Temperature correction of energy statistics

Prepared for the Department of Energy and Climate Change

January 2011

Sumit Rahman

Methodology Advisory Service,

Office for National Statistics

Summary and Recommendations

Summary

1. Energy consumption statistics are a key output of the Department of Energy and Climate Change (DECC). The monthly statistics for each fuel type are currently corrected for temperature before being seasonally adjusted. The purpose of temperature correction is to help users better understand underlying trends in energy consumption, which can be affected by fluctuations in temperature, by producing a monthly series which can be interpreted as the consumption that would have happened if temperatures had been at their average for the month. Then changes in the series from one year to the next are not due to changes in temperature.

2. When a month is cooler than its long term average, the temperature correction adjusts its energy consumption figure downwards to account for the fact that more energy was required to keep rooms warm. Conversely when a month is warmer than its average, consumption is adjusted upwards. For each degree Celsius that the monthly temperature deviates from its average, a temperature correction factor is applied which adjusts the consumption figure by a certain percentage (so the percentage is proportional to the temperature deviation).

3. Temperature correction is effective for the natural gas and coal series, in that some variation in these series can be explained by temperature fluctuations around the monthly average, and the seasonal adjustment is improved.

4. The temperature correction for coal can be improved by using a new set of 12 temperature correction factors, one for each month. The correction factor is multiplied by the deviation in a month's average temperature from its long term average and the result tells you by what percentage the consumption figure should be adjusted.

5. These new factors were found by adding 12 variables as permanent prior adjustments to the time series in the X12 seasonal adjustment software package and running a regression on them. A set of factors was found for the gas series, but this set was not as effective as the current temperature correction, which is performed by the gas data supplier for DECC. It appears that the data supplier might stop performing this temperature correction for DECC so it is worth trying to produce a set of factors for gas anyway.

6. Effective temperature correction factors could not be found for the petroleum and nuclear energy series, and the petroleum series was not even judged to be seasonal.

7. The new factors that were estimated by the software package were adjusted so as to form a series of factors that vary smoothly throughout the year, with two major peaks (representing months where consumption is most sensitive to changes in average temperature) in Spring and Autumn.

8. New models were found making use of the concept of 'heating degree days'. These are the sum of temperature differences (between actual temperature and some fixed 'baseline temperature') over time, and improve on the sole use of fluctuations in average monthly temperatures by only taking account of when the temperature is below the baseline and not when it is above. There is a slightly different method used by Eurostat for calculating heating degree days, compared to how they have been calculated historically in the UK.

Models were fitted making use of the Eurostat concept and the conventional degree day (using two different baselines).

9. The model using the Eurostat concept was found to give the best results, and the models using the conventional degree days the worst, with the new sets of temperature correction factors in-between. A number of approximations were used to produce a suitable data set for the regressions on the conventional degree day measure, and this might account for the poor performance of these factors.

Recommendations

10. DECC should consider moving to using the Eurostat heating degree day for its temperature correction of monthly natural gas and coal consumption series. The correction factors should be applied by using the exponential formula (the proportional percentage changes are an approximation to this).

11. If DECC would prefer to stick with using monthly average temperatures then the new set of proposed temperature correction factors for natural gas and coal should be used.

12. The petroleum series does not need to be temperature corrected or seasonally adjusted (at least, not from 1995).

13. If DECC wishes to use the conventional heating degree days measure then the traditional baseline temperature of 15.5° should not be used; a baseline temperature of 18° is more appropriate.

Contents

1 Introduction

1.1 The Department of Energy and Climate Change (DECC) publishes monthly estimates of energy consumption by fuel types. As well as 'raw' figures the Department publishes adjusted series which are temperature corrected and seasonally adjusted. The latter are the key estimates for users interested in energy efficiency at a macro level, and the Department judges that increasing attention is being given to year on year changes in the temperature adjusted data.

1.2 DECC also produces quarterly estimates of final energy consumption by fuel type and broad sector (e.g. 'transport', 'domestic', 'iron and steel'). These are not currently adjusted for temperature or seasonal effects.

1.3 The current method of temperature correction is to compare the temperature in a given month (using Met Office data) with a long term average for that month, using degrees Celsius, to arrive at a temperature deviation. For each fuel type and month there is a given factor which is multiplied by the deviation to give the percentage increase or decrease to apply to the consumption estimate. For example, in April 2009 the observed temperature is calculated from the Met Office data to be 1.8° higher than the long term average for April[1]. The factor for coal in April is 2.1% so that the 'raw' April figure for coal needs to be adjusted by 1.8 × 2.1% = 3.7% – meaning the consumption figure is increased by 3.7 per cent. (The method assumes that energy consumption increases in colder weather and reduces in warmer weather, so that the temperature correction adjusts the consumption figure up in this example.)

1.4 After the temperature correction has taken place the series can undergo the seasonal adjustment process using the X12 software package in the usual way as recommended for the Government Statistical Service by the Office for National Statistics.

1.5 The temperature correction takes place in this way for coal and petroleum. For natural gas the temperature correction is performed in a different way by the data supplier, and DECC then seasonally adjusts it. However the method used by the data supplier is about to change and DECC judges that the new methodology will be unsuitable for its purposes.

1.6 Coal, petroleum and natural gas account for about 90% of total consumption. The nuclear series is not temperature corrected because nuclear production takes place without responding directly to demand. For similar reasons the wind and hydroelectric series is not temperature corrected. Both are seasonally adjusted however. Finally the net imports of electricity is not adjusted at all.

1.7 This report is focussed on the time period 1995 – 2008. Since 2009 all the monthly series are collected and reported on a calendar month basis. Before 2009 some of the series, such as coal, were collected on a statistical month basis ('4, 4, 5': i.e. two months of four weeks followed by a single month of five weeks, with the occasional 'leap week' to bring the statistical months back in line with the calendar).

1.1

---

[1] Degrees are degrees Celsius throughout this report.

2 Does correcting for temperature add value for the user?

Summary of main points

- *For coal and natural gas, the current methods of temperature correction do improve the seasonal adjustment and so temperature correction should continue*
- *The petroleum series is not seasonal and the current method of temperature correction does not lead it to become seasonal*

2.1 Intuitively there seems to be a reasonable case for performing some kind of temperature correction to the consumption series, given that a large proportion of energy consumption is used for domestic heating. One would expect that in cold months more energy is required to warm up houses than in warm months, but this is accounted for by seasonal adjustment. The temperature correction that is applied is to adjust for changes in the month's temperature in comparison to that month's long term average. So an unusually warm January should be adjusted up, even if it is still a low temperature in comparison with the whole year.

2.2 The temperature corrected series then is interpretable as what the consumption would have been had the temperature in each month been the long term average for each month.

2.3 Monthly temperature is calculated as the mean monthly temperature for 17 weather stations across Great Britain (with four stations weighted twice as much as the other stations). For a weather station the mean monthly temperature is the mean daily temperature, where the daily temperature is the mean of the minimum and maximum temperature for the day. DECC receives the daily minimum and maximum temperatures from the Met Office every month and calculates the monthly temperature using spreadsheets.

2.4 The long term average is calculated for each month and relates to the period 1971 – 2000. For the purposes of this report we do not consider other 'bases', although DECC will eventually move on to a 1981 – 2010 basis. The use of 30 year periods for long term averages is widespread and this is unlikely to be criticised by users.

2.5 In order to decide on the value of a temperature correction method, we can consider the impact it has on the seasonal adjustment of the series. The canonical decomposition of a time series, on which the X12 method depends, expresses it as the product of a 'trend' series, a seasonal factor and an 'irregular' factor[2]. We can judge the effectiveness of a temperature adjustment by considering the impact on the irregular. If the irregular is noticeably improved (i.e. it is 'smaller' in some sense, such as contributing less to the variation in the series) then the temperature adjustment is worth doing. We can also

1.1

---

[2] The seasonal factor is removed to produce the seasonally adjusted series.

judge the effectiveness by looking at the standard diagnostic statistics produced by X12 (the M statistics) and see if they are affected by the temperature correction.

GAS

2.6 We consider first the gas series. This series is clearly seasonal:

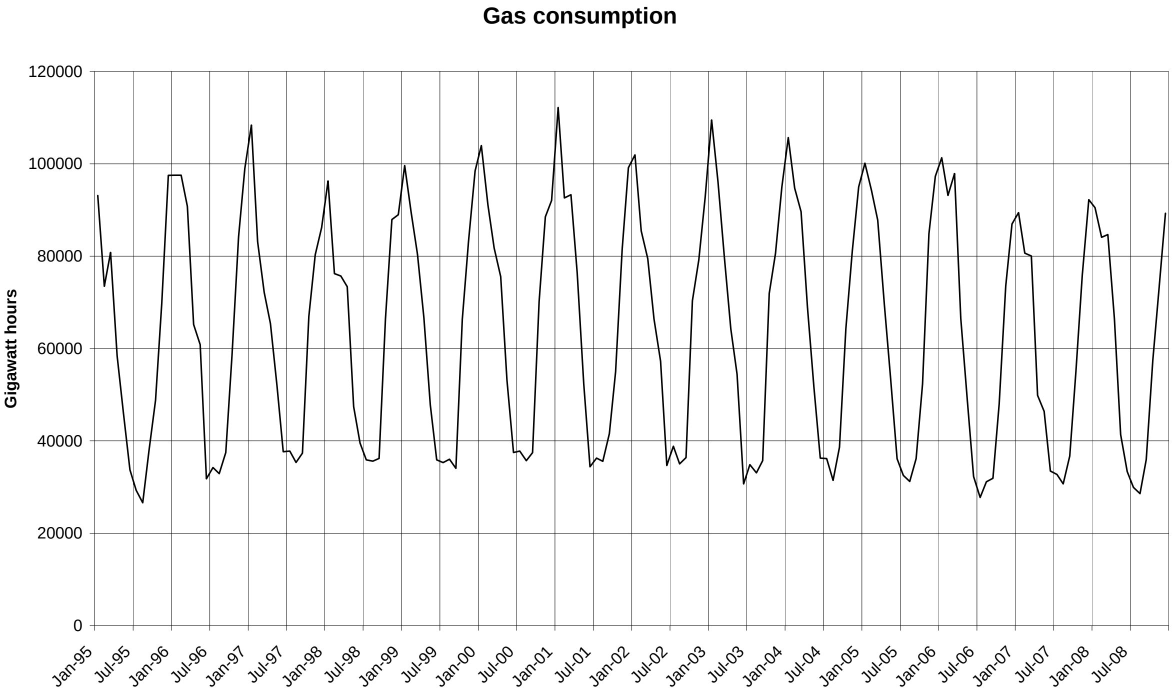

Figure 1 – Gas consumption (no temperature correction)

2.7 It is worth commenting that the temperature correction does not take place on the whole series – there are minor changes which are taken out before these adjustments and then added back in. Similar steps occur in the coal series. In this report we focus on the series in the state before temperature and seasonal adjustment, so the ‘raw’ series may not exactly correspond to the published ‘raw’ series (but the differences are minor).

2.8 We can run the seasonal adjustment program on this series without any temperature correction and then with the correction, and examine the irregulars: have they become smaller?

2.9 One way to determine this is to look at the percentage annual changes in the series before seasonal adjustment takes place, and see how these are affected by the temperature correction. Looking at annual changes (i.e. comparing each month’s figure with the same month the year before) largely removes the impact of seasonal effects (assuming no sharp intervening seasonal breaks), and the F tables in the standard X12 output summarises these changes.

