Standard Approaches to Load Forecasting and Review of Manitoba Hydro Load Forecast for Needs For and Alternatives To (NFAT)

Wayne Simpson

University of Manitoba

Douglas Gotham

Purdue University

TABLE OF CONTENTS

INTRODUCTION

The Manitoba Hydro report on Needs for and Alternatives to Business Case (NFAT, August, 2013) provides load forecasts to 2032 that constitute the starting point for their assessment of alternative development plans. Elenchus Research Associates have now provided a review of Manitoba Hydro's load forecast for the Public Utilities Board (Elenchus, January, 2014) that describes the forecasting process in detail and provides an assessment of its accuracy and reliability. In summary, Elenchus (iv) finds that the NFAT load forecast is reasonable but deficient in terms of the alternative economic and population scenarios considered and the methodology used to assess its reliability. This report attempts to avoid duplication with the work done by Elenchus to provide additional discussion of the accuracy and reliability of Hydro's load forecasts in the NFAT. It is meant to be read as a companion to the Elenchus report. In particular, this is not an attempt to repeat the detailed description of Hydro's load forecast methodology already provided by Elenchus.

This report is structured in two parts. Part 1 provides a general overview of load forecasting techniques and the significance of forecast accuracy for resource planning purposes. Part 2 looks specifically at Manitoba Hydro's load forecasting in the NFAT.

PART 1 - Standard Approaches to Load Forecasting

A number of different approaches to load forecasting exist, each with its own advantages and disadvantages. In general, load forecasting methods can be classified as bottom-up, top-down, or a hybrid of those two. Bottom-up forecasting involves producing projections at the individual customer or device level and summing across the various customers and/or devices. Top-down forecasts are produced by aggregating the customers into larger groups and projecting usage at the group level. Hybrid approaches use features of both bottom-up and top-down methods.

In general, forecasting models use what is known about the past to predict what will happen in the future. The level of detail and sophistication of the model may vary considerably, as can the model's ability to capture fundamental changes in the future. In many cases, different forms of forecasting models will be used in conjunction. For example, an estimate for one of the drivers of an econometric model may be derived using a regression model for that driver.

While there are a number of approaches to load forecasting, they are not all equally appropriate. In its whitepaper on forecasting methodologies, the Mid-continent Independent System Operator (MISO) identifies the qualities of a good forecasting system as understandability, credibility, accuracy, reasonable cost, maintainability, and adaptability. MISO also provides lists of acceptable and unacceptable forecasting methods MISO, ("Peak Forecasting Methodology Review," 2013-12-06, https://www.misoenergy.org/_layouts/miso/ecm/redirect.aspx?id=98923). The inclusion of different methods on those lists is provided with the description of the approaches.

Top-down forecasting

Trend analysis

Load forecasting using trend analysis (also referred to as time series or regression analysis) relies solely on the historical load with no consideration of the factors that affected the amount of energy used. In essence, regression models determine a mathematical equation that explains historical usage and extrapolated to future usage using that equation. Perhaps the simplest form of this model is to assume the future value will stay at the historical average. The most common form of a trend model is a linear trend. In a linear trend, the historical data is fit to a straight line (as best as possible). The slope of the line then provides the change in value from one period to the next in the future. The line fit is usually determined using the ordinary least squares method.[1] It is also possible to use non-linear regressions (such as polynomial[2] or exponential[3]) in a trend model.

The major advantage of trend analysis is simplicity. It requires no data beyond the historical observations of the value that is being forecasted and the regression can be calculated using the statistical functions of commercial spreadsheet software.

The major disadvantage of trend analysis is inaccuracy. Trend models do not account for changes in the economic, climatic, and demographic factors that may change energy use. It may not be possible to obtain a regression with a good fit to the historical data, particularly if there is a lot of variability in the data. Furthermore, the choice of historical data can influence the results. For instance, a load forecast based on the past five years, which saw a significant economic recession, would produce a very different result than one based on the last twenty years.

Load forecasts based on regression models have been largely discredited. In its forecasting review whitepaper, MISO states that “any statistical extrapolation of historical trend using only data from the series to be forecast is unacceptable as the primary forecasting technique.”

