National Grid's Response to Ofgem's Consultation Document:

Electricity Capacity Assessment: measuring and modelling the risk of supply shortfalls (Ref 132/11)

National Grid owns and operates the high voltage electricity transmission system in England and Wales and, as National Electricity Transmission System Operator (NETSO); we operate the Scottish high voltage transmission system. National Grid also owns and operates the gas transmission system in Great Britain and the distribution system in the heart of England serving approximately eleven million offices, schools and homes. In addition, National Grid owns and operates significant electricity and gas assets in the US.

Key Messages:

Within our response we have included a detailed statement of options for the treatment of wind and demand. However, we note that it is impossible to set out a definitive recommendation as to the appropriate statistical approach until we have access to the wind resource dataset for the 2011/12 study. In addition optimal methodologies may also evolve in future years as more data becomes available, and as further research results on relevant statistical methodologies become available.

While we recognise that a de-rated margin can offer some guidance on security of supply, in future systems with high renewables penetrations, the primary measure of risk will be indices such as LOLE (Loss of Load Expectation) and EEU (Expected Energy Unserved). LOLE is the most commonly used index internationally in adequacy studies while other more detailed metrics include frequency and duration indices (e.g. estimating the return time of events of a given severity). However, for the latter the amount of data required far exceeds that available so we will consider a number of "added value" metrics which can be calculated relatively easily, and can be informative when explaining the risks involved. An example would be the capacity value of wind generation or imports from interconnectors which would be required to ensure a given calculated risk level.

Wind data for the adequacy study will be based on historic meteorological records, converted to wind power output. It is of critical importance that data from the last two winters is included in the study; during these, extreme demands were coincident with low transmission-metered wind outputs (although whether this implied a poor wind resource across the whole of GB and relevant offshore waters remains uncertain). We are currently in negotiation with parties to procure sources of such data which if successful will enable us to determine the appropriate methodology.

The "Base Case" generation background to be utilised in the modelling has gone through an extensive stakeholder consultation via the Offshore Development Information Statement and the Transporting Britain's Energy consultations. Hence National Grid's data on transmission connected generation is fit for purpose. However, time series data for embedded generation is very limited and consequently raw demand data with embedded generation added back on cannot be reconstructed with exception of embedded wind. Consequently, embedded generation (with the exception of wind) will be accommodated by utilising Transmission level demand which is net of embedded generation. This lack of reliable data will be addressed in future studies by requesting data from DNOs.

The Cheviot boundary currently accounts for two thirds of all constraint payments making it essential to model. Over time other boundaries will become important as new generation is connected and should then be incorporated but until such time a two area system will suffice.

Contact: Duncan Rimmer, Energy Scenarios Manager,
duncan.rimmer@uk.ngrid.com,

Date of Response: 7th December 2011

Responses to Specific Questions:

The attached Appendix contains a single coherent statement of our proposed treatment of wind and demand rather than spreading it across various responses to different questions; however, we will include some comments within our responses to individual questions where appropriate and making reference to the appendix when more detail is required.

Chapter 3. Measuring the risk of supply shortfalls

Question 1: Do you agree that the de-rated capacity margin is a good indicator of future capacity adequacy?

For clarity, we define de-rated margin as the margin, over a given measure of underlying peak demand measure (e.g. ACS peak or 1 in 20), of the installed capacity scaled by an appropriate de-rating factor (mean availability for conventional plant, a capacity credit for variable generation such as wind).

This has proved to be a good adequacy indicator for a mostly conventional system, as demonstrated by its use in the Winter Outlook Report. However, looking forward to future systems, de-rated margin will become a less complete measure for two reasons.

Firstly, if the overall reliability of the generation fleet changes, then a de-rated margin will only account for the change in mean of the distribution of available capacity, and not the change in its width (e.g. if the profile of the conventional fleet changes so that average percentage availabilities increase, then not only is a reduced total nameplate capacity needed, but the required de-rated margin would also decrease due to decrease in the width of the distribution of available capacity.)