2.10 For the gas series without any temperature correction, the average 12 month change is 7.4% (taking the absolute value of the percentage changes, so that positive and negative changes do not offset each other). The equivalent figure for the temperature corrected series is just 3.8%. This suggests that the temperature correction helps smooth the series (at 12 month intervals) and removes some 'noise'.

2.11 Another measure is described by X12 as the 'relative contributions to the variance of the percent change in the components of the original series' – this allows us to see, after the seasonal adjustment has taken place, how much of the variation in the 12 month percentage changes is attributable to the changes in the trend of the series and how much is unexplained (attributed to the 'irregular' series).

2.12 Temperature correction for the gas series reduces the unexplained variation in the 12 month growth figures from 55% of the total variation to just 22%. The variation explained by the trend series increases from 38% to 76%.

2.13 So we can say that temperature correction as currently used by DECC for the gas series reduces the variation in the annual growth figures and attributes far less of it to the irregular series. This is good evidence to justify performing the temperature correction.

2.14 For such a strongly seasonal series it is hard to tell if this reduction in the irregular is much of an impact – is the temperature correction worthwhile? We can look at the M statistics to help us. We focus on M1, M7 and Q. Values below 1 indicate good seasonal adjustments. Even without the temperature correction the seasonal adjustment is very good. However, after the temperature correction performed by the data supplier, each of these diagnostic statistics improves.

	Uncorrected	Corrected
M1	0.049	0.007
M7	0.112	0.044
Q	0.66	0.35

2.15 These statistics show that, although the uncorrected series leads to a seasonal adjustment that would be considered satisfactory in most contexts, the temperature correction still improves it.

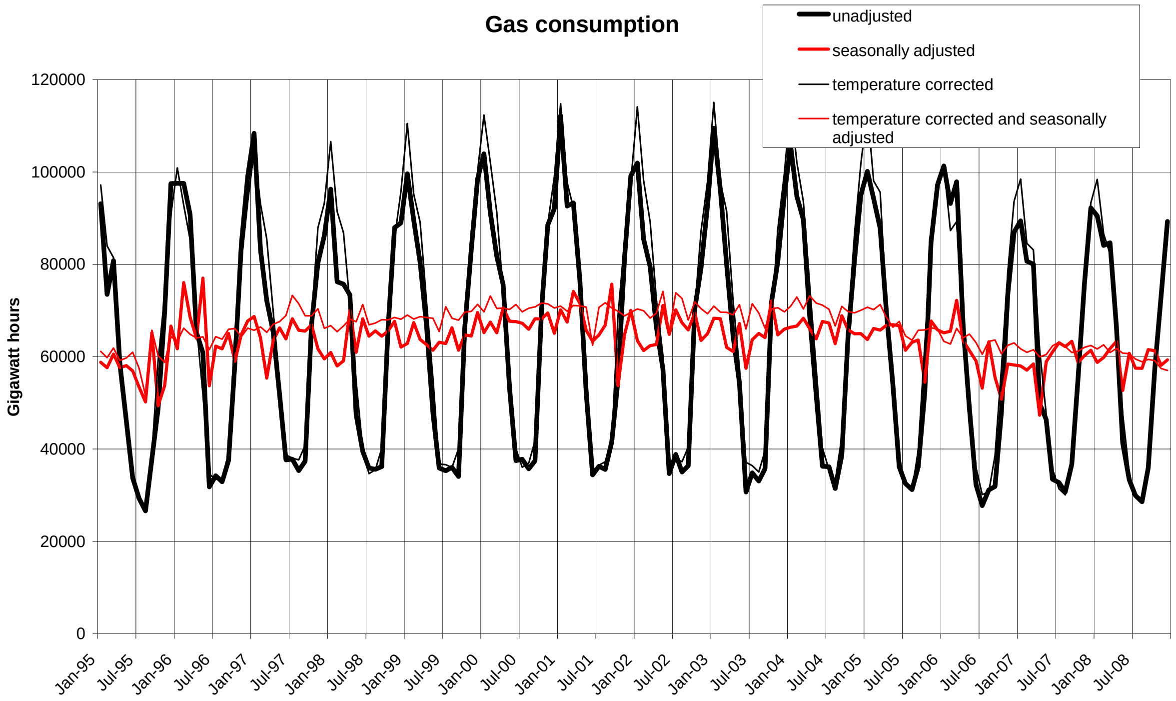

Figure 2 – Gas consumption, with the adjustments as currently performed

2.16 Figure 2 shows that the final adjusted series is probably slightly smoother than the seasonally adjusted, non-temperature corrected series. It seems fair to conclude that the temperature correction is 'adding value' for users.

COAL

2.17 Next we consider the coal series. This does not exhibit the seasonality as clearly as the gas series did:

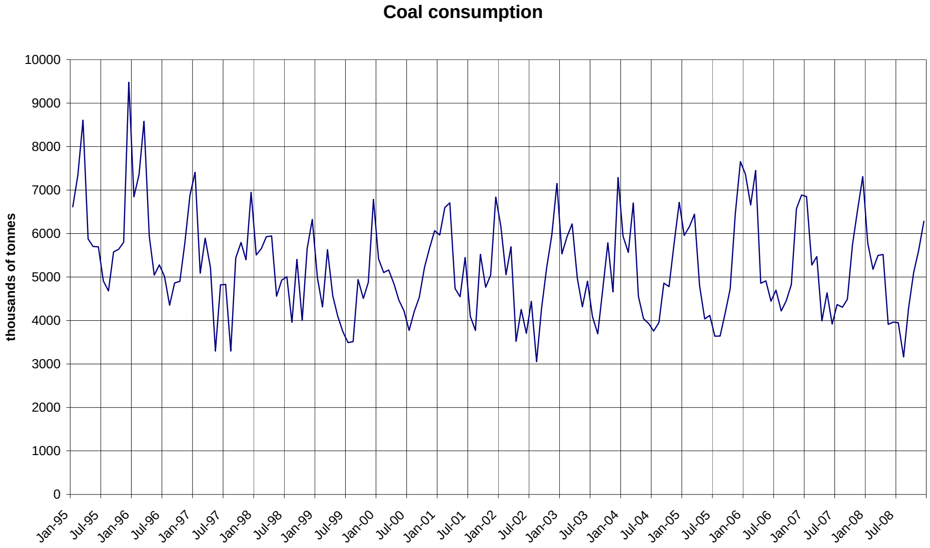

Figure 3 – coal consumption (no temperature correction)

2.18 We can see a pattern of peaks around January and troughs around July, but the irregular component is clearly playing more of a role here than it did in the natural gas series.

2.19 Looking at the F tables we see that temperature correction as currently employed by DECC has a positive impact on the 12 month changes. The mean 12 month change (ignoring sign) reduces from 12.4% to 11.8%. The variation in the 12 month change explained by the irregular reduces from 49% of the total variation to 43% (and the variation explained by the trend series increases from 51% to 56%. The M statistics also show an improvement with the temperature correction – in particular the M1 statistic falls to below 1.

	Uncorrected	Corrected
M1	1.157	0.837
M7	0.335	0.329
Q	0.95	0.91

2.20 It is not clear from Figure 4 that the temperature correction has helped the seasonal adjustment, but the statistics in paragraph 2.19 are enough to satisfy us that the temperature correction is worth doing. Of course we will look at improving the temperature correction method and see if the resulting seasonal adjustment also improves.

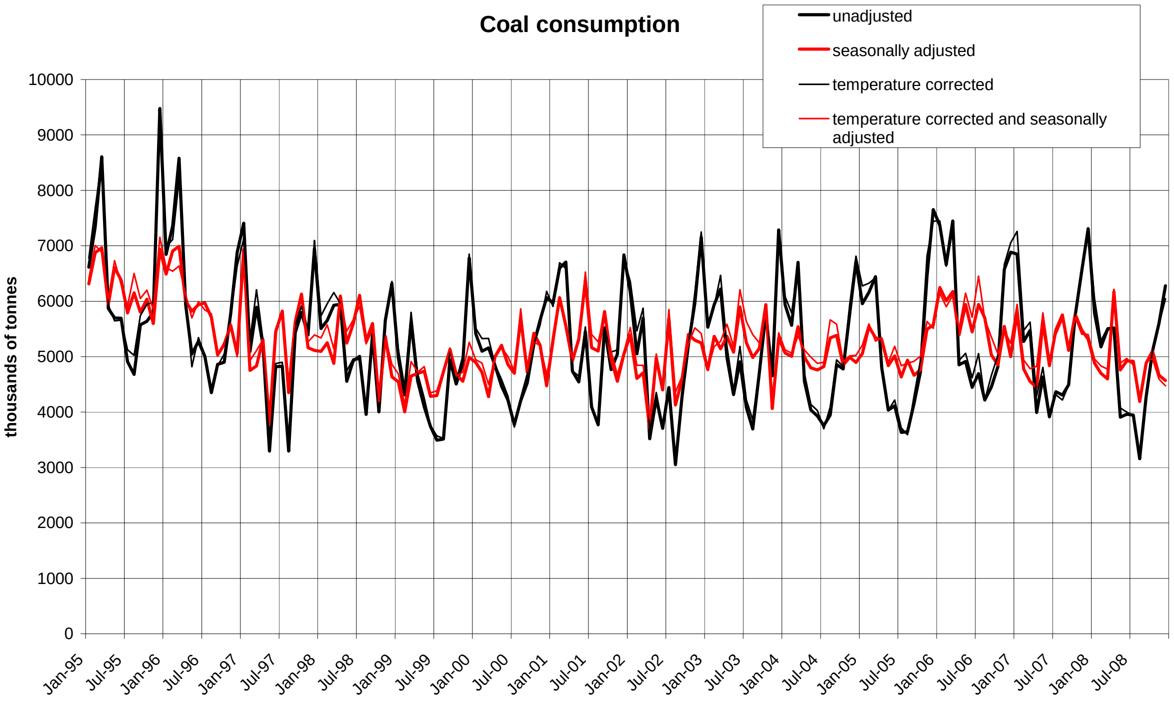

Figure 4 - coal consumption, with the adjustments as currently performed

PETROL AND NUCLEAR

2.21 Next we examine the petrol series. The graph of the raw series (Figure 5) does not suggest there is a seasonal pattern.