Econometric models

Econometric models attempt to quantify the relationship between the parameter that is being forecast (the output variable) and a number of factors that affect the output variable. These factors are commonly referred to as explanatory variables or drivers. Each explanatory variable affects the output variable in a different way. For instance, manufacturing output may be positively correlated with energy use in that they tend to go up and down together, while electricity price may be negatively correlated with energy use. The mathematical relationships (aka sensitivities) are determined

[1] The ordinary least squares method is used for estimating a line that is as close as possible to the historical data. For each time period for which the data is collected (for annual data this would be for each year), the difference between the historical value and the corresponding value on the line is determined. These differences are then squared and summed across all points. The line with the best fit is the one that has the lowest value for that sum.

[2] A polynomial function is a mathematical expression where the variables are raised to an integer power. The simplest form is when the variable is squared ($y=x^2$). An extension of the square function is the quadratic ($y=ax^2+bx+c$). Higher order polynomials include cube functions and beyond.

[3] An exponential function is one where the variable is an exponent of some constant ($y=a^x$). An exponential function will increase or decrease by a fixed percentage as opposed to a fixed value in a linear function.

simultaneously and can be calculated via any of the methods used in time series forecasting, such as linear, polynomial, and exponential. Thus, an equation is derived that includes the relationship of each driver to the output. Projections of the values for the drivers are then used to determine the output variable for each forecast period.

The appropriate explanatory variables may differ from one utility or region to another. They may also change over time as factors change. Common explanatory variables include population and demographics, employment, economic output, personal income, weather, and the price of electricity and competing energy sources.

Econometric forecasting is more time and resource intensive than trend analysis. In addition to the development of the model, it requires the acquisition or development of projections of the drivers. These projections may be produced in house, using another econometric model or a regression model, or they may be produced by commercial vendors or by government entities.

The major advantages of econometric forecasting are the potential for improved accuracy, the ability to analyze the impact of scenarios that are more optimistic or pessimistic, and a greater understanding of the factors that affect the forecast uncertainty.

The major disadvantage of econometric forecasts is that it is difficult to account for factors that will change the future relationship between the drivers and the output variable. A common example of this is changes in energy efficiency, either through utility demand-side management programs or through government codes and standards.

MISO includes econometric forecasting on its “Acceptable List” of forecasting methodologies.

Bottom-up forecasting

Survey-based forecasts

Survey-based (aka informed opinion) forecasts use information from a select group of customers regarding their future plans as the basis for the forecast. This is most commonly done with the largest consumers of energy, since those customers have the greatest impact on the forecast and are often a source of considerable uncertainty. Information is collected regarding the customers’ future production and expansion plans. Sources for this information can be from direct contact with the customer, public announcements, or discussions with an industry expert. The forecast then reflects the information regarding future plans. Thus, if a facility is expected to maintain current production levels, the forecast will indicate a constant load. Similarly, an increase in production or an addition of new facilities will result in a forecast load increase. Conversely, if a customer is expected to scale back production or close facilities, forecast load will drop.

The major advantage of survey-based forecasts is the ability to account for expected fundamental changes in customer demand for large users, especially in the near-term when customer plans are relatively firm. It may be difficult to capture these changes in an econometric model.

The major disadvantage of survey-based forecasting is the lack of information regarding customers' plans in the long-term. Most industrial facilities do not know what their production levels will be five or ten years into the future. Similarly, while some customers will cease operations in the long-term, very few are currently expecting to do so in the future. New facilities will likely be added in the long-term, but the forecaster will have no knowledge of them. Thus, survey-based forecasts are inherently inaccurate in the long-term. A second disadvantage of survey-based forecasts is the lack of transparency. Conversations between large customers and utility representatives are typically held in confidence.

Entities that rely on this type of forecasts will sometimes rely on it only for the early period of the forecast and use another method, such as econometrics, for the later period.

MISO includes informed opinion forecasting on its "Unacceptable List" of forecasting methodologies.

End-use models

End-use forecasts look at energy use at the individual device level. The consumption of energy is categorized into a number of different activities which provide a desired service or end use. Examples of these include lighting, refrigeration, space heating, and cooling. End use models start with a catalogue of the existing stock of devices for each end use. This includes the vintage, or age, fuel source, and efficiency of the devices. For each forecast period, the model assumes that some of the existing stock will fail, with failure rates being a function of the vintage of the device. When failure occurs, the device can either be repaired or replaced. Additionally, new devices will be added to the stock as the number of homes and businesses increase. In some cases, old (but still functioning) devices may be replaced by new ones as well. New devices, along with replacement of existing devices, are chosen from the available options. This provides a new "existing" stock to be used for the next forecast period. The forecast is then derived from the energy used by all of the devices in each forecast period.