Secondly, the only systematic way of including wind generation in an assessment of de-rated margin is via a risk-based capacity value. As this capacity value assessment would require a full calculation of LOLE (or another index), then any statement of de-rated margin would in reality be this risk calculation result converted to a MW scale by inverting the relevant cumulative distribution function.

We therefore conclude that de-rated margin can offer some guidance, but that particularly in future systems with high renewables penetrations, the primary measure of risk will be indices such as LOLE and EEU.

Question 2: Are there any measures of risk other than LOLE and EEU that we should report and what are their comparative advantages?

Commonly used planning timescale measures of risk divide into expected value indices such as LOLE and EEU, and frequency and duration indices.

LOLE is the most commonly used index internationally in adequacy studies (e.g. the relevant IEEE Working Group is actually called the “LOLE WG”). Calculation methods for EEU are very similar; it is sometimes thought of as a more detailed index as it gives some consideration to severity of shortage events.

Within the general framework of LOLE/EEU, a number of “added value” metrics can be calculated relatively easily, and can be informative when explaining the risks involved. An example would be the capacity value of wind generation or imports from interconnection which would be required to ensure a given calculated risk level.

Other more detailed metrics include frequency and duration indices (e.g. estimating the return time of events of a given severity). Given an infinite amount of data we would look to calculate these. However, as discussed in our methodology appendix, due to limited experience of very high demands, there is great uncertainty associated with estimates of even the simplest indices such as LOLE. As a consequence, we are very extremely doubtful as to whether it is possible to make meaningful estimates of frequency and duration indices which depend sensitively on the range of possible peak period load shapes (and the wind profiles which might be coincident with these).

We further note that LOLE and EEU only give an expected value (in the mathematical sense) estimate for what might occur in the future scenario under study; they give no indication of the variability of outcomes about this central figure due to statistical fluctuations (which would be present even given perfect data). For instance, in a viable system under present planning and operation philosophies, in many years there would be no adequacy-related demand reduction at all. In principle, again given an infinite volume of data, one could calculate a probability distribution of the number of periods of shortage, or the unserved energy. However, once more we doubt that the volume of available data will support this beyond providing a list of possibilities without probabilities attached.

Finally, we note briefly that, given the relevant probability distributions, expected value indices do not depend on serial associations in the various time series for demand, wind and available conventional plant. On the other hand, frequency and duration indices do depend on serial associations, as does the sampling uncertainty in expected value indices. This issue is discussed further in the methodology appendix.

Chapter 4. Modelling and data requirements

Question 3: Are there any additional key input assumptions that we should consider in the modelling?

There are a number of key input assumptions most notably demand data, embedded generation data, “base case” generation background and wind speed/output data.

Dealing with these in turn:

- A time series of transmission metered demand data is available on a consistent GB basis since BETTA was introduced in 2005. Prior to that data for E&W and Scotland from the NETA era would require some minor adjustments but this would be a minor uncertainty in any calculations. Prior to NETA (i.e. pre 2001) the data would be less reliable, and in any case may well be irrelevant due to different underlying demand patterns (different residential/commercial/industrial split, less electronic load further back etc). We also note that adding more years of data might not add what is really needed to improve calculations, i.e. additional experience of truly extreme demand periods. To summarise, more relevant data is always helpful in statistical estimation; however low quality, irrelevant or unrepresentative data is of little benefit. A key judgement in the adequacy assessment project will thus be how far back aggregate transmission metered demand remains relevant for the study (2001 to the end of the study period in 2015/16 covers 15 years of technological and economic change); we may be able to draw on the experiences of our short term and longer term forecasting processes in making this judgement. Finally, while a longer historic time series may be of very limited benefit for quantitative risk calculations, it might e.g. provide valuable qualitative information on typical wind conditions at times of extreme demand.
- To enable the use of Embedded or Distribution connected small and medium generators output data, a complete time series is required. Unfortunately, on investigation this is not currently available; consequently, its inclusion will have to be indirectly via modelling Transmission demand which is net of embedded generation i.e. embedded generation moves from the generation side of the equation to the demand side of the equation. Note that for Embedded Wind we are moving it to the generation side (under the assumption that it will generate if the wind blows). In principle embedded solar and tidal can be treated the same way (assuming that we have a per-location time series of resource availability) – but as the capacity of these other renewable sources is very low (<100MW capacity) and in the case of solar won’t be generating at times of high demand, we can ignore them for the purposes of this assessment. The lack of reliable data could be potentially addressed in future studies if a time series becomes

