NATIONAL GRID COMMENTARY ON OFGEM DEMAND FORECASTING INCENTIVE

We first analyse issues relating to the accuracy part of the proposed scheme, and then issues relating to the bias part of the scheme.

Use of settlement data

- The main users of Demand Forecasts are the ENCC System Operator control room and the Market. The security of the system relies on the System Operator control room receiving accurate forecasts of required generation (referred to within the industry as National Demand or Transmission demand).
- The System Operator control room responds to observed real-time demand – Operational demand – and to the frequency signal of the network to ensure that correct level of generation is being supplied every second.
  - Although we have two customer classes, the security of the system relies on the accuracy of forecasts supplied to the System Operator control room
  - It is more important to the operation of the system that the System Operator control room are provided with accurate operational forecasts, than that market participants are provided with forecasts of the settlement demand after a financial resolution process
  - The incentive for deducing the settlement demand value from the operational demand value should be the financial incentive provided by the cash-in/cash-out prices that market participants are exposed to.
- There is an enduring difference between demand as measured by operational meters and the value of demand determined by the settlement process. On average, operational demand is higher than settlement demand, but there is a time of day structure to the difference, and at some times of day operational demand is usually lower than settlement demand.
- In order to recreate a settlement version of National Demand or Transmission Demand it is necessary to use operational metering, as not all generation included in the definition of these demands are half-hourly metered.
- The difference between operationally defined demand and settlement defined is on average about 200MW, but can be as much as 1800MW.
- If these large differences were entirely down to operational metering errors, the electricity network would have an excess supply during these periods, large enough to cause large frequency deviations. These large deviations do not occur when the differences between settlement and operational measures of demand are large, and hence do not reflect the physical reality of what is occurring on the network.
- Because of this, the System Operator control room cannot safely receive forecasts based on settlement data. If SO has to produce forecasts based on settlement data for the market, it will also have to produce separate forecasts based on operational data for the ENCC System Operator control room. There will then be a risk of two different views of what demand is likely to be, which could well cause confusion between system participants.
- In addition to this, the SO uses short term (day head and in-day) operational data to refine its demand forecasts. This is not possible using settlement data, which has a lag of at least 7 days between real time and availability of the initial estimate, the Settlement Information Run.
- Different estimates of settlement demand data are made on the different settlement runs. The initial view at a lag of 7 days is the Information Run. The R2 run is usually fairly definitive, but for this there is a lag of 4 months. The final run RF has a lag of 14 months.
- Using the R2 run as a base, the mean absolute difference between the Information run and the R2 run is 200MW. We take this as the inaccuracy of (7 day lag) settlement data. It is of the same order of magnitude as the absolute error in operational data.

- The proposed scheme does not specify which settlement run would be used as the basis for calculating National Demand. The longer the lag between forecast and availability of data, the more the scheme depends on (chance) corrections in settlement data.
- What causes differences between operational and settlement data is not currently known. What is clear is that when differences are large, because of the physical consequences that would be apparent on the network, settlement data is more suspect than operational data.
- There are no clear data to support the idea that the first available settlement data is more accurate than operational data.

Restricted and unrestricted demand

- Restricted demand is the demand seen on the transmission network. It may be either restricted National Demand or restricted Transmission demand. Unrestricted demand is the demand that would have been seen on the transmission network if customers did not take uninstructed actions to reduce their demand.
- At present, the only difference between restricted and unrestricted versions of demand occurs during November – February inclusive, when Triad rules for half hourly demand customers are in operation.
- Up to and including day-ahead time scales SO forecasts unrestricted demand. This is to ensure the proper functioning of the triad system. Demand customers need to know what the demand would be if no customer took any triad avoidance actions, so that they can determine whether or not they should take actions. If restricted demand forecasts were published, the triad system would not function properly.
- It would be possible to publish both restricted and unrestricted versions of demand forecasts, but this could lead to confusion for customers.
- Accuracy of triad forecasting at time horizons of day-ahead, or further out, time scales is difficult, because it depends critically on the history of all darkness peak CP outturns up to the previous day. Small fluctuations in the demand outturn on the day before can have up to 2000MW impacts on the next day's outturn.
- It is not possible to measure an outturn value for unrestricted demand, as such a demand is counterfactual and never observed. The SO estimates triad avoidance each day and hence estimates the unrestricted demands. This is critical for our demand forecasting procedures, as demand forecasting models for darkness peak CP are built using historic unrestricted demand values. But they are estimates not measurements.
- If the accuracy of unrestricted demand forecasts is measured against restricted demand outturns there will necessarily be an "error" and a "bias", even if the forecast is 100% accurate.
- The proposed scheme does not deal with these issues, and needs to be amended.

