE value for the last available year
- 'off the line' taking the value off a regression line fitted to the weather normalised history (discussed below).

Generally speaking if the two options are close to one another the preferred approach is to take the starting point 'off the line'. However, if the line and point are 'far' from one another the preferred option is to start 'off the point' unless no reason can be found to justify doing so.

It is also important to remember that this choice will be rendered largely obsolete by the reconciliation process.

2.4 Select the initial growth rate

Growth rates are chosen based on the regression developed in choosing the starting point, though again some judgement is required.

The first step is to use the growth rates implied by a regression of historical weather normalised maximum demand and a time trend.

The second step is to sense check this projection with local area experts.

In these cases it may be necessary to modify the growth rate suggested by the regressions or to substitute a growth rate selected manually, for example by substituting the growth rate from a nearby CP or from the system forecast. The decision to do this, and the reasons the particular changes were made, should be recorded.

2.5 Baseline forecasts

At this stage baseline forecasts can be computed by applying the growth rate to the starting point and adding anticipated block loads and future network transfers.

It may be appropriate at this point to make adjustments to these forecasts to account for policy changes, though this will depend on the nature of the policy and the available data.

2.6 Post model adjustments

The baseline forecasts are now adjusted to account for changes in demand that have not otherwise been accounted for.

This is to ensure that factors that *are* expected in future but *are not* incorporated into the system or baseline forecasts are reflected in the final forecasts.

The most frequent post model adjustment is accounting for known block loads, either increments or decrements, to be made at given connection points.[4]

Perhaps the next most frequent cause of post model adjustment is a change in Government policy, or factors caused by such a change, that impact maximum demand. For example, the uptake of solar photovoltaic (PV) systems driven by the various feed-in tariffs and other policies that have been developed in the last five years would have justified a post model adjustment when preparing forecasts of that time. Another example is increased use of demand side management.

Another possible cause is increasing energy efficiency, though energy efficiency is essentially targeted at average demand. Its relationship with *maximum* demand is complex.

---

[4] Adding these block loads to the CPs where they are expected will improve the *allocation* of growth between CPs but, because it is done before reconciliation, will not allow *total* growth to exceed growth in the system forecast. This is important as applying block loads after reconciliation would lead to double counting because the block loads are 'in' the system growth.

The appropriate way to forecast the impact of a policy change will vary with the particular policy in question. The key issue is to focus on *changes* in maximum demand attributable to that policy.

Some post model adjustments are inherently difficult to prepare and contentious. By their nature they rely on assumptions which are difficult or impossible to verify. They often rely on expectations of future Government policy, which can change. It is imperative that the adjustments made and the assumptions and methodology used to develop them are stated explicitly and available for review. To the maximum extent possible the forecast impacts should satisfy the same principles as the forecasts themselves.

2.7 Reconciliation to system forecast

The final stage in the demand forecasting process is to reconcile the CP forecasts to an independently prepared system forecast. In this context 'system' refers to forecasts of maximum demand at the NEM region level.[5]

The purpose of doing this is to 'import' the likely impact of drivers of electricity demand into the CP forecasts. Some drivers, such as economic activity are not well understood or forecast at the CP level and lend themselves more to incorporation into a higher level forecast. Depending on their nature, some policy changes will be more appropriately modelled at the system level and 'imported' into the CP forecasts.

In doing this it is necessary to account for system losses and for diversity.[6]

5 Note that this is a notional forecast. There is no single physical point where this demand will be observed.

6 Diversity is the possibility that maximum demand at a given CP will occur at a different time than system maximum demand.

3 Verification of methodology

ACIL Allen's methodology for forecasting maximum demand at the connection point (terminal station) level consists of seven steps as described in chapter 2 of this report and in more detail in ACIL Allen's methodology report.

3.1 Verification of methodology

In summary, ACIL Allen is satisfied that AEMO applied the methodology outlined in the June 2013 report appropriately. Certain minor modifications were made in consultation with ACIL Allen in response to advice from Frontier Economics. Those modifications were consistent with the methodology and are reasonable.