available through data obtained from DNOs. However, as the forecast of Transmission demand in the "Base Case", which would be used to scale historical demands, assumes some growth in non-wind embedded generation it would be reasonable to assume the demand profile over the next 4 years of this new load would be similar to that of existing sites. The "Base Case" assumes just over 1GW on new non-wind embedded generation will connect over this time period.

- The "Base Case" generation background assumes a number of changes over the next four years including station closures, e.g. LCPD opted out and nuclear magnox, new CCGTs and wind farms along with the mothballing of plant. However, since this "Base Case" was developed there have been a small number of market changes e.g. Tilbury converting to biomass from coal. Consequently, it would be prudent to incorporate these known changes by adjusting the "Base Case" but to either leave the other less definitive uncertainties as sensitivities/scenarios or investigate the implications of moving the generation plant distribution either way by 1 or 2GWs.
- Wind data for the adequacy study will ideally be based on historic meteorological records, converted to wind power output in collaboration with our short term forecasting team. It is of critical importance that data from the last two winters is included in the study; during these, extreme demands were coincident with low transmission-metered wind outputs (although whether this implied a poor wind resource across the whole of GB and relevant offshore waters remains uncertain due to the limited geographical extent of current transmission metered wind). We are currently investigating sources of such data and have received "ball park" quotes from the Met Office and Meteo Group. If for some unforeseen reason these prove inappropriate then a fall back position would to approach Poyry Consulting who could provide aggregated wind power output data given a set of capacity assumptions; however, the potential lack of a complete audit trial for this approach is less desirable. It may also be possible to contract an experienced meteorological group in academia to produce a wind resource database for us using publicly available worldwide "reanalysis" datasets. We would need to explore relative costs, quality and delivery times of all options to ensure best value.

An additional input variable that could be considered is the potential sterilisation of capacity behind a network constraint e.g. thermal or wind stations in Scotland. However, even before further reinforcement of the Cheviot boundary, at times of high demand this sterilisation would most likely be associated with high wind generation in Scotland. As this would commonly mean at least some wind generation availability in E&W also, we do not anticipate that the effect of including network constraints to be very great, particularly as installed wind capacities increase further.

Question 4: Do you agree that the use of stochastics (probability distributions) to model short-term variation of key input variables is the best available method? Do you agree with the use of scenarios and stress tests for capturing long term uncertainty in key input variables?

Stochastic Modelling:
Demand, wind and available conventional plant (given a scenario of installed capacity) should all be treated probabilistically; each of these exhibits short term variability, and there is appropriate data available to attempt estimation of the necessary probability distributions. We discuss each of these briefly here; a more detailed presentation of our proposed approach may be found in the methodology appendix.

- ***Demand and wind.*** These are discussed in more detail in Question 5 and the methodology appendix. We note here however that it is not possible to give a definitive methodology until we have obtained the wind data which will be used for this study, and analysed what statistical estimation approaches this data can support.
- ***Conventional generation.***

- Given a scenario of installed capacity, we will build a distribution of available capacity using a standard capacity outage table approach (i.e. convolving distributions for each individual unit to give a distribution for the aggregate).
- Uncertainty over what conventional plant will be installed will be treated through scenario analysis. In principle it would be possible to assign subjective probabilities to different scenarios based on expert knowledge, however any choice of probabilities would be open to challenge, and we believe that this approach would be less transparent in interpretation than the scenario approach.