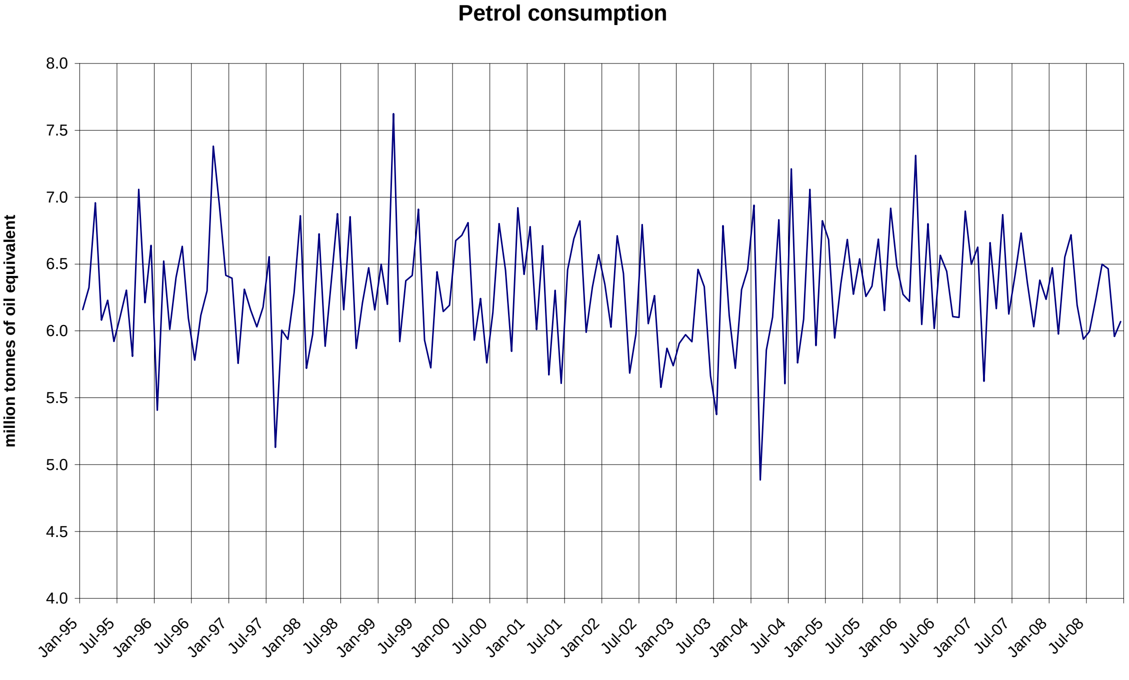

Figure 5 – Petrol consumption (no temperature corrections)

2.22 The lack of seasonality is confirmed when running the raw series through X12, as it fails the combined test for seasonality. There is some evidence for seasonality in that it passes two of the three separate tests, but the final M1 statistic for example is 3, which indicates an unacceptable seasonal adjustment.

2.23 The temperature correction as currently performed does not improve the seasonality – the temperature corrected series does not pass any of the tests for seasonality and the M statistics do not improve. The average 12 month change actually increases slightly after the temperature correction, as does the relative contribution of the irregular to this variation. (This latter increases from 86% to 87%.)

2.24 Figure 6, and the statistics in paragraphs 2.22-23 suggest that there is no real benefit to the adjustment to the raw series as currently undertaken by DECC.

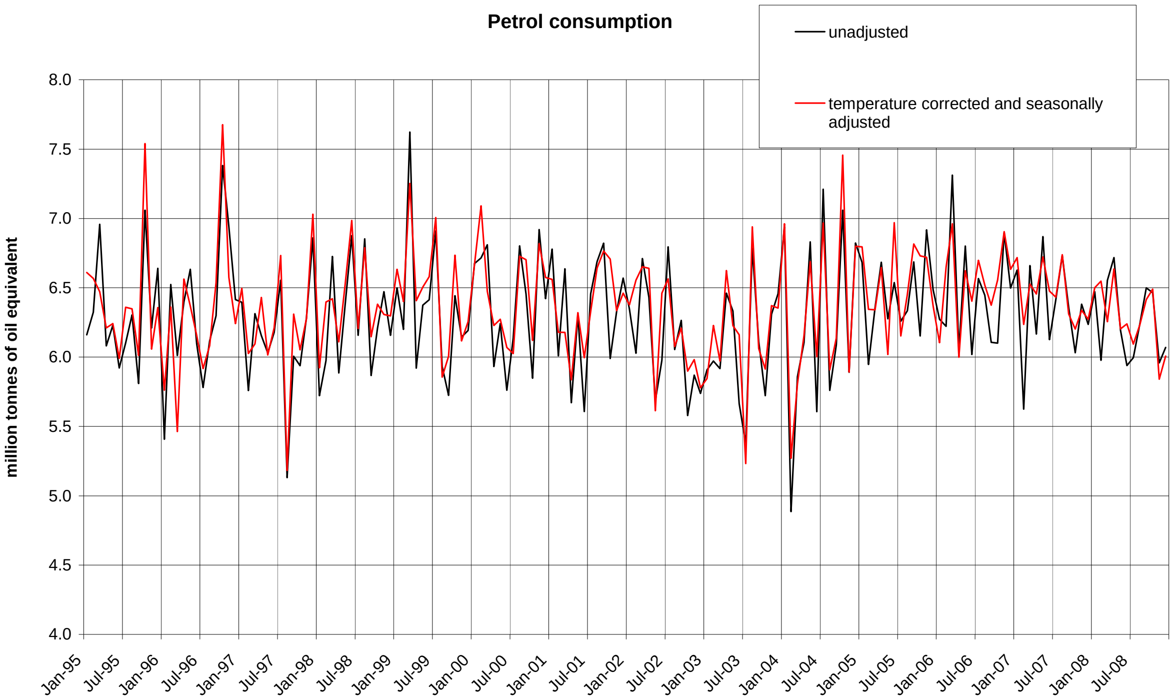

Figure 6 - petrol consumption unadjusted and adjusted

2.25 The nuclear series is not corrected for temperature as it is part of the baseload of energy production (so does not respond to short term changes to demand in the domestic energy market). The series seems to be seasonal – there appear to be peaks in the final month of each quarter, and the combined test for seasonality performed by X12 concludes that the series is seasonal. The seasonally adjusted series appears to be smoother, although the M statistics are not good (the overall Q value is 0.87, indicating a reasonably good adjustment, but the M1 statistic is 2.252, much higher than we would like). On balance, it seems that the seasonal adjustment of the nuclear series should continue to be performed by DECC.

3 Adjusting the temperature correction factors

Summary of main points

- *We can amend the monthly temperature correction factors for coal to arrive at an improved seasonal adjustment*
- *Using monthly correction factors for gas is not as effective as the current method in terms of the seasonal adjustment, but if the new temperature corrected series from the data supplier are not suitable for the Department then the proposed monthly factors still result in an excellent seasonal adjustment*
- *We cannot improve the seasonality of the petrol series by using monthly correction factors and so recommend that the petrol series not be temperature corrected or seasonally adjusted*

3.1 The temperature correction factor for coal in April described in the introduction to this report (see paragraph 1.3) was 2.1% – for each degree of deviation in April from the long term average April temperature the series is adjusted by 2.1%. In fact for coal the correction factor is 2.1% for every month. We can look to vary these factors by month and see if we can improve the seasonal adjustment further.

3.2 For gas, the method used for temperature correction is effectively using a daily correction factor. However the Department has discovered that the methodology is changing so that the basis for the long term average is no longer a 30 year period in the past but one centred on the present year and making use of forecasts. There is then a need to produce a new temperature correction method for the Department to apply by itself to the unadjusted series.

3.3 The correction factors for petrol vary by month: for nine months it is 1.8% and for June, July and August it is 0.7%. Since this correction does not lead to a series that can be usefully seasonally adjusted we shall look to change these factors as well.

3.4 We will also look at the nuclear series, which is currently not corrected for temperature, to see if a temperature correction helps the seasonal adjustment.

3.5 The approach we take is to treat these correction factors as 'permanent prior adjustments', in the language of X12. The regARIMA part of the X12 package can estimate these by regressing 12 variables (one for each month, with the value of the variable equal to the temperature deviation in the relevant month, or zero if it is not the relevant month). The regression will produce estimates for the coefficients of each month, and these will give us the new correction factors.

COAL

3.6 For coal the regression was run and the model was found to be highly significant (χ^2 = 38.6 with $df = 12$, $p < 0.005$). The estimates for each month's correction factor are in the following table:

Variable	Estimate	Standard error	t-value
Jan	-0.044	0.022	-1.94
Feb	-0.045	0.019	-2.34
Mar	-0.063	0.019	-3.36
Apr	-0.091	0.028	-3.26
May	0.002	0.024	0.10
Jun	0.036	0.040	0.91
Jul	-0.012	0.021	-0.55
Aug	-0.026	0.023	-1.16
Sep	-0.018	0.034	-0.53
Oct	-0.057	0.024	-2.38
Nov	-0.021	0.028	-0.74
Dec	-0.009	0.026	-0.34

3.7 We can interpret these coefficients as follows: because X12 treats the series using a multiplicative decomposition, the regression is on $\log Y_t$. If we write $\log Y_t = B_t X_t + Z_t$, where Y_t is the unadjusted consumption series, B_t the coefficients estimated above (repeating with period 12), X_t the temperature deviations and Z_t the prior adjusted series, then the temperature corrected series is given by $Y_t \times \exp(-B_t X_t)$. So if we observed a temperature deviation in January of 0.5°, the temperature correction would be $\exp(0.044 \times 0.5) = 1.022$, i.e. an increase of 2.2%. This fits with our expectation at the end of paragraph 2.1, that a warmer than expected month leads to an increase to the unadjusted data.

3.8 Since $-B_t X_t$ will be small (deviations are typically between -1 and 2, and the B_t estimates are all less than 0.1), we can use the approximation $\exp(x) = 1 + x$. This means that the temperature corrected figure is $Y_t \times (1 - B_t X_t)$ and we can interpret these B coefficients as temperature correction factors as we do in the current DECC methodology, except for the change in sign. Note that use of this approximation in the above example leads to the same percentage change to 1 decimal place ($0.044 \times 0.5 = 0.022$ = 2.2%.) For this report temperature corrections will be calculated using the full formula and not using the approximation (except in paragraphs 3.21-26). The approximation nevertheless allows for this easy to explain interpretation to be used in describing the temperature correction method.

3.9 The estimates as given in paragraph 3.6 need more consideration. The existence of positive coefficients in May and June is problematic, as this implies the temperature correction is applied in the wrong direction. But note that these coefficients are not significant. Running the model fitting process again with these coefficients set to zero leads to a model with no positive coefficients, but the non-significant coefficients remain so. A model can also be fitted which uses just the five significant months and when this is done the five months (January to April, and October) remain significant.

3.10 The M statistics below allow us to compare the seasonal adjustments when no corrections are made for temperature; when the current correction is made; using the model with May and June set to zero; and the reduced model described in the previous paragraph. The graph compares the seasonally adjusted series resulting from the reduced model with the current method's series.

	Uncorrected	Current temperature correction	Using the model with 10 months	Using the model with 5 months
M1	1.157	0.837	0.717	0.652
M7	0.335	0.329	0.293	0.280
Q	0.95	0.91	0.91	0.92

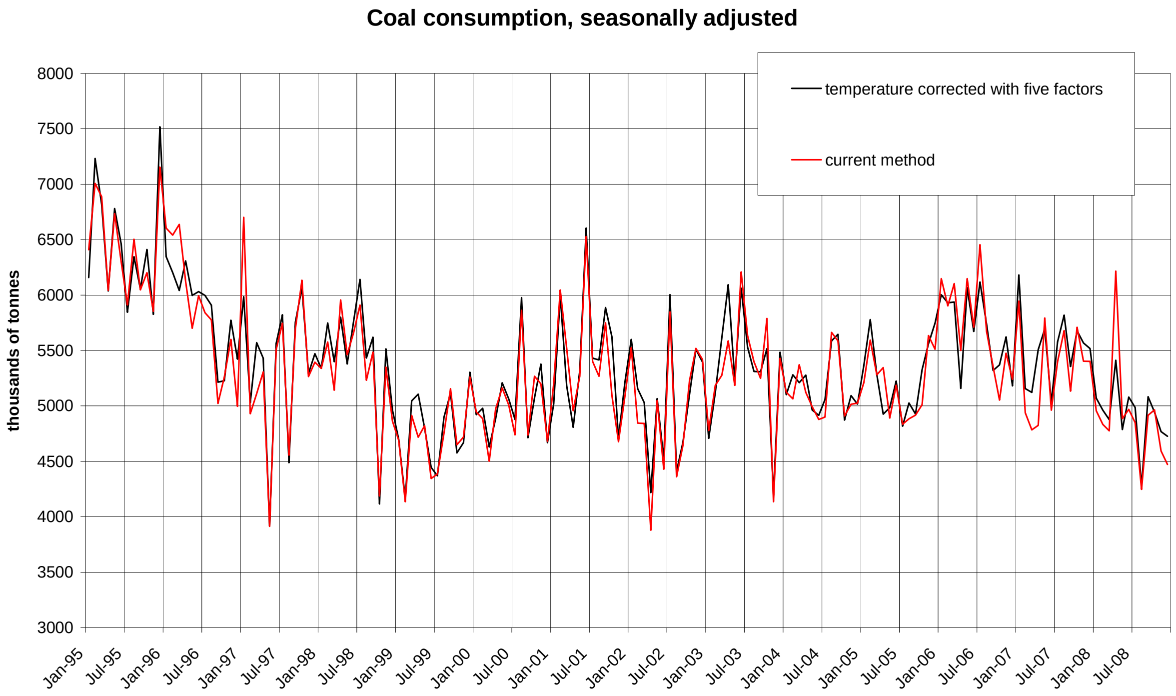

Figure 7 - coal, seasonally adjusted after the current temperature correction and after a new temperature correction

3.11 After discussion with the Department, the statisticians felt that this five month model would be difficult to justify and interpret, with large factors for the five significant months and no correction at all for the other months. It would be preferable to have a set of correction factors that change more smoothly throughout the year. To help guide our selection of a new set of factors we can start looking at the gas series and see what the correction factors look like in this case.