3.2 Implementation summary

Table 1 provides a step by step comparison of ACIL Allen's methodology and AEMO's implementation.

Table 1 **Summary evaluation of implementation**

Methodology step	AEMO implementation	ACIL Allen conclusion
Prepare data (demand data)	AEMO obtained demand data from its internal databases at the CP level.	ACIL Allen has not directly reviewed the demand data. The source of demand data was appropriate.
Prepare data (weather data)	AEMO relied on temperature data collected from the Bureau of Meteorology, which is the recommended source.	The source of temperature data was appropriate and consistent with ACIL Allen's methodology as was AEMO's method for selecting the appropriate weather station, though in some cases required judgement.
Prepare data (other data)	ACIL Allen understands that AEMO had numerous discussions with the relevant DNSPs to clarify and explain various aspects of the forecasts and to test AEMO's preliminary forecasts with local area experts.	The source and extent of metadata was appropriate and consistent with ACIL Allen's methodology.
Normalise	AEMO committed considerable effort to weather normalisation, including by considering a wide range of candidate weather stations. AEMO applied the weather normalisation procedure as described in ACIL Allen's methodology report. However AEMO identified the issues discussed in section 4.3 for further testing in future. It appears that demand at some CPs was not temperature sensitive or responds to variables other than temperature. In these cases a 'constant only' model was used. AEMO was not able to account for this variability directly in all cases, which is neither unreasonable nor surprising.	The approach to normalisation was appropriate and consistent with ACIL Allen's methodology.
Select starting point	AEMO used a combination of statistical tests and judgement to determine whether forecasts should be taken from 'the point' or 'the line'.	AEMO's approach to selecting the starting point is consistent with ACIL Allen's methodology in principle However, as discussed in section 4.1, ACIL Allen considers the usefulness of statistical testing to determine whether to take forecasts off 'the point' or 'the line' to be limited.
Select growth rate	Growth rates were typically taken from a time trend in normalised historical demand. The mix of demand in the area supplied by each CP was considered based on metadata obtained from the relevant DNSP. In regions where demand is mainly residential the growth rate was compared with population growth rates for regions that are mainly residential using data from Government sources.	The selection of growth rates was appropriate and consistent with ACIL Allen's methodology.
Baseline forecast	Baseline forecasts were computed by applying growth rates to starting points.	The preparation of baseline forecasts was appropriate and consistent with ACIL Allen's methodology.
Post model adjustments	Post model adjustments were made for the impact of solar Photovoltaic (PV) systems and Energy Efficiency.	ACIL Allen did not review the post model adjustments.

Methodology step	AEMO implementation	ACIL Allen conclusion
Reconcile to system forecast	AEMO computed coincidence factors as described in ACIL Allen's methodology and then 'trimmed' the sum of coincident maximum demand so that it was equal to maximum demand as forecast at the system level using an independent model. It then used the same coincidence factors to compute non-coincident maximum demand at each connection point.	The approach for reconciling the forecasts to the system level was appropriate and consistent with ACIL Allen's methodology. ACIL Allen has not reviewed the system forecast used for reconciliation.

4 Frontier Economics red flag reviews

This chapter provides an overview of the advice ACIL Allen provided AEMO in response to Frontier Economics' two red flag reviews. There were four issues, which are discussed in turn below:

1. methods for formalising the selection of:
   a) starting point (section 4.1)
   b) growth rate (section 4.2
2. several issues relating to the detail of the weather normalisation procedure (section 4.3).
3. the appropriate treatment of CPs where demand does not appear to vary with temperature and whether and how to produce forecasts at different POE levels.

4.1 Statistical testing to assist in choosing the starting point

The issue was whether forecast growth rates should be applied to a regression line fitted to:

1. (weather normalised) historical demand data (the line)
2. the last weather normalised value (the point).

Frontier recommended that forecasts should be made 'off the line' if two conditions are met:

1. the historical data fit a linear trend model
2. the last weather normalised value is not an outlier from that model.

Otherwise, Frontier recommends that the forecast should be applied 'off the point'.