Questions of demand versus calendar ordering will be discussed further in the methodology appendix.

Scenarios and stress test:
We believe that scenarios and stress tests provide good pragmatic means of studying the sensitivity of security of supply to a number of uncertainties e.g. interconnector behaviour. Any attempt to model interconnectors "bottom up" via risk or economic models of the relevant systems would reduce to the question "what is the probability that when the GB system is in trouble the French / Dutch / Belgian / Irish systems would be in trouble as well?" As the uncertainty surrounding this question is so great, we believe that such bottom up modelling will not be productive. A further option is once more to use subjective probabilities to quantify beliefs; however, for the same reason as for installed capacity our view is that this will not be appropriate. For a more detailed explanation of this please refer to our response to Question 9.

There are various uncertainties which we plan to treat through scenario analysis, and whose effect (to leading order) is to shift the distributions of supply and demand relative to each other; a further example is demand growth (which might be based on demand scenarios published in the Statutory Security of Supply Report). We will therefore explore the possibility of treating multiple scenario uncertainties in a unified way by examining the sensitivity of model results to shifts in the distributions.

Question 5: Do you agree with the proposed approach to modelling wind availability?

In response to the consultation document paragraphs 4.19 and 4.20, we add that we will consider the use of both historic metered data and turbine specifications in the conversion of wind speed data to wind farm power time series. In particular, we will draw on the expertise and access to data used by our short term forecasting team. Typically 'wind farm power curves' look rather different from individual turbine curves, and hence access to metered historic data is invaluable when making this conversion. However, we acknowledge that there is very limited direct historic experience of offshore wind farms.

Our planned approach to deriving probability distributions of available wind capacity is described in detail in the methodology appendix. We note here that the direct historic experience of high demands combined with poor wind resource (the critical regime for adequacy in a system with high wind penetration) is very limited. Moreover, even if the relevant data were available, techniques for assessing any statistical relationship between demand and available wind capacity are still at a fundamental research stage (if indeed they have been developed at all). Preliminary uncertainty assessments of the commonly used hindcast approach (i.e. using directly in the risk calculation the bivariate historic time series of demand and wind resource) indicate that with the available data, it will be inappropriate for this study due to huge sampling uncertainty; the results are dominated by data from a small number of distinct historic periods.

An alternative is to estimate a probability distribution of available wind capacity based on historic records which are deemed relevant to times of high demand; this may consist (e.g.) of winter afternoons/evenings, all winter data, or data from times when demand exceeded a certain level. Within the risk calculation, the available wind capacity would then be assumed independent of demand. This would account for seasonal variation in the wind resource, and to an extent for any diurnal variation. We accept that ideally the calculation would account more explicitly for any wind-demand relationship; however, as stated above statistical methodologies for doing so would have to be developed from scratch for this project.

However, we note once more that it is impossible to make a definitive recommendation as to the appropriate statistical methodology for this study until we have access to the final wind dataset. Further to this, due to limited time for this initial project we anticipate that the methodology for treating wind and demand will evolve in future years in response to new research results on statistical approaches for adequacy assessment.

We will treat transmission-connected and embedded wind on a common basis, as embedded wind clearly behaves more like transmission connected wind than other components of transmission demand. Moreover, historic data on installed capacities and wind speeds allows us (unlike other important embedded generation technologies) to reconstruct an historic time series of embedded wind power output.

Question 6: Do you agree with the proposed use of NGET's existing data and assumptions, regarding, in particular, commissioning and decommissioning dates and embedded generation?

The "Base Case" generation background to be utilised in the modelling has gone through an extensive stakeholder consultation via the Offshore Development Information Statement and the Transporting Britain's Energy consultations. Details of this scenario and others were also published in November 2011 in National Grid's "UK Energy Scenarios" document NG_Futureenergyscenarios.