Underlying Systematic uncertainties

- In setting targets for forecast accuracy it is important to understand the sources of uncertainty in demand outturn. Uncertainty is not fixed, but a function of the structure of the electricity system as a whole.
- The main sources of uncertainty are:
    - Uncertainty in day-to-day fluctuations in level of underlying demand (a mixture of economic decisions by businesses and individual human behaviour)
    - Error in weather forecast variables affecting human behaviour
    - Error in weather forecast variables affecting amount of weather-driven unmetered generation (wind and solar generation)
    - Fluctuations in amount of demand met by unmetered non-weather-driven generation (all other unmetered distributed generation)

  - Error in measurement of observed demand (whether measured from operational data or settlement data)
  - Residual statistical error of models employed
- As installed capacities of unmetered generation, wind, solar and other, increases, so does the overall uncertainty, which will inevitably decrease accuracy of future forecasts.
- The SO has traditionally been able to build statistical models with very high explanatory power (R^2 values of 0.95-0.97) and low residual standard error (300-400MW, ~0.5% of maximum daily demand).
  - These R^2 values are considered exceptionally high by academics who have reviewed the SO's forecasting techniques
- With the impact of increased renewable generation the explanatory power of the forecasting models has decreased to around R^2 = 0.95, but the amount of variability has increased, leading to higher standard errors, and lower forecasting accuracy. This is a consequence of a system that allows an increasing proportion of generation to be unmetered (in recent years the unmetered proportion has increased from 7% to 16%).
- The underlying variability will increase as the installed capacity of unmetered generation on the electricity network increases. This will impact forecast accuracy, and must be taken into account when setting targets for accuracy.

Irreducible weather forecast errors

- As the installed unmetered weather-driven (wind and solar) generation increases, the impact of errors in the weather forecasts used to forecast estimates of the component of demand met by these generation sources will increase
  - This is particularly significant for solar generation. The weather forecasting of incident solar radiation is particularly difficult
  - This is true for all weather providers
  - Weather provision is put out to tender every 3 years, and the forecasting service with best overall accuracy in terms of impacts on demand is chosen
  - The problems originate in the physics of clouds
    - There are dozens if not hundreds of research groups working on this around the world
    - It is of huge importance and interest as it is a key factor in determining climate sensitivity to greenhouse gas emissions
    - There are no imminent breakthroughs on the ten-year horizon
    - Given the level of research interest around the world, this is not susceptible to a quick fix
  - No amount of financial investment into weather forecasting research can stop the equations of weather evolution from having a mathematically chaotic nature
  - From the position of the SO weather forecasting errors are an irreducible error
- Although no fundamental breakthroughs into improving weather forecasting processes are on the horizon, the SO is working closely with its weather forecasting providers and with academia to mitigate the shortcomings stemming from these irreducible errors
- The increased contribution of weather forecasting errors as installed unmetered weather-driven generation increases also needs to be taken into account when setting targets for accuracy.

Use of relative error as a performance measure

- Proposed scheme states that relative error will be used, but does not specify how relative error will be calculated.
- In uses where the method of calculation is not specified, it is usually assumed that relative error is error expressed as a percentage of outturn value.

- This would place more value on errors when the demand is low, rather than placing more value on errors when the demand is high. The SO should not be incentivised to place more value on errors when demand is low.
- Several variants exist: relative error as a percentage of maximum daily error, maximum weekly error (avoids incentivising accuracy at weekends more than weekdays), maximum monthly error.
- As the proportion of demand met by unmetered generation increases, using any form of relative error makes the targets tighter in terms of MW (because demand supplied by transmission network is falling) at the same time as uncertainty is rising and demand accuracy (in MW) is falling. This is an unreasonable incentive design.
- The quantity of interest to users of the forecasts is MW demand, not demand as a percentage of some quantity.
- Targets should be based on absolute MW errors, not relative errors.