GAS

3.12 We use an identical method to arrive at monthly correction factors for gas. The resulting model is highly significant, more so than with the coal model ($\chi^2 = 545$ with $df = 12$, $p < 0.005$). This time each coefficient is significant and negative.

Variable	Estimate	Standard error	t-value
Jan	-0.042	0.009	-4.50
Feb	-0.055	0.007	-8.05
Mar	-0.074	0.008	-9.64
Apr	-0.089	0.010	-9.36
May	-0.090	0.011	-8.16
Jun	-0.049	0.016	-3.01
Jul	-0.039	0.007	-5.29
Aug	-0.027	0.010	-2.58
Sep	-0.063	0.011	-5.57
Oct	-0.074	0.007	-11.0
Nov	-0.049	0.011	-4.51
Dec	-0.051	0.009	-5.62

3.13 The M statistics for the seasonal adjustment following this model are given below, with those already seen for the seasonal adjustments for the uncorrected series and the series corrected by the more complex model. We see that the new method is not as good as the current one in terms of these statistics, but the resulting seasonal adjustment is nevertheless very satisfactory.

	Uncorrected	Corrected using current method	Corrected using the 12 monthly factors
M1	0.049	0.007	0.016
M7	0.112	0.044	0.056
Q	0.66	0.35	0.52

3.14 Since all these factors are negative and significant, we can return to our model for coal and try to adjust the correction factors so that they vary in a similar way to these gas factors.

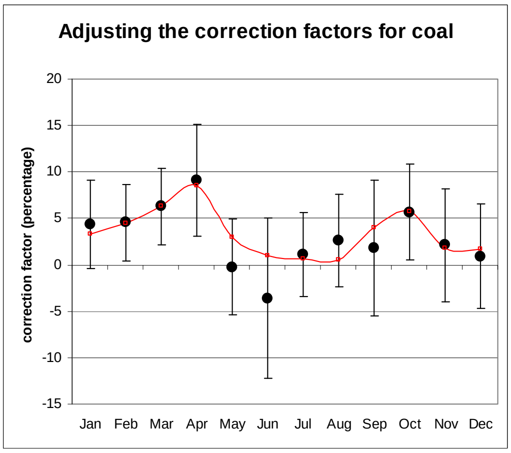

Figure 8 – new correction factors for coal

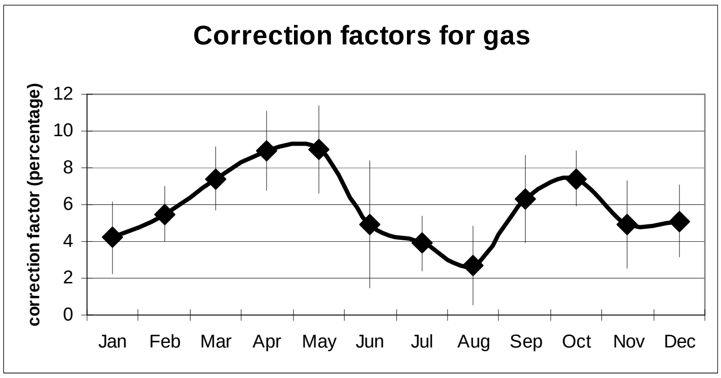

Figure 9 - new correction factors for gas

3.15 Figure 8 shows the estimated factors for coal in the model described in paragraph 3.6 (black circles with 95% central confidence intervals), with the red line showing the adjusted factors, so that they vary in a similar way to the factors for gas (Figure 9). The negative coefficients in the model have been graphed as positive correction factors. Note that no factors have been moved beyond their 95% confidence interval, and that the factors are all the same sign.

3.16 There is a case to be made for adjusting the factor for May in the gas model. The Department is sceptical that there is as much sensitivity to temperature deviation from the long-term average in May as there is in April and in the coal model the confidence interval for May barely overlaps with April's. Adjusting the May factor to 6.9 (the mean of the April and June factors) produces a curve which is similar in shape to the coal factors' curve and still keeps all factors within their confidence intervals.[3]

3.17 As well as looking at the standard errors we also considered the reliability of these estimates by running the regressions on different spans of data. The full regressions were on the 14 year time period 1995 – 2008; the shorter spans were 1995 – 2006, 1996 – 2007, 1997 – 2008 (three spans of 12 years) and 1995 – 2007 (a 13 year span). The

1.1

---

[3] In section 4, when models are fitted using the concept of heating degree days, there is clearer evidence of a peak in April.

gas estimates were generally stable, with the full regression estimates all within 0.002 of the median of the estimates derived from the five different spans (except for September where the difference was 0.004). The coal estimates were less stable, but even so for eight of the months the estimate from the 14 year span was within 0.002 of the median of the estimates from all five spans. The results of this sensitivity analysis were taken into account when arriving at the adjusted factors in paragraph 3.14.

3.18 Something that the sensitivity analysis confirmed was that good quality estimates for regression models do rely on the data coming from a suitably wide range. The span which produced the most anomalous estimates was the 1997 – 2008 span. This is the only span which does not include 1996. This is significant because for the whole period 1995 – 2008 1996 was the coldest year (in fact the only year whose average temperature was below the long-term average 1971 – 2000). The starkest example is the estimate for the coal correction factor for May. For the other four spans the estimates ranged between 0.2% to 0.5% (with standard errors all below 0.3%), but for this span the estimate was 8.8% ! If we look at the temperatures for the statistical month of May, during this span the temperature deviations range from 0.0 to +1.9, but the deviation in 1996 (which is not in the span 1997 – 2008) is -2.2. DECC needs to be aware of this issue when it moves to a new long-term average base (e.g. 1981 – 2010) and produces a new set of deviations. When this happens the Department will want to make sure that any new regressions make use of a good range of temperature deviations (or when they do not, that the resulting estimates are treated warily).

FINAL FACTORS FOR COAL AND GAS

3.19 So the final recommended correction factors for the coal and gas models are:

Month	Correction factor (gas)	Correction factor (coal)
Jan	-0.042	-0.033
Feb	-0.055	-0.045
Mar	-0.074	-0.063
Apr	-0.089	-0.085
May	-0.069	-0.030
Jun	-0.049	-0.010
Jul	-0.039	-0.007
Aug	-0.027	-0.005
Sep	-0.063	-0.040
Oct	-0.074	-0.057

Nov	-0.049	-0.018
Dec	-0.051	-0.017

3.20 If DECC wishes to make these new factors the *only* changes to its current methodology, then it will be very simple to introduce: on its processing spreadsheets there is a column for coal with the correction factors for each month. The 12 factors need to be put into one year and then copied to the other years. A similar column would be needed for the gas spreadsheet. (The way that the spreadsheets are currently set up, the January factor for coal, for example, which is "-0.033" in the table above, would need to be entered as "3.3"; similarly the January gas factor would need to be entered as "4.4" if the same formula is copied from the coal spreadsheet.) This means that the interpretation of these factors is exactly as described in paragraph 1.3, and not approximately (as explained in paragraph 3.8).

3.21 The two graphs below compare, for coal and for gas, the seasonal adjustment using the current correction factors and using the proposed new factors.

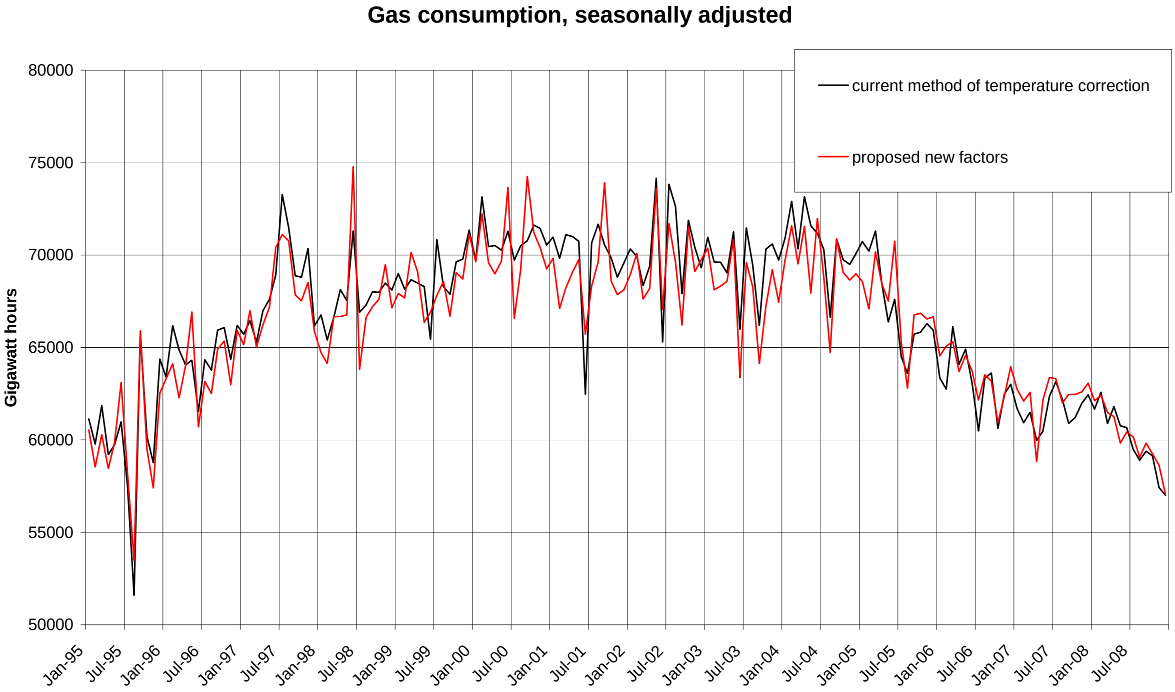

Figure 10 - gas, seasonally adjusted after the current temperature correction and a proposed new temperature correction