Hence National Grid's data on transmission connected generation is fit for purpose. However, time series data for embedded generation is very limited and consequently raw demand data with embedded generation added back on cannot be reconstructed with exception of embedded wind. This lack of reliable data could be potentially addressed in future studies if a time series becomes available through data obtained from DNOs.

Question 7: Do you believe that Ofgem should require industry stakeholders to submit up-to-date data with regard to commissioning and decommissioning dates and embedded generation? Which industry process will ensure the confidentiality of information provided?

National Grid through its industry consultations and contract negotiations around new connections has a good source of new station information; however, information about decommissioning or TEC reductions is much more limited due to a combination of Grid Code obligations only giving 6-12 months notice and that station operators themselves are unsure about closure dates beyond the immediate time horizon. So it is unclear whether any more information would provide a greater insight as various stakeholders commented on during the 5th September workshop. However, if such information is forthcoming then we would endeavour to incorporate it within the analysis either by adjusting the "Base Case" or through scenarios.

With regards to embedded generation data is currently limited to annual and peak demands with little or no time series data. Whilst this shortfall can be covered by using transmission demand which is net of embedded generation it would be worth asking stakeholders e.g. DNOs what data could be made available either via a Grid Code modification or more realistically via a direct request.

Question 8: What are your views on how best to model LCPD opt-out plants' restricted running regimes?

LCPD opted-out plants have two important elements to include in the modelling; firstly, their mode of operation and secondly, their date of closure.

Since LCPD came in place in January 2008, these plants have operated around the time of high demand and high electricity prices and we would expect no change in this operation into the future. Analysis has been carried out to compare these plants against opted-in coal-fired generation; which suggests that they have different behaviours during periods of high demand as well as across the year. Load factors are 10 – 15% lower for opted-out plants and therefore any modelling of these plants should take account of this difference to opted-in plants.

The other main uncertainty surrounds the closure dates. We have produced analysis to estimate the plants' closure dates based on their operating patterns since January 2008 (start of LCPD). The "Base Case" would assume closures dates based on existing operation patterns and generators portfolio developments i.e. their overall demand and generation balance. Variation in these closure dates could be addressed through either sensitivity analysis e.g. LOLE increases by X if LCPD plants close early or by examining the sensitivity of model results to shifts in the distributions.

Question 9: Which of the two approaches for modelling electricity interconnection flows will provide the most realistic flows? If you favour the scenario based approach, what are your views on reasonable scenarios to run?

Endeavouring to model flows across interconnectors around the time of system stress is essentially about estimating the probability that other systems will be in trouble at the same time as GB. Whilst interconnector flows have in the past correlated well with price differentials there is huge uncertainty around future energy prices. Consequently, it is virtually impossible to forecast flows accurately at times of system stress due to a combination of the following four effects:

- For a capacity assessment, we are only interested in the interconnector flows at times of system stress – i.e. when demand is high and generation is limited.
- As the demand in Europe is correlated with demand in GB, and wind power output is similarly correlated: the interconnectors will not necessarily flow towards GB. (This is because weather systems are about the size of continents, and therefore GB and Europe often have similar short term weather effects).
- The third energy package will cause interconnector flows to respond more quickly to price signals (and therefore shocks to either market). The third package also removed use of system charges fro exports thus changing the market and invalidating the reliability of using historical data.
- Finally, unexpected policy decisions can have dramatic effect on capacity and therefore force in the interconnector flows e.g. changes in Government policy similar to that of nuclear power in Germany

Hence we believe a scenario based approach for interconnectors is the best way forward covering the full range of import/export, float and halfway between. Although we will keep the option open of using the functionality within Plexos to determine potential flows at this stage we aren't planning to utilise it.

Question 10: Under what conditions would users respond by curtailing their demand and how would you go about modelling this? Is it worth Ofgem requesting data from DNOs on self-interruption and interruptible contracts?