In other words, Frontier recommended determining whether the historical data exhibit a linear trend and, if so, whether the last value is an outlier.

Frontier recommended two statistical tests for this:

1) add a quadratic term to the regression model. If the coefficient on the quadratic term is statistically significant, reject the linear trend model (and therefore forecast 'off the point')
2) include a dummy variable in the last year for which data are available. If the coefficient on that dummy is significant, conclude that the last point is sufficiently far away from the regression line that it is an outlier (and therefore forecast 'off the point').

The underlying idea of Frontier's suggestions is consistent with the base methodology as described in the following passage from ACIL Allen's methodology report, which Frontier quoted in the red flag document.

> If no valid reason to support starting 'off the point' can be found, it is reasonable to conclude that the difference between the 'point' and the 'line' is due to randomness in the data. In this case, starting 'off the line' is preferred over assuming that the same random outcome will be repeated in every forecast year.

That is, it is generally preferable to forecast off the line where possible.

However, it must be remembered that, by definition, demand at the level shown by 'the point' was observed recently. As Frontier noted, this might be due to 'idiosyncratic factors that will not recur in the future'. If it is, forecasting off the line is preferable.

However, the difference may also be due to a structural change that has occurred recently and that will recur in the future, in which case forecasting off the point may be preferable.

In the base methodology ACIL Allen emphasises the importance of drawing on the judgement and experience of people with direct knowledge of the areas in question. Caution should be used in relying on an automated statistical approach that may overlook local detail. However, by the same token, it is important that the forecasts do not become too subjective.

As we said in the base methodology report (p. xviii):

> Starting the forecast 'off the line' may force a step change in demand between the last actual year and the first forecast year. Whether this is appropriate depends on the reason for that step change.

In relation to the suggested use of the quadratic term, this would distinguish between a linear and non-linear relationship if sufficient data were available.

However, in this particular application there will typically be very few data points (perhaps 10 at most). With a data set this small we would be concerned that the quadratic approach would lead to over fitting. That is, it would reject linear trend models where they are more appropriate by fitting a curve to the particular sample that has been observed.

In relation to selecting the starting point, in the base methodology report we said that forecasts should start 'off the line' if the difference between the line and the point can be attributed to randomness. Frontier's proposal is a statistical way of evaluating whether the difference to which we refer can reasonably be attributed to randomness or whether there is likely to be another cause. Therefore, it is consistent with the base methodology.

However, we are concerned that the use of the dummy variable is too much of a 'blunt club'. This approach will put all of the variability in the last year onto the dummy variable. The coefficient on the time trend will be the same as if the last year was excluded from the model. There is a risk that this will overcorrect for the problem.

4.2 Selecting the growth rate

The issue is the growth rate at which forecast should proceed, whether that is from 'the line' or 'the point'.

Frontier recommends that if the linear trend model is accepted the growth rate in that model should be continued. Otherwise, the population growth rate should be applied.

The recommendations are reasonable.

Further, in areas where a large part of the electrical load is non-residential the population growth rate may be less suitable. In this case it may be more appropriate to fit a regression model to historical population and to use population projections to drive the forecast. It may also be worth including other variables in the base model rather than reverting directly to population growth, though this would add to the complexity of the forecasting task.

More broadly, though, we note that there is no discussion in Frontier's red flag document of reconciling the bottom up forecasts discussed here to an independently prepared system forecast, which is a fundamental part of the base methodology but had not been done at the time the red flag reviews were conducted. In the base methodology, the system level forecast overrides the individual spatial forecasts at the last step. Forecasting the growth rates individually is really about forecasting the proportion of total load that will be experienced at each CP.

We consider Frontier's recommendations to be reasonable, but emphasise that the forecasts should be reconciled to an independent system forecast, which reduces the significance of the growth rates applied to the spatial forecasts in the first place.

4.3 Recommendations regarding weather normalisation

Frontier Economics proposed amendments to the weather normalisation procedure in ACIL Allen's methodology. In ACIL Allen's methodology, ACIL Allen recommended that forecasts of demand in each year are calculated using:

1. a constant term, which provides information on the non-temperature sensitive level of demand
2. a slope, which describes the relationship between temperature and temperature sensitive demand.