The repair/replace and new purchase decisions are based on the purchase and operating costs of the available options for the end use, along with the customer payback period. Alternatively, a model may have a distribution of payback periods to reflect differences in consumer behavior. Thus, the model will choose between options with low purchase costs and those that are more efficient but cost more to buy. Also, end-use models can reflect the competition between different energy sources, such as electricity vs. natural gas.

The major advantage of end-use models is the ability to directly capture changes in efficiency, through both government codes and standards and incentive programs. In the case of a changing standard, such as the phase out of incandescent lamps, the model simply does not include the less efficient option as a possibility for new stock once the standard is in place. For incentive programs, the purchase price of the efficient option can be adjusted to reflect the rebate or tax exemption.

Disadvantages of end-use models include being very data-intensive, the potential to miss energy consumption from devices that have yet to be invented or adopted, and the inability to capture changes in customer behavior. It is not feasible to know the exact number of devices that are in use. Populating the initial stock of devices is usually done via customer surveys and care must be taken to ensure that the surveys are representative of the overall mix of customers. While most end-use models include a miscellaneous category for devices that either do not use much energy or are not widely used, over time new end uses evolve that are often not adequately captured. Early end use models did not include personal computers and other such devices. Most current end use models do not include electric vehicles, which could be a significant user of electricity in the long-term. Finally, end use models generally do not account for changes in customer behavior that may affect the amount that a device is used. Once they have installed a higher efficiency device, some customers may use the device in a different fashion than they used the old one. A customer may adjust the thermostat to a more comfortable setting with a high efficiency air conditioning or space heating system.

MISO includes end-use forecasting on its "Acceptable List" of forecasting methodologies.

Hybrid forecasting

Hybrid forecasting models employ facets of both top-down and bottom-up models. The most common of these is the statistically-adjusted end-use (SAE) model, which embeds econometric formulations within an overall end-use model. In reality, most end-use models are hybrid to some degree in that they rely on top-down approaches to determine the growth in new devices. Other types of hybrid models are possible, such as using a survey-based model for the short-term which feeds into a longer-term econometric or regression model.

In general, hybrid approaches attempt to combine the relative advantages of both model types. This usually comes at the cost of increased model complexity.

MISO includes hybrid forecasting on its "Acceptable List" of forecasting methodologies.

Forecast accuracy for resource planning

Regardless of the methodology used to develop the load forecast, having an accurate forecast is an important factor in resource planning. An inaccurate forecast can have significant reliability and cost implications. For instance, if the forecast is too low (load ends up being much higher than was predicted), the utility could end up having insufficient resources in the future. This may force the utility to rely on options that can be implemented with a short lead time (such as a market purchase) that could be more expensive than the options that could have been used if the forecast had been more accurate. Similarly, if the forecast is too low, the utility will acquire too many resources (and earlier than necessary). This also results in expenses that are higher than they would have been with an accurate forecast. While a perfectly accurate forecast is unattainable, it is important to avoid a forecasting methodology and assumptions that are likely to introduce a bias in either direction.

PART 2 - Report on Manitoba Hydro Load Forecasting in the NFAT

Context

Elenchus provides an extensive discussion of a variety of possible scenarios that could impact electricity demand, both domestic and imported, including the development of competitive alternative energy sources. The report, however, spends less time assessing the load forecast on its own terms in the absence of the arrival of alternative energy competition. Although it identifies Hydro's lack of analysis of alternative population and economic growth scenarios, it does not deal with other important limitations of the Hydro load forecasting methodology. In particular, it does not consider the important effects of rising Hydro rates projected in the NFAT apart from a limited discussion in section 2.1.3 and on page 46.

Forecast Methodology

Manitoba Hydro's load forecasting methodology can be described as one that is evolving but that remains a blend of existing approaches that is at times difficult to understand. It uses a variety of approaches to forecast load that preclude any assessment of within-sample reliability, an important component of any evaluation of prospective forecast error.[4] It also provides limited discussion of its methodology that makes it difficult to assess how Hydro has constructed its models and evaluated them against alternative approaches.