Currently we see around 500MW of triad avoidance demand side response; however, very little of this is officially notified to us as per Grid Code. Consequently, while we may have reasonably accurate data around the peak demand periods; for other high demand periods very little data is available. We also have some contracted demand side response via STOR (Short Term Operating Reserve).

In terms of STOR, we currently have accepted contracts for ~1500MW of Non BM STOR, which includes units that are planned or under construction. For the current season, we are seeing 700-750MW of that 1500MW available at any one time.

The majority of the Non BM STOR is provided by demand side generation, with only ~100MW (variable) from true demand reduction services, there is a further 120MW provided by aggregator sites that have some demand reduction mixed with generation in their aggregated portfolio.

National Grid uses STOR in operational timescales economically i.e. National Grid will choose not to synchronise a conventional unit for reserve because STOR is available. This has the effect of slightly reducing the reserve utilisation as the decision to synchronise can usually be deferred.

A second way National Grid has visibility of demand side response is via FCDM (Frequency Control by Demand Management) and the FFR (Firm Frequency Response) services, in combination these services provide 200-250MW of demand reduction. These services only deliver when the frequency has already moved away from 50Hz, therefore they should not be included as possible demand side response for the purposes of a capacity assessment.

Question 11: Is historical data of scheduled outages a good indicator of future patterns of scheduled maintenance timings?

Historical data gives useful information on the broad picture of maintenance scheduling under current generation running regimes. Looking at this might however be misleading, as operating patterns change over time with the increasing penetration of wind generation. We anticipate however that it will be some years before resulting changes in maintenance patterns have a substantial influence on adequacy risk.

Including maintenance in risk calculations in a meaningful way is very challenging indeed, due to the possibility of flexing maintenance schedules on an operating timescale if margins are looking tight in the short term. This breaks the usual assumption of independence within these calculations between available conventional capacity and all else. To our knowledge, where explicit consideration of maintenance schedules has been made within probabilistic risk models used in existing applied studies, this is based on unrealistic rigid maintenance schedules. As this will seriously overestimate the adequacy risk, we propose instead to make a qualitative assessment of any issues surrounding maintenance, based on maintenance requirements of different technologies, the length of the maintenance season (essentially British Summer Time) and de-rated margins. Particular attention will be paid to maintenance over-running into the start of winter time, when the early evening darkness peak first appears each autumn.

Question 12: Will treating half-hour periods independently have significant effects on our estimates of the de-rated capacity margin and risk of supply shortfalls and how should the model take into account half-hourly cross-correlations?

Whether one needs to account explicitly for serial statistical associations in generation availability and demand depends on the index being calculated. Expected value indices such as LOLE or EEU do not depend on such serial associations, as the expectation value of a sum of random variables is the sum of the individual expectation values, whether or not the variables are independent.

However, frequency and duration indices do clearly depend on serial associations, as does the variability of out-turns about the central expected value estimate. As discussed elsewhere, our current best assessment is that there is so much uncertainty over estimates of expected value indices (such as LOLE) that attempts to estimate more detailed indices will not be worthwhile. This assessment will be reviewed as we have access to our final dataset.

The uncertainty in estimates of LOLE and EEU do depend on serial associations in generation and demand, and we will assess this effect by bootstrap resampling as described in the methodology appendix.

Question 13: Are there any boundaries other than Cheviot that may significantly affect the risk of supply shortfalls?

The Cheviot boundary currently accounts for two thirds of all constraint payments making it essential to model. Over time other boundaries will become important as new generation is connected, particularly with Connect & Manage (C&M), and should then be incorporated in any future modelling. There are currently a number of other boundaries included in Plexos which can be monitored on an ongoing basis as part of the C&M governance, keeping track of the expenditure due to each of the SYS boundaries. However, extra boundaries in the model (i.e. more areas in the capacity assessment) have the effect of increasing the amount of data required to perform the assessment. Hence finding the correct balance between detail and minimising the errors due to inaccurate estimates of boundary limits will be a hard task to parameterise. Consequently care would need to be taken when assessing what trigger levels should be employed in deciding when to include additional boundaries.