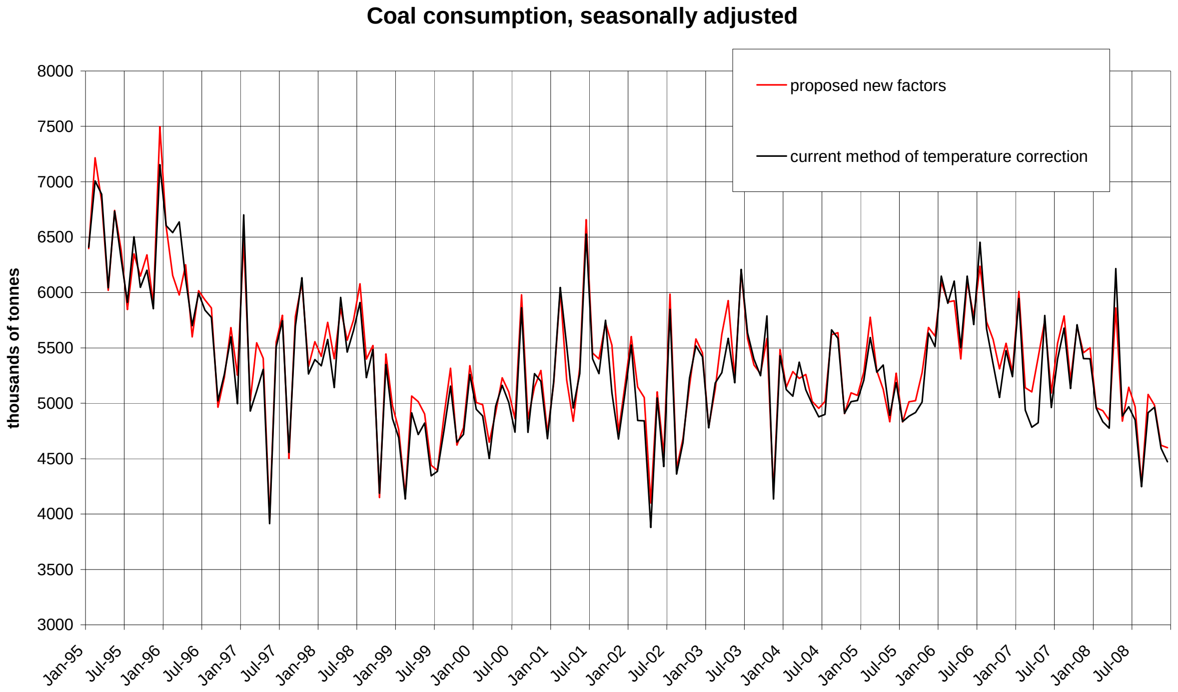

Figure 11 - coal, seasonally adjusted, after the current temperature correction and a proposed new temperature correction

3.22 The M statistics comparing the seasonal adjustments for coal for the uncorrected series, the current method of temperature correction and the method in 3.19 are as follows:

	Uncorrected	Current temperature correction	Using the factors in 3.19
M1	1.157	0.837	0.741
M7	0.335	0.329	0.298
Q	0.95	0.91	0.92

3.23 These statistics for the recommended factors show an improvement on the current method, and compare well with the 5 factor model (see 3.10).

3.24 In paragraph 2.19 we commented that the mean 12 monthly change in the coal series went down from 12.4% to 11.8% with the temperature correction, meaning that the temperature correction was helping to smooth the series. For these new correction factors the mean change reduces further (slightly) to 11.2%. However for the proportion of variation in 12 month growth rates, these new factors do not improve on the 42% ‘score’ of the current temperature correction, being higher at 49% (which is no different from the uncorrected coal series).

3.25 The M statistics comparing the seasonal adjustments for gas for the uncorrected series, the current method of temperature correction and using the factors in paragraph 3.19 are as follows:

	Uncorrected	Current temperature correction	Using the factors in 3.19
M1	0.049	0.007	0.015
M7	0.112	0.044	0.056
Q	0.66	0.35	0.51

3.26 These statistics show that the proposed factors improve on the uncorrected series, so that the temperature correction is worthwhile. It is not as successful as the current method but still produces an excellent seasonal adjustment, and the minor change to the factors (between 3.12 and 3.19) has barely changed the quality of the adjustment. There is a similar pattern when we consider the mean 12 month change, with the 7.4% average for the uncorrected series reducing to 4.2% with the proposed temperature correction: an improvement, but not as good as with the current method where the average change reduced to 3.8%. Variation in the 12 month growth figures that is attributed to the irregular component has reduced from 55% of total variation (in the uncorrected series) to 36% with these correction factors (it is just 22% using the current temperature corrected series).

PETROL AND NUCLEAR

3.27 A model was fitted to the petrol series but it was not significant ($\chi^2 = 12.9$ with $df = 12$, $p < 0.38$). This model did not lead to a reduction in the average 12 month change, which stayed at 7.5%. We can see if the seasonal adjustment of petrol as it is currently done helps with the indirect seasonal adjustment of the total series. This would offer us one reason to continue with the seasonal adjustment of petrol. As Figure 12 shows, there is no real benefit to the total consumption series when the current temperature and seasonal adjustment takes place. There is clearly a level shift (because of the temperature correction and the fact that the period 1995 – 2008 was warmer on average than the period 1971 – 2000).

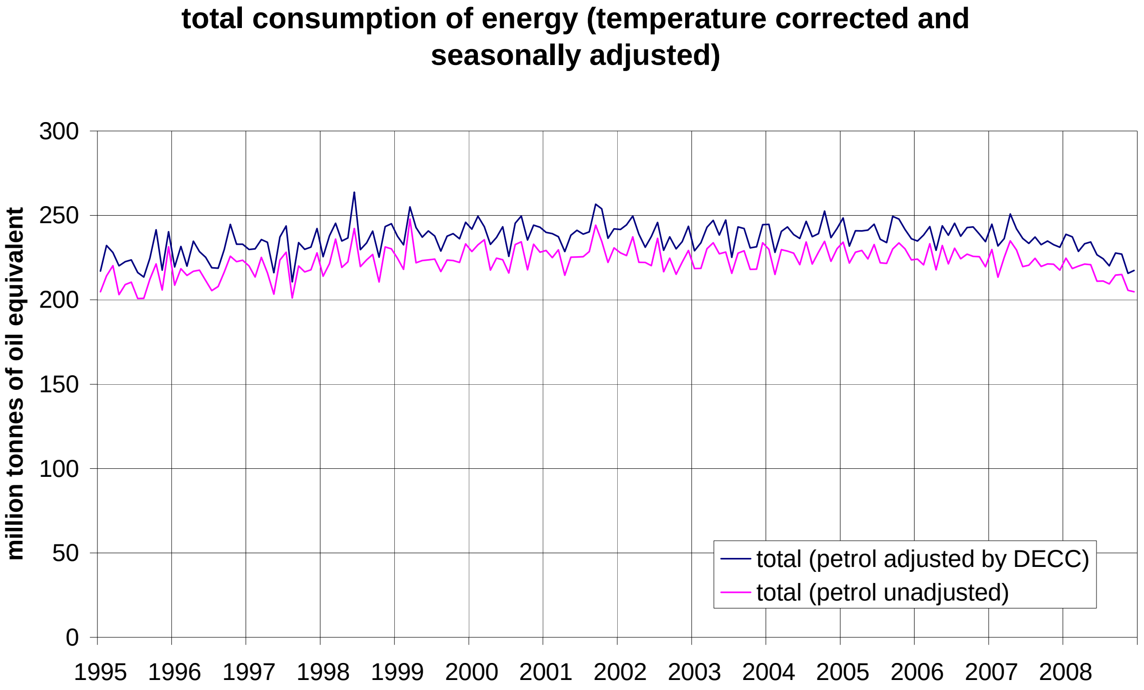

Figure 12 - the total energy consumption series, with petrol consumption unadjusted and adjusted

3.28 We recommend that the petrol series is no longer seasonally adjusted or corrected for temperature.

3.29 As discussed in Section 2, the nuclear series is not currently corrected for temperature as it is part of the baseload production. A quick attempt to fit a regression model shows that there is no significant set of correction factors and we recommend that DECC maintain its current method for dealing with the nuclear series (do not correct for temperature, but continue to seasonally adjust, as discussed in paragraph 2.25).

4 Heating degree days

Summary of main points

- *Using the degree days as published by Eurostat leads to estimated correction factors for gas that have smaller standard errors than the regression estimates from the previous section*
- *The use of 'heating degree days' with a base temperature of 18 degrees is a close approximation to the Eurostat concept and can be estimated quickly from data that DECC already has on average monthly temperatures*
- *The more usual heating degree days basis that is used for the UK (15.5 degrees) results in less precise estimates (although more work needs to be done to determine if this is simply to do with the use of Hitchin's formula to approximate the heating degree days instead of directly calculating them)*

4.1 The concept of heating degree days appears to date from the 1870s, when Lt-Gen Sir Richard Strachey used them to measure the length of a crop growing season, and they are still used to assess crop growing conditions. The use of heating (and cooling) degree days has now become prominent in energy management for buildings.[4]

4.2 Heating degree days (we shall refer simply to degree days from now on unless explicitly making a contrast with cooling degree days) are essentially the sum of differences between actual temperatures and a pre-specified baseline temperature. If temperatures are being measured continuously then a measurement of a degree day amounts to finding an area beneath a curve (the area between the temperature curve and the baseline temperature). There are a number of approximations available for calculating degree days if temperature measurements are available on a daily or monthly basis.

4.3 There are at least two ways in which using degree days to perform a temperature correction on the energy consumption series might be conceptually better than the current method which simply looks at the monthly average temperature.

4.4 First, degree days are unidirectional with respect to the baseline temperature. This means that when the temperature is below the baseline, there is a positive contribution the degree days measure (the number of extra degrees of heating required to get up to the baseline), but when the temperature is above the baseline there is no offsetting contribution to the degree days measure. In contrast, using the average temperature for the month means that warm days and cool days cancel each other out. We think that the degree days approach models domestic users' behaviour better: during a day that is

1.1

---

[4] Much of the background information on heating degree days in this section has come from the technical memorandum published by the Chartered Institute for Building Services Engineers (*Degree-days: theory and application – CM41:2006*; Prof. Tony Day is the principal author). This volume is referred to as CIBSE for the remainder of this report. It is available by searching for 'degree days' at www.cibse.org

colder than the target room temperature we are more likely to put the heating on (or set the thermostat higher) but when the temperature is already above the target room temperature we can only switch the heating off and throw open the windows (there is no 'negative consumption' of energy). In some climates there is likely to be a case for expecting increased energy consumption during hot days (e.g. to run ceiling fans or air conditioning units) and this is what cooling degree days can be used for, but we do not think it likely that there is enough of this type of energy consumption in the UK for cooling degree days to be of much use in this project.

4.5 The second conceptual advantage is related to the first. Because degree days measure periods when the temperature is low but do not include days when heating systems are not used, they better account for extreme events such as a cold snap that takes place for only part of a month. The CIBSE publication says (page 2) "this makes them more reliable in estimating energy consumption, particularly in the milder months, but also in those periods with extreme cold snaps where they capture both magnitude and duration of an event".

4.6 As a simple example, using daily average temperatures, imagine a month with 15 days of average temperature 16° (Celsius) and 15 days of 12°. If we take the baseline temperature to be 15.5° (a common baseline temperature historically for UK degree days databases), then the total degree days for the month is 15 × (15.5 – 12) = 52.5 and the average daily degree day is 1.75. This is close to the average deviation from 15.5°, which is 1.5. Suppose now that five of those 16° days were actually 20°. Then the degree days calculation is unchanged, but the average temperature deviation from 15.5° is now 0.83, less than half the average degree day measure for the month.