Residential Basic Forecast

The residential load forecast uses an "end use" methodology common in the industry that divides the customer base by dwelling type, area and heating type. The process forecasts residential customers via a consensus (simple average of forecasters) forecast of residential population divided by some past average of people per household (about 2.8 since 1997) and then forecasts the proportion using electric heating for each customer group. This latter forecast of electric heating share used a variety of regression[5] techniques until 2013, when the regression approach was abandoned completely in favour of an ad hoc approach involving an adjusted five-year moving average. This "bottom up" ad hoc approach is not compared to any sort of "top down" econometric approach, such as a set of regression models for the customer groups that would include population, income (GDP or household income measures), prices, weather, and other factors. [Elenchus, (16) also notes the "lack of consideration of alternative models and methods, such as top-down econometric approaches, to test the reference case scenario."] Thus, we have no idea whether the Hydro approach provides superior forecasts to such alternatives, as is implied in the NFAT. There is also no natural assessment of the within-sample reliability of the forecasting technique that would follow from the use of regression methods (e.g. R^2 as a measure of within-sample forecast error). In short, there is no rationale for the forecasting method that is chosen and its obvious deficiencies in providing estimates of prospective forecast reliability.

Manitoba Hydro assumes the number of customers will change proportionately with population. This relies on the assumption that the number of people per household will not change. This has not been true in the past and is unlikely to hold true in the future. The number of occupants per household will be affected by not only the number of people, but the relative ages of the population. For instance, if the fastest growing segment of the population is over 50, there will usually be fewer people per household in the future. Another factor affecting the number of occupants per household is personal income. As income increases, the number of occupants per household generally decreases. In the housing model used by the State Utility Forecasting Group (SUFG) for the state of Indiana, headship

---

[4] Within-sample forecasting error refers to a measure of the differences between the forecast and actual outcomes in the data available to the forecaster, such as the coefficient of determination (R^2) in econometric forecasting models. The measure provides an indication of the extent to which the forecasting methodology can predict known outcomes and, as such, is an indicator of the expected accuracy of the forecast in the short term.

[5] See the section on trend analysis in Part 1 for an explanation of regression-based forecasting.

rates (the inverse of occupants per household) are projected using a logit model that is a function of age, income, marital status, and the prior year's headship rate.

Manitoba Hydro projects the number of dwellings that use electricity for heating from a five-year average and then uses that as an exogenous assumption to the end use model. This nullifies one of the major benefits of end-use modeling, which is the ability to simulate the economic trade-off of different technologies and fuel sources based on the capital and fuel costs of the different options. Ideally the number of new dwellings would be an exogenous input and the fuel choice decision would be handled endogenously by the model.

General Service Mass Market

The forecast of growth of the General Service Mass Market has employed regression models, but the model specification has changed from year to year without any explanation of the rationale (Elenchus, 18). In the current version, only GDP growth and residential customer growth are components of the model, but the regression methodology does permit an assessment of within-sample forecast reliability and a consistent method to choose the appropriate forecasting model going forward. Whether that model selection methodology has been used in the past is unclear, since the basis for the choice of the current forecasting elements (GDP and residential customer growth) rather than alternative specifications is unclear.

This sector is forecast with an econometric formulation to determine the number of customers, using GDP and population as the drivers. The electricity utilization per customer is then assumed to stay constant at the most recent 5-year average. In reality, utilization per customer will not stay constant, especially when real electricity prices are changing.

The SUFG forecasting methodology for Indiana is a little different in that customers are separated into commercial and industrial classifications (as opposed to combining them and separating out the largest customers), but the experience is still informative. Indiana experienced a period of declining (in real terms) rates from 1985 to 2005 and has experienced increasing real rates from then on. During the period of declining rates, the commercial sector saw intensity (in utilization per unit of floor space) increase at 2.4% annually. With the start of rate increases, we start to see declines in intensity of 0.4%. In the industrial sector, intensity (in utilization per real manufacturing GSP) has been dropping since the mid-1980s. During the earlier period of declining rates, intensity fell by an average of 1.2% annually. More recently, the decline has been 1.9% per year.

General Service Top Customers

The forecast of Top Customers is based on assessments from Hydro's own economic experts and account representatives. While Hydro argues that regression techniques are inappropriate for this customer segment, its own methodology has had a consistent upward bias on the order of 5% (Elenchus, 23). Also, there is no justification that this approach is superior to appropriately crafted regression modelling in terms of forecast accuracy, nor is there any assessment in the NFAT of the limitations of the forecasting methodology used. Furthermore, this approach relies on two forecasting methods (informed opinion and trend analysis) from MISO's list of unacceptable methods.