Methodology Appendix:
- National Grid's proposed treatment of wind and demand

We make here a single detailed statement of options for the treatment of wind and demand; spreading this across the various related questions above would result in a less coherent description. We note once more that it is impossible to set out a definitive recommendation as to the appropriate statistical approach until we have access to the wind resource dataset for the 2011/12 study. The optimal methodologies may also evolve in future years as more data becomes available, and as further research results on relevant statistical methodologies become available.

Qualitative features of historic datasets and statistical estimation methods
We first note several qualitative features of the available historic datasets (or those which could conceivably become available).

Firstly, there is very limited historical experience of the extreme demand periods which dominate probabilistic generation adequacy assessments in GB. Indeed, from the last 10 years (for which we have detailed transmission demand data for the system) there have been just two distinct truly extreme periods (in the last two winters), which cannot supply more than two fully independent items of data. Indeed there is doubt as to whether data from consecutive winters may be regarded as independent. Even assuming that weather in separate winters is independent and that underlying demand patterns do not change (both of which would be false assumptions), for extreme demand events we have a list of possibilities, rather than the ability robustly to assign probabilities or estimate return times for events of a given severity.

Secondly, estimating the distribution of demand for a future winter is a simpler task than quantifying any statistical association between the demand level and available wind resource at times of extreme demand. Once more, the historic record supplies a list of possibilities, rather than an ability to assign probabilities without making further assumptions.

Thirdly, given these substantial data limitations, any estimates of probability distributions or risk indices must be associated with an estimate of uncertainty in those quantities; a central estimate taken in isolation will be essentially meaningless. Importantly, when calculating expected value indices such as LOLE and EEU *given the necessary probability distributions*, it is not necessary to consider serial associations in the relevant time series (as the expected value – in the mathematical sense - of the sum of two random variables is the sum of their individual expected values). However, when assessing uncertainty in these estimates, consideration of relevant serial associations is vital.

Finally, historic weather records exist over many decades. This may provide assistance in understanding qualitatively the types of weather systems and quality of wind resource typically associated with extreme demand conditions. However, the reality of a non-stationary climate and underlying demand patterns means that historic data beyond a certain age is not suitable for quantitative estimation of the distribution of demand, or of any wind-demand relationship – a longer historic record may however give valuable qualitative information on the typical wind resource at times of truly extreme demand.

Statistical approaches: investigating demand in isolation
Historic Load Duration Curve (LDC)
The simplest option is to use the historic empirical load duration curve as the distribution of demand. For comparability between winters, demands would be rescaled according to historic out-turn and future projected ACS peak levels. In a system without wind generation, the estimate of LOLE would then be

$$I^{\mathrm{LOLE}} = \frac{1}{n}\sum_{t} P(X \le d_t)\,, \qquad (1),$$

where d_t is the (rescaled) demand at historic time t (a known data item), X is the available conventional capacity at a particular time in the peak season in the future scenario under study (a random variable), n is the number of years of historic data used, and the sum is over the historic data. Depending on the resolution of the time series, this LOLE estimate might be in settlement periods per year, hours per year, etc.

Using this approach, the question of whether historic weeks are calendar-ordered or demand-ordered is not relevant. The question of ordering of weeks only becomes meaningful if one wishes to consider separately the peak week, second peak week, etc as PJM does. We expect available data to be too limited to make meaningful separate considerations of different weeks in a year, particularly due to statistical associations within and between weeks of a given historic year.

Time series modelling
A further alternative would be to build a statistical time series model for demand. However, we are not aware of any such time series techniques which are appropriate for adequacy (where the model must place particular focus on extremes of demand). Given the complexity of such a modelling task if all relevant daily/weekly/seasonal cycles are to be considered, we do not propose to attempt the development of such techniques ourselves on the timescale of this project.