4.7 The choice of a baseline temperature is important. Converting a degrees day figure from one baseline to another is not a straightforward matter (there will not be a simple ratio because of the unidirectional nature of the calculation)[5]. In building energy management, the baseline temperature is the temperature that it needs to be outdoors so that no energy is required for heating the building. For this project it seems reasonable to think of it as the 'target' temperature that people want to heat their homes to. Historically in the UK degree days have been calculated for a baseline temperature of 15.5° (e.g. by the Met Office). But this baseline is based on American dwellings data from the 1920s and seems rather implausible as a target temperature now.

4.8 There is a slightly different measure used by Eurostat for its energy statistics. These degree days work with two thresholds, 15° and 18°. When the daily average temperature is 15° or less, then that day's contribution to the degree days measure is 18° minus the average temperature. But if the temperature is more than 15° then there is zero contribution to the degree days measure. So any single day's degree day measure will be zero, or a number that is greater than or equal to three.[6]

4.9 The Eurostat degree days are available for the UK on a monthly basis from the Eurostat website a few months in arrears. We have not been able to establish who supplies this

1.1

---

[5] The CIBSE manual (Appendix A3) describes how to do this.

[6] see http://epp.eurostat.ec.europa.eu/cache/ITY_SDDS/en/nrg_esdrg_esms.htm for the reference metadata for the Eurostat heating degree day concept.

data to Eurostat but the Met Office clearly produces data that could be used to calculate these degree days (daily average temperatures).

GAS

4.10 Three regressions were run on the gas series, one using the Eurostat degree days measure, one using the usual degree day measure with baseline temperature 15.5° and one with the baseline temperature of 18°.

4.11 The variable to use in the regression is the deviation in the mean daily degree day for each month from that month's long term mean daily degree day. A deviation of 1 for January 1995, for example, means that in January 1995 on average an extra degree of heating was required per day, compared to the average amount of heating required in Januaries between 1971 and 2000. (This means that January 1995 would have been cooler than the average January.)

4.12 For the regression on the Eurostat degree days, we needed to estimate the long term averages for each month for the period 1971 – 2000. These were not available on the Eurostat website, so were estimated by performing a simple linear regression of Eurostat degree days on average monthly temperatures for the UK. Data were available for the period January 1991 to May 2010.

4.13 There is a very strong linear relationship between mean daily Eurostat degree days and mean monthly temperatures, with a correlation coefficient of -0.998. So the regression equation $18.291 - 0.98x$, where x is the temperature, was used to convert the long term average temperatures for each month into long term average Eurostat degree days for each month. These were subtracted from the actual daily degree days for each month (1995 – 2008) to arrive at the deviations to feed into the regression.

4.14 For the regressions on the usual degree days measures, we did not have daily temperature readings for the calculation of the degree days for the period 1995 – 2008 (or to calculate the long term average). Instead we used Hitchin's formula[7] to convert the mean monthly temperatures into estimated daily average degree days for each month during the period 1995 – 2008, and also to convert the long term average temperatures into long term average daily degree days.

4.15 Hitchin's formula approximates degree days in the UK well, although there is a greater chance of inaccuracy when the average temperature is close to the chosen baseline temperature. The formula is $\frac{\theta_b - \theta}{1 - e^{-0.71(\theta_b - \theta)}}$, where θ_b is the baseline temperature, θ the mean monthly temperature (assumed different from the baseline) and 0.71 an empirically chosen constant that is considered to produce good results for estimating degree days for the UK. Figure 13 shows the close approximation between Hitchin's formula with a baseline temperature of 15.5° and an average of degree days we have

1.1

---

[7] See CIBSE, section 2.4 for more details on Hitchin's formula.

calculated using the BizEE website's database[8] (using a set of weather stations as close as possible to those used by DECC to calculate UK-wide average temperatures).

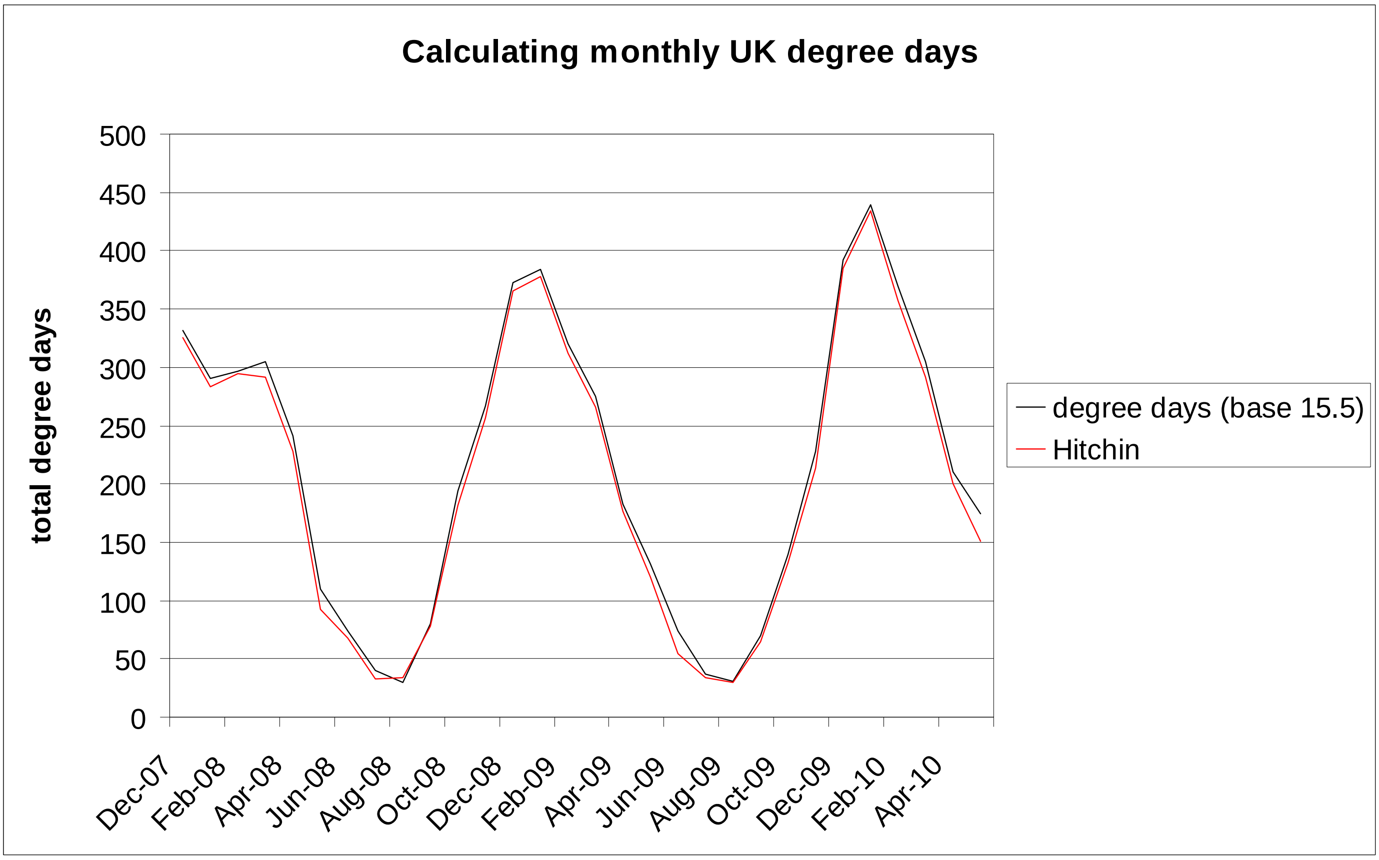

Figure 13 - using Hitchin's formula to estimate monthly degree days

4.16 The formula was used to calculate degree day deviations to feed into regressions; once with the 15.5° baseline temperature and once with 18°.

4.17 When comparing the estimates from the three regressions it is important to remember that the various approximations employed, in different ways, mean that comparisons need to be made with some care. In the case of the Eurostat degree days, we have used a simple linear regression to calculate the long term averages, and the Eurostat data have (presumably) not used the same weather stations that DECC uses to arrive at the UK wide figures.

4.18 The graphs below summarise the estimated factors and their estimated standard errors from these three regressions. Also included are the estimates from the regression done in Section 3 – these are the lilac bars. (Note that the factor for May is the original one estimated in 3.12, and that the factors from Section 3 have been made positive to aid comparisons. Factors for models using the degree days concepts are positive because a positive deviation corresponds to a cooler than average month – the 'other way round'

1.1

---

[8] www.degreedays.net

compared to temperature deviations.) ‘H15.5’ is the model using degree days to a baseline temperature of 15.5°.