Missing Elements

Electricity is a standard product whose demand should be understood as part of customer demand (the residential customer component) and as an input to production (the General Service Mass Market and Top Customer components). The principle factors in a conventional analysis of demand for a product of this nature would be: the price of the product (electricity), the prices of related products (especially alternative energy products available to residential and commercial customers), income (household incomes and the value of production (GDP)), population, and factors such as weather.

While Hydro spends a great deal of time examining the variation in demand related to weather (cooling and heating days), this is largely a short-term phenomenon unless there are dramatic changes in weather patterns (climate change) that are relevant to the forecasting horizon. Weather variation may account for some of the fluctuation in load demand within each year and, to a lesser extent, across years, while the other factors (and weather trends related to climate change) will account for movement in the trend or average load over time. This trend constitutes the expected long-term forecast about which weather will cause minor variations. In short, there should be less concern about adjustments for weather (which are, in any case, of dubious value in the NFAT according to Elenchus (27-29)) and more concern about the limitations of the trend forecasting methodology.

Elenchus (ii-iii, 30-31, 42-43) makes a similar point in referring repeatedly to the need for a wider range of population and GDP scenarios, since potential variation arising from population and economic growth is ignored in the risk analysis in the NFAT. What is missing in the Elenchus report is some indication of how much alternative population and economic growth scenarios might matter to the comparison of alternative plans, something that will not be resolved directly by the alternative DSM scenarios Hydro is now running. These DSM scenarios may, however, provide some indication of the implications of reduced load projections for the comparison of alternative development plans.

It is also a puzzle why the load growth forecast for Manitoba (1.6%; NFAT, ch.4, p.12) exceeds the load growth forecast for the U.S. (0.9%) despite similar population growth forecasts in Manitoba and the U.S. and higher GDP growth forecasts for the U.S. compared to Manitoba. This was not resolved in the interrogatories. This is an issue about their forecast trend, however, rather than potential variability about the trend arising from population and economic growth uncertainty.

The major missing factor in the load forecast is prices. The NFAT (Exec Summary, 9) admits that energy prices matter but makes no attempt to incorporate what amount to fairly substantial projected rate increases into its load forecast. Moreover, Hydro indicates that it does not pay attention to what is a fairly robust literature on the impact of prices on electricity demand from other jurisdictions. In response to the interrogatories GAC_CAC/MH II-001a and b, Manitoba Hydro did produce some correlations of prices with customer usage, but the results are based on a small number of points and a simple regression analysis that ignores the other important factors in the determination of customer demand. A more detailed analysis, or the application of results from better analyses elsewhere, is needed.

Manitoba Hydro indicates that prices would increase by about 4% per year in nominal terms, or about 2% per year in real terms (NFAT, Appendix D, 55). This should result in a reduction in utilization per customer for a number of reasons: it results in reduced disposable income for customers so they purchase fewer energy using devices, those purchases that customers make are more likely to be made with energy efficiency in mind, customers may opt to use energy sources other than electricity where possible (conversion from natural gas to electricity for space heating is less likely to occur), and customers may change their behavior (adjusting temperature settings, turning off lighting when not in use, etc.). If electric rates have been stable for some time, it is possible that the forecasting model estimation would not capture the impact of price elasticity.

Some illustrative "back of the envelope" calculations might indicate the potential size of the price effects on load forecasting in NFAT. Take the U.S. estimates that a 10% increase in the price of electricity can be expected to reduce household load by around 5% in the long run (http://www.e3network.org/ElasticitySurvey2_matt.pdf). Compounded annually, the projected 2% (real) increase in electricity prices over 30 years amounts to a whopping 80% increase in rates over and above general price inflation. Apply this increase only to the residential sector, which accounts for about 1/3 of load. Hydro projects a load increase of 1.6% per annum for this sector, of which 1.2% is attributed to population growth and 0.4% to increased energy usage (NFAT, ch.4, 12). Over 30 years, this implies load growth of about 60%, with about 45% attributable to population growth and 15% attributable to increased usage per household, ignoring the impact of price increases. Applying the U.S. price elasticity estimates, however, implies that the 80% increase in prices would reduce load by 