Parametric estimation
A further question is whether some form of parametric distribution estimation should be performed, as in the PJM approach. There is certainly no fundamental reason to assume any particular distribution family for this purpose. However, where statistical smoothing is required one could potentially use specialised extreme value statistical techniques for the distribution tail, and general non-parametric methods elsewhere.

Uncertainty quantification
Whatever statistical approach is used, some assessment of uncertainty in modelling results is required. Given that estimation approaches are likely to be non-parametric, an appropriate method for quantifying sampling uncertainty is bootstrapping. This attempts to answer the question “how much could the estimate have varied given a different n years of data” by creating a large number of further data samples of the same size by through resampling (with replacement) from the original time series. The distribution of calculation results obtained using these samples then provides the uncertainty assessment.

This quantitative approach should also be complemented by intuitive qualitative assessments such as examining which parts of the historic data is driving the final calculation results. If data from a small number of distinct periods dominates the results, then without the need for any formal uncertainty analysis techniques it is clear that those calculation results cannot be at all robust.

Inclusion of wind generation
Hindcast
Hindcast extends the ‘LDC’ approach for demand by using directly coincident historic wind resource data. The LOLE estimate is then:

$$I^{\mathrm{LOLE}} = \frac{1}{n}\sum_t P(X + w_t \le d_t)\,, \tag{2}$$

where w_t is an assessment of what the available wind capacity would have been at historic time t, given the scenario of installed capacity under study. Hindcast is a popular approach in the literature, as it attempts to account for statistical association between wind and demand without requiring any advanced statistical technology. However, our initial assessment (including sampling uncertainty assessment by bootstrapping) is that there is insufficient historic data available for this approach to deliver meaningful results. This effect will be magnified at very high installed wind capacities, hence even if it were deemed to be suitable for a 2011/12 scenario hindcast is not in any way future-proofed. Indeed, we expect that the additional volume of data required to reduce the sampling uncertainty to acceptable levels would be so great as to be unobtainable even under strong assumptions about stationary demand patterns and climate.

Time series approaches
Some form of statistical smoothing is therefore required. As discussed above, we are not aware of appropriate non-parametric time series modelling approaches for treating demand for adequacy purposes, never mind a bivariate time series modelling considering wind and demand together (along with their statistical association).

Static estimate of joint distribution
We are also not aware of appropriate statistical methods which have been developed for estimating a static joint distribution of demand and available wind capacity suitable for use in calculating expected value indices such as LOLE. However, we will investigate whether it is possible to develop such approaches in the time available for the 2012 report. A further option would be to take a hindcast time series of net-of-wind demand, and apply statistical smoothing using univariate extreme value methods in the tail and non-parametric methods elsewhere.

Separate wind and demand distributions
The final candidate approach is to derive a distribution of available wind capacity based on a subset of records deemed representative of times when the system when the system might be under stress (e.g. all winter records, or all winter afternoon/evening records). This will give much reduced sampling uncertainty, in exchange for the modelling assumption that all historic records used are equivalent. Within the probabilistic model, this wind distribution would then be considered to be statistically independent of demand. There would remain a question as to what degree of statistical smoothing should be applied to the separate wind and demand distributions.

Conclusion
We have described here a number of candidate approaches for the probabilistic modelling required for the capacity adequacy assessment. We emphasise once more that we make no final recommendation of an approach, as we have not yet procured the wind resource dataset for this study, and believe strongly that the choice of statistical approach must necessarily be driven by the available data. It is very possible that the statistical methodology used may evolve in subsequent adequacy reports from 2013, as more relevant research results become available.

We suggest that quantification of any statistical association between demand and available wind capacity (and indeed to what extent the available data can enable this) is an area which is neglected in the literature, and where further research is required. We would be interested to hear through this consultation of any techniques from the literature of which we are unaware.