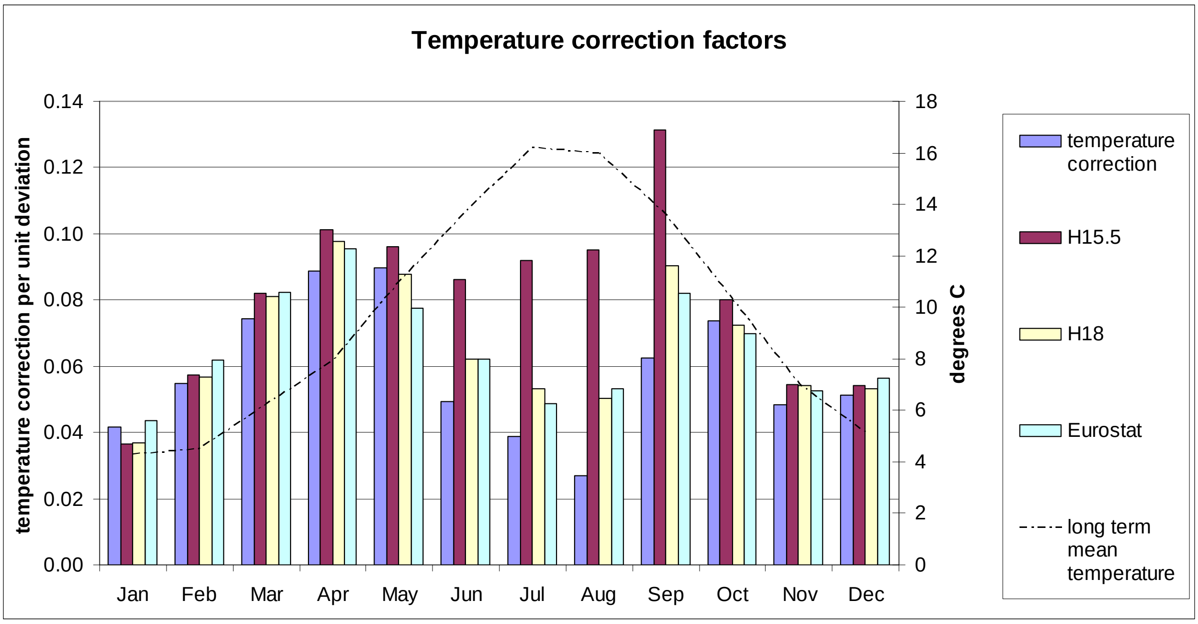

Figure 14 - different models' estimates for gas correction factors

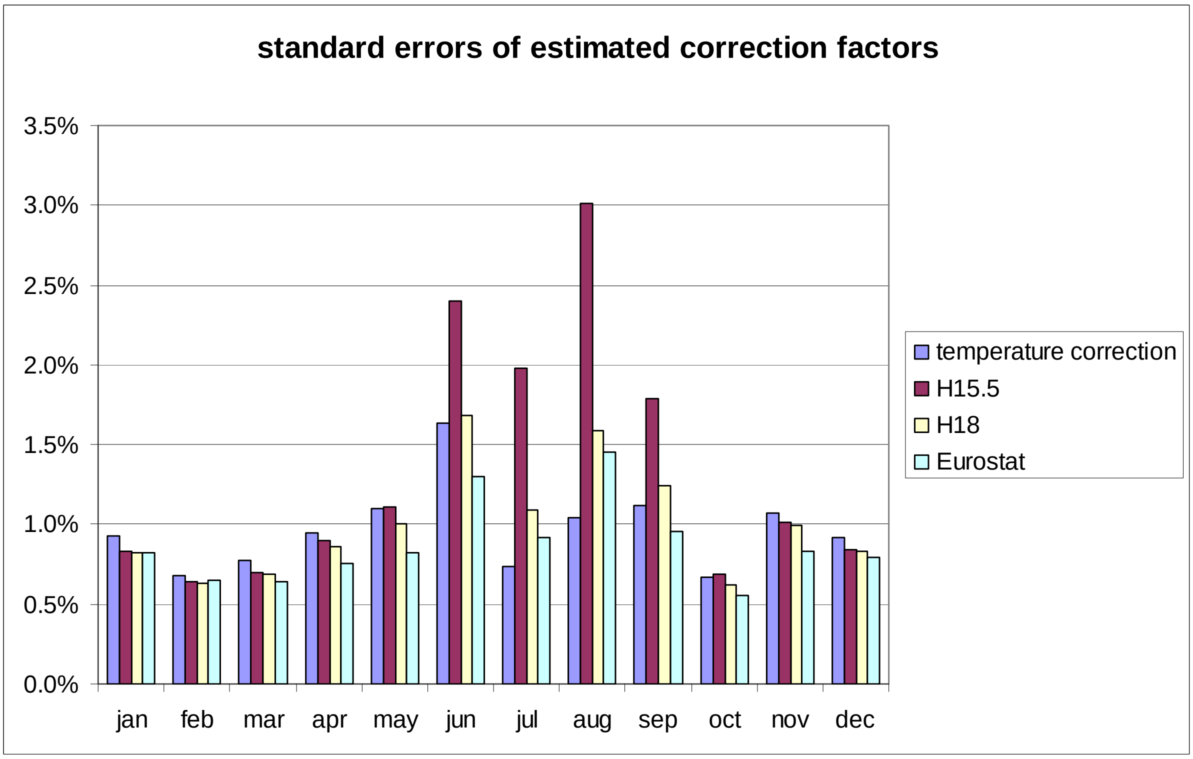

Figure 15 - standard errors for the estimated gas correction factors

4.19 It is clear that the H15.5 model is anomalous, but only for months where the average long term temperature is above 10°. This could be down to the fact that the baseline temperature is unrealistically low, and that energy users in the UK prefer to keep their rooms at a temperature closer to 18° than 15.5°. There is perhaps also an issue with the use of Hitchin's formula to convert average temperatures into average daily degree days: the formula becomes less accurate when temperatures are close to the baseline, and in these warmest months there will be many instances where the average temperature is close to 15.5°.

4.20 For the warmest months the temperatures can get close to 18° and if there is a 'formula effect' then we would expect an effect on the H18 model as well. There is some evidence of this, in that the two warmest months (July and August) are the months where the standard errors of the H18 differ most from those of the temperature correction model. These two months have averages of 16.2° and 16.0°, close to the baseline temperature of 18°.

4.21 It might well be the case that, without this formula effect, there is little to choose from between the H15.5 and H18 models. Of course one way to remove this formula effect (if it is real) is to make the baseline temperature very high, but then this would remove any advantage the degree day concept has over simply using average temperatures. It remains the case that the baseline of 18° is a more realistic one than the lower value of 15.5°, even if we were to use a more direct method of measuring the degree days (i.e. without using Hitchin's formula).

4.22 From the graphs we see that, ignoring the H15.5 model, the correction factors move in similar ways, but the Eurostat model has smallest standard errors. Its chi-squared value is 746, indicating a better fit than the H18 model (624) and the temperature correction model (545) (all on 12 degrees of freedom, so all are very significant). Looking at the M statistics to see the impact on the seasonal adjustment of the gas consumption series we see that this Eurostat model improves on all our other models so far, but not the temperature correction that is currently performed by DECC's data supplier.

	Current temperature correction	Using the factors in 3.19	Eurostat	H18
M1	0.007	0.015	0.012	0.015
M7	0.044	0.055	0.056	0.055
Q	0.35	0.52	0.48	0.50

4.23 The model using the Eurostat degree days is the best for DECC to use if the current temperature correction method is to end (see paragraph 1.5). However the Department will not be able to produce its monthly output using the Eurostat website, as the data there are four months in arrears and so not timely enough. It would be straightforward to produce degree days following the Eurostat concept with the data the Department currently receives from the Met Office, since the temperature data are broken down daily. The Department's spreadsheets for processing the data would need an extra column and a simple =IF formula to convert a day's average temperature to a Eurostat degree day. The only difference from the published Eurostat data would be due to different weather stations being weighted together to arrive at a UK average. Such differences are likely to be very minor and it is unlikely that the temperature correction factors would need to be changed.

4.24 The correction factors when using deviations in daily Eurostat degree days are as follows:

Variable	Correction factor	Standard error
Jan	0.044	0.008
Feb	0.062	0.006
Mar	0.082	0.006
Apr	0.096	0.008
May	0.078	0.008
Jun	0.062	0.013
Jul	0.049	0.009
Aug	0.053	0.015
Sep	0.082	0.010
Oct	0.070	0.006
Nov	0.052	0.008
Dec	0.056	0.008

COAL

4.25 We were not able to fit a model to the coal data using Eurostat degree days because we did not have the degree day data available on the statistical months basis that the coal consumption data were collected on. It would be easy (though perhaps time consuming) for DECC to calculate the degree days for each month and then run the regression.

4.26 We could use the relationship described in paragraph 4.13 to convert each statistical month's average temperature into a daily Eurostat degree day, as well as the long term averages, but then this would in effect be the same model as the original regression on temperature deviation performed in Section 3. Indeed, since we would calculate all Eurostat degree days using $18.291 - 0.98x$, the factors for each month would be the same as those for the temperature deviation models multiplied by -0.98 (and the standard errors would be multiplied by 0.98, so that the t-values would be identical except for a change in sign and the overall significance of the model would be identical). This was all confirmed by actually producing the estimates and comparing.

4.27 Using the same ideas as described in 3.14 onwards, the coal model's factors have been adjusted so that they vary in a similar way to the gas model's, while paying due attention to the 95% confidence intervals and the sensitivity of the estimated factors to the different timespans. The biggest difference now compared to what was done in paragraph 3.14 is that more 'weight' has been given to following the shape of the curve of the Eurostat gas model, given that it was overall a better fit, with smaller standard errors, than the temperature deviation gas model in Section 3.

4.28 The factors for the coal and gas series using the Eurostat degree days can be seen here. Again, the black dots represent the original estimates for the coal model and the red line the proposed adjusted factors.

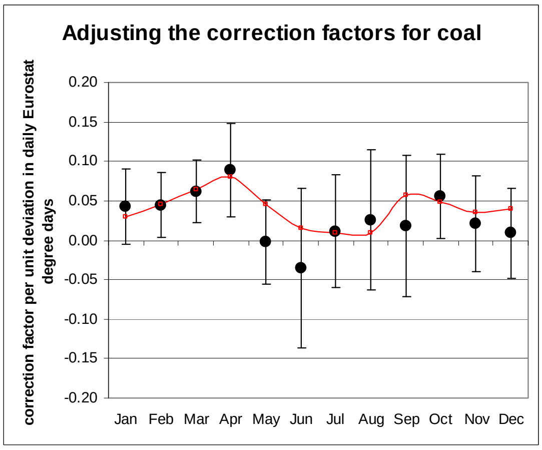

Figure 16 - adjusting the correction factors for coal in the Eurostat model

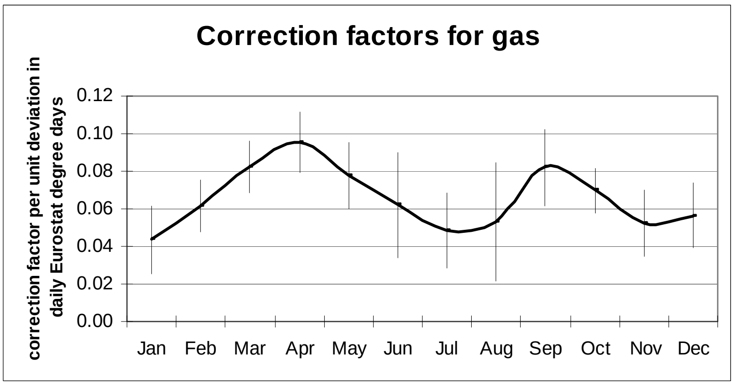

Figure 17 - correction factors for gas in the Eurostat model

4.29 The *pattern* of the black dots is identical to that seen in 3.14, since these factors are just 2% smaller than the temperature deviation model's (as explained in 4.26). But the gas model's factors' pattern is different. The new model's curve is clearly a bit smoother, there is a clearer single peak in April (with May now looking rather like what we recommended adjusting the gas model to in 3.16), the autumn peak is in September and the summer months are a bit higher.

FINAL RECOMMENDED FACTORS FOR COAL AND GAS – EUROSTAT DEGREE DAYS

4.30 The recommended factors for the gas and coal models, then, when using the Eurostat degree day measure, are:

Month	Correction factor (gas)	Correction factor (coal)
Jan	0.044	0.030
Feb	0.062	0.045
Mar	0.082	0.064

Apr	0.096	0.080
May	0.078	0.045
Jun	0.062	0.015
Jul	0.049	0.010
Aug	0.053	0.010
Sep	0.082	0.057
Oct	0.070	0.048
Nov	0.052	0.036
Dec	0.056	0.039

4.31 The M statistics for the resulting seasonal adjustments, given in the following table, compare well with the values from other models. The values for coal are better than with the current temperature correction method, and just marginally better than with the improved temperature correction factors. For gas, these statistics indicate that the Eurostat model is slightly preferable to the proposed temperature correction method, but once again the current method for gas is best.

	Gas	Coal
M1	0.012	0.725
M7	0.057	0.309
Q	0.48	0.92

4.32 The average 12 monthly change for coal is 11.2%, the same as for the improved temperature correction factors. The percentage of variation in 12 monthly growth that is attributed to the irregular is 46.1%, a bit worse than we saw for the improved temperature correction (48.8%). For gas these figures are 4.3% and 33.1% respectively for the Eurostat model and 4.2% and 35.8% for the proposed temperature correction method.

4.33 As an example to help interpret these factors, consider January 2010. The Eurostat degree days figure for January 2010 is 504.878. Divide this by 31 to get the average daily degree days for January 2010 – this is 16.286. The long term average daily degree day for January during the period 1971 – 2000 is 14.108 (paragraph 4.13 explains how this was calculated), so the deviation is 16.286 – 14.108 = 2.179. Because the correction factor for January is given (in 4.30 above) as 0.044, which is 4.4%, we can multiply this deviation by 4.4% to get 9.6%. This means that the gas consumption figure for January 2010 should be adjusted down by (about) 9.6%. In fact the true correction is performed by multiplying the consumption figure by exp(-2.179 × 0.044) which is 0.909, implying the true correction is actually 9.1%.