Setting of accuracy incentive targets

- Targets should be set on a scientific basis, rather than in an ad hoc manner.
- Using raw historic forecasting accuracy data in an environment in which, as a result of government policies, uncertainties are increasing with a consequent worsening of forecast accuracy is not reasonable
    - Requiring the SO to do better in the future in a significantly worsening forecasting landscape than it has ever been able to achieve historically, when uncertainties were less, is not 'challenging'. It is impossible.
- The SO and the Regulator should agree on a scientific model that reflects the expected increase in uncertainties, and sets targets in the light of this model.

Bias incentive

- We do not believe that the additional incentive to incentivise under or over-forecast is required as the objective is already facilitated by wider licence conditions. However we should like to comment on the specific proposals, and suggest alternatives should the decision be to retain this part of the incentive.
- The SO has traditionally focused on its in-day and day-ahead demand forecasts, and given less detailed attention to its week ahead forecasts
    - This is partly as a result of there being less information available at the week ahead stage on which to base more detailed forecasts
- In recent years the demand supplied by the transmission system has been falling on an annual basis. The current forecasting practice has been to blend current underlying demand levels into underlying levels from the previous year over the course of the 11 week ahead forecasting time horizon
    - This leads to a structural bias in the week ahead forecasts
    - This has been partly driven by the forecasters feel that being short of electricity is more concerning than being long
    - This does not appear to be a requirement of either the market or the ENCC System Operator control room
- Now that this has been raised as an issue, as a responsible operator, the SO will adapt its forecasting procedure to correct this, regardless of whether or not the SO is incentivised to do this.
- The proposed scheme is split into two parts for calculation purposes:
    - One part looks at the bias in all CPs at a given time horizon within a month
    - The second part effectively looks at the worst performing CP forecast at that time horizon (in terms of bias)

- If the worst performing CP hits the trigger level of 70% bias, the bias scheme for that time horizon is set at the collar for that month
- Such a scheme is not fit for purpose:
  - The proposed scheme does not take into account the number of CP forecasts for a particular CP. In the shoulder months a CP may only exist for one or two days, and will inevitably be biased
  - Even if forecasting were as good as it possibly could be (with the proviso that consistent 0MW error is impossible) then this part of the scheme means that the bias incentive hits the collar 43% of the time
  - A forecasting bias scheme that even under perfect conditions means the SO is losing the entire incentive value 43% of the time simply because of random fluctuations incentivises no-one
- Forecasting practice is to have an intuitive understanding of a day's, or at relevant time scales, a week's underlying demand level. It is unusual for any single CP to perform worse in terms of bias than any other CP
  - If a forecast has a bias in one CP at a given time horizon it is likely to have similar bias in other CPs
- Consequently, we suggest that the second part of the scheme is not needed. In any case it cannot be used in its current proposed form
- A realistic bias scheme for all CPs would recognise that a certain range of random fluctuation is inevitable.
- Given that bias in CPs is highly correlated within a target day, there are of the order of thirty independent bias measurements a month. Small fluctuations in these 30 or so measurements should not dominate the calculation.
- We suggest that 'small fluctuations' should be seen in the context of metering error. Whether operational or settlement metering is concerned, metering errors in the short term are of the order of +/- 300MW. It is unreasonable to incentivise the SO based on the basis of fluctuations in metering error
- The proposed bias scheme is sensitive to very small changes in demand data. There should be a deadband representing the inaccuracy of demand metering, giving a quantitative value to 'small changes', the size of which would be different for different forecasting time horizons, as what would count as a small fluctuation would depend on how far ahead the forecast has been made.

  We propose:

  - +/- 100MW is suitable for day ahead forecasts
  - +/- 150MW at 2 day ahead; and
  - +/-300MW at week ahead

  Bias should only contribute to the performance if it lies outside this deadband.
- This still exposes the SO to the risk that atmospheric conditions may introduce bias into weather forecasts for a prolonged period of days, or even weeks
  - Current bias in a forecast does not have a persistence property
  - It is still the case that, even observing a bias for a number of days in a weather forecast, the best estimate for the following days weather is the mean of the forecast
  - Weather forecast providers adjust any bias that they can explain through manual processes – all information has already been used