In a given year 't' , this function could be written as shown in equation (1):

$$Y_t = \alpha_0 + \alpha_1 Temp_t + e_t \tag{1}$$

where Y_t is the demand in time period t, $Temp_t$ is the temperature in period t, and e_t is the idiosyncratic error term. α_0 and α_1 are coefficients estimated through regression techniques.

ACIL Allen recommended estimating this equation for each year separately.

In its red flag reviews, Frontier noted that this would lead to volatility in the slopes. In turn this could lead to misleading linear trends, affecting the forecasts.

There appears not to be a fundamental difference between Frontier and ACIL Allen in relation to this issue. Both consultants noted that temperature sensitivity changes over time. The real question is how much to constrain this in the modelling.

Frontier's first suggestion was to pool all of the available data and estimate equation (2):[7][8]

$$Y_t = \delta_1 I_1 + \delta_2 I_2 + \delta_3 I_3 + \delta_4 I_4 + \delta_5 I_5 + \delta_6 Temp_t + e_t \tag{2}$$

where Y_t is the demand in time period t, $Temp_t$ is the temperature in period t, the I_n's are dummy variables set to 1 in year 'n' and 0 otherwise and e_t is the idiosyncratic error term. δ_1, δ_2, δ_3, δ_4, δ_5, and δ_6 are coefficients estimated through regression techniques.

This amounts to assuming that the slope in the relationship between temperature and demand is the same each year.

The pooled approach has the advantage that it increases the sample size, and hence may lead to increased statistical significance of particular parameters in some instances. In cases where there are insufficient data to produce a reliable estimate of the slope ACIL Allen recommended this approach.

7 It should be noted that this understanding is based on discussion at a meeting rather than any written document from Frontier Economics.

8 This example is based on an assumption that there are five years of data. Frontier Economics' recommendation was more general, pooling all the available data.

However, this approach was not recommended in ACIL Allen's methodology unless the data are insufficient to avoid it. The reason was ACIL Allen's expectation that the slope is the same every year will usually not be an appropriate assumption. The reason is that, in heavily residential areas load growth is generally associated with more customers and, therefore, more air-conditioners. This means that temperature sensitivity, when measured in terms of MW increase per degree, will increase over time.

This assumption could be relaxed with pooled data by estimating equation (3):[9]

$$\begin{aligned} Y_t = \beta_1 I_1 + \beta_2 I_2 + \beta_3 I_3 + \beta_4 I_4 + \beta_5 I_5 + \ & \beta_6 Temp_t I_1 \\ & + \beta_7 Temp_t I_2 + \beta_8 Temp_t I_3 + \beta_9 Temp_t I_4 \\ & + \beta_{10} Temp_t I_5 + e_t \end{aligned} \quad (3)$$

where I_1 to I_5 and β_1 to β_{10} are as defined for equation (2).

This specification would yield the same coefficient estimates as estimating the equation year by year. However, it also assumes that the standard error of the regression is the same each year, which is significant for the simulation step in the process.

In this model the assumption that the temperature coefficients are the same over time could be tested by determining whether the temperature coefficients estimated for each year are statistically distinguishable from one another (i.e. β_6 to β_{10}). For example, given that Frontier's model is linearly constrained within a pooled version of ACIL's model, an F-test could be used to assess whether Frontier's assumption of non-varying temperature sensitivity is appropriate.

In the second red flag review Frontier recommended a middle ground where the slopes would be calculated as a rolling average over time or by partial pooling.

ACIL Allen advised AEMO that this is essentially a trade off. AEMO needs to trade off the volatility of coefficients in annual models with the loss of resolution in pooled models.

Without reviewing individual data it is difficult to make general statements, though the second approach of pooling data three years at a time seems reasonable. If this approach was taken it would be important to include a dummy variable for the three years to allow for underlying growth.

9 As above this assumes five years of data.