4.34 Overall, there is very little to choose between the Eurostat model and the others. A crude ranking based on the five diagnostic statistics that we have been using in this report (the two statistics from the F table that summarise the influence of the irregular factor on 12 month growth rates, and the three M statistics summarising the quality of the seasonal adjustments) shows that the Eurostat models perform best (ignoring the current temperature correction on gas) and the H18 model described below performs worst, with section 3's updated factors in-between. (This overall ranking is reasonably robust to different plausible weightings of these five statistics.) We recommend that DECC move to using the Eurostat method not just because it performs best in this comparison, but also because it represents a conceptually more convincing basis for temperature correction than the current system of using monthly average temperatures. We do not take a view over the merits of the Eurostat degree day concept over the more standard degree days (with baseline 18°), but the temperature correction using the Eurostat concept is clearly more satisfactory. This might be merely a matter of having a better dataset to use for the Eurostat data. Even so, it is probably worth using the Eurostat degree day if Eurostat is committed to harmonising the degree day across the EU for energy consumption statistics, since it is producing satisfactory results.

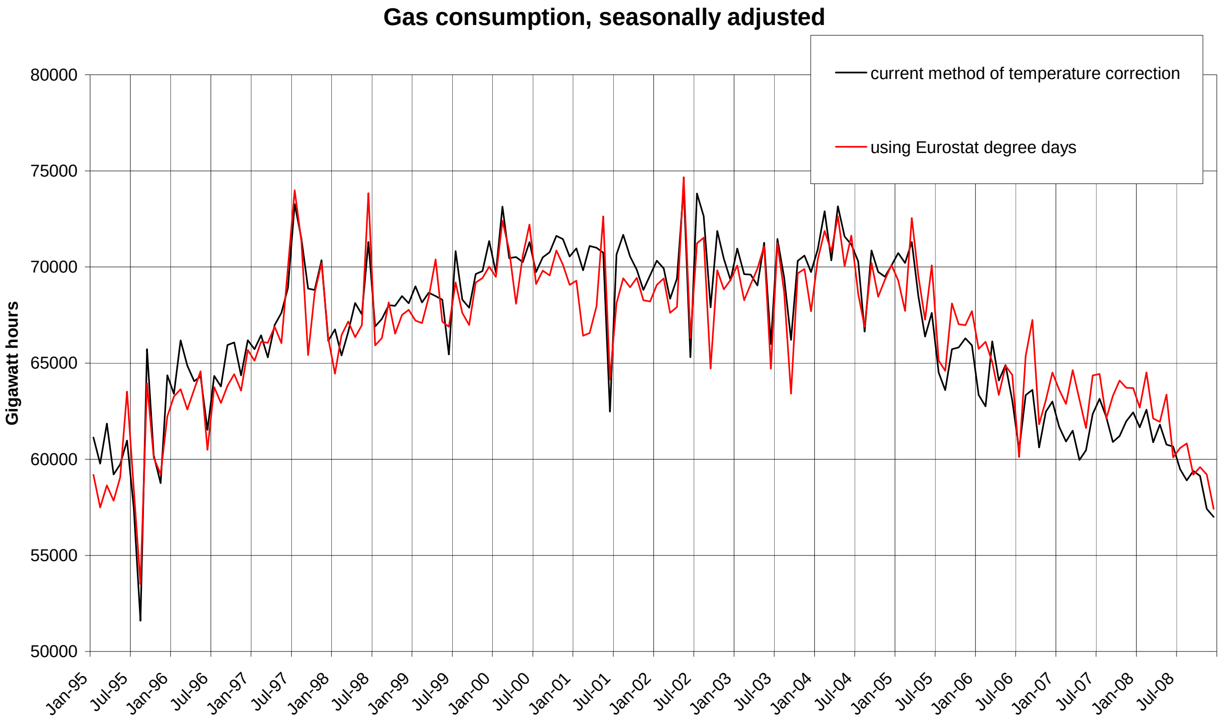

Figure 18 - Gas, seasonally adjusted, after the current temperature correction and the new, recommended temperature correction using Eurostat degree days

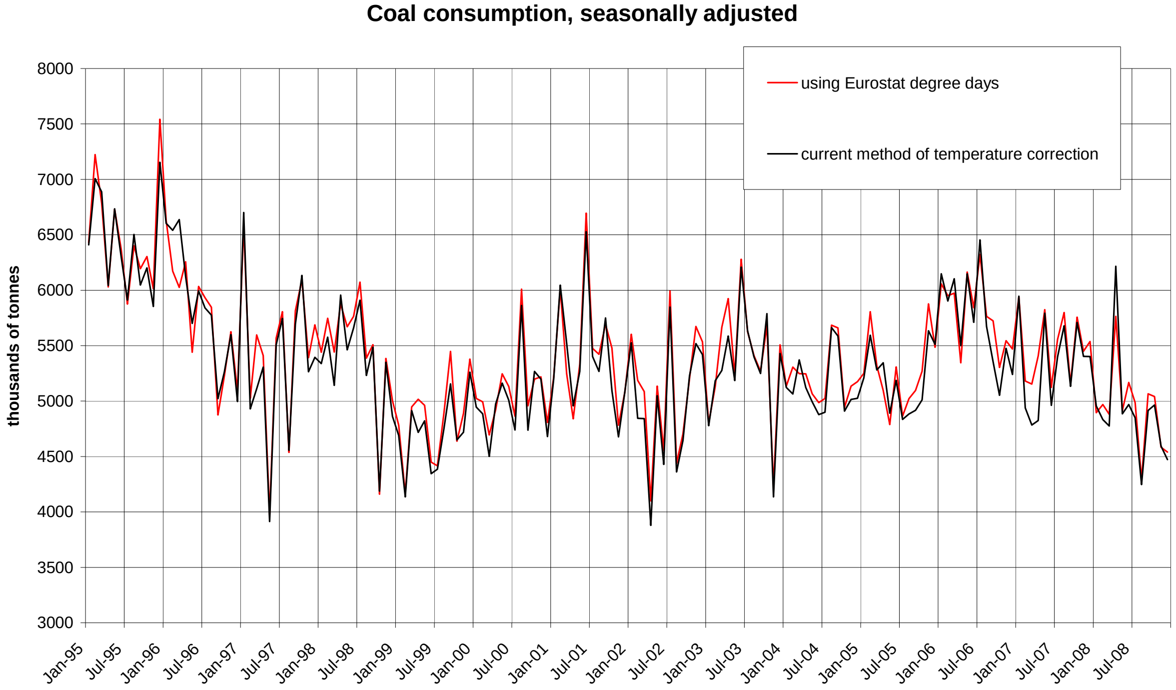

Figure 19 - coal, seasonally adjusted, after the current temperature correction and after the new, recommended temepráture correction using Eurostat degree days

CONVENTIONAL HEATING DEGREE DAYS

4.35 Finally, if DECC decides to use the usual degree days measure instead of the Eurostat concept (if, say, users are more comfortable with the usual measure) then we recommend the following correction factors, working with a baseline temperature of 18°.

Month	Correction factor (gas)	Correction factor (coal)
Jan	0.037	0.030
Feb	0.057	0.046
Mar	0.081	0.066
Apr	0.098	0.092
May	0.088	0.051
Jun	0.062	0.035
Jul	0.053	0.025
Aug	0.050	0.025

Sep	0.091	0.074
Oct	0.073	0.059
Nov	0.054	0.044
Dec	0.053	0.043

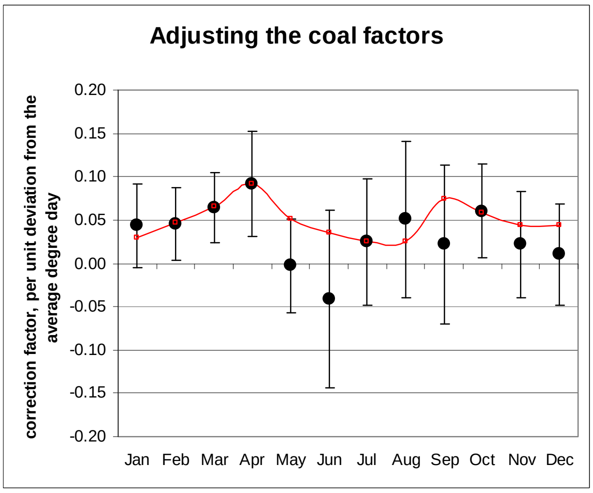

Figure 20 - adjusting the correction factors for coal in the H18 model

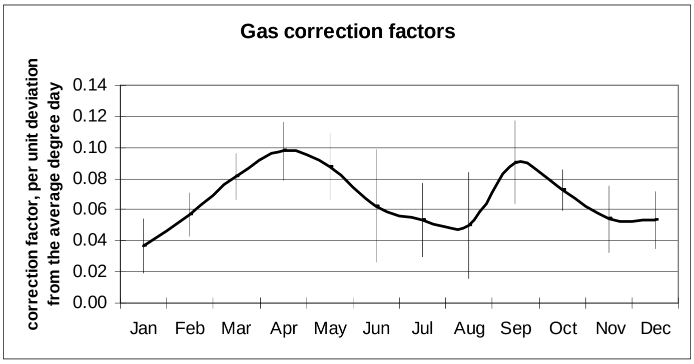

Figure 21 - correction factors for gas, for the H18 model

4.36 DECC may choose to use Hitchin's formula to convert older monthly temperature data into degree days, but can easily calculate future degree days with the daily temperature data it currently receives from the Met Office. There are three options:

4.37 This would be easy to implement into the Department's current set of spreadsheets for processing the Met Office temperature data. For each weather station, for each day subtract the average temperature from 18° unless the average temperature is 18° or more, sum to get a total for the month, obtain the weighted average to get a UK figure, and divide by the number of days in the month to get the average daily degree day for the month for the UK.

4.38 The second method is perfectly feasible because DECC receives the daily maximum and minimum temperatures for each station. This method uses the so-called Met Office formula (also known as the 'McVicker formula' or the 'British Gas formula'). This method makes use of the fact that temperatures usually follow an approximate sinusoidal curve throughout the day, which can potentially make the method of using just the day's average temperature inaccurate when the minimum and maximum are on different sides of the baseline temperature. For each day, the degree day is calculated as follows[9]:

1.1

---

[9] section 2.2 of the CIBSE manual gives more details for the interested reader.

- Let T_{max} and T_{min} be the day's maximum and minimum temperature and let T_b be the baseline temperature (e.g. 18°)

- If $T_b \geq T_{max}$, then the degree day is $T_b - 0.5(T_{min} + T_{max})$

- If $T_b > T_{min}$ and $(T_b - T_{min}) > (T_{max} - T_b)$, then it is $0.5(T_b - T_{min}) - 0.25(T_{max} - T_b)$

- If $T_{max} > T_b$ and $(T_{max} - T_b) > (T_b - T_{min})$, then it is $0.25(T_b - T_{min})$

- Otherwise $T_{min} \geq T_b$ and the degree day measure is zero

4.39 The third option is to see if the Met Office can provide degree days as part of its service to DECC. If it can provide degree day data with a baseline temperature of 18° then DECC will not need to calculate the data itself. (There is a method for converting between degree day measures of different baselines[10], but it is not straightforward to implement on a spreadsheet and is not recommended for DECC to consider, given how easy it is for DECC to calculate degree days to a baseline of 18° anyway.)

4.40 If DECC does not wish to use the Eurostat degree day, and can produce conventional degree day data using one of these options, then it would be worthwhile running the regressions again and seeing if this improves on what has been suggested in paragraph 4.35.

1.1

---

[10] CIBSE Annex 3 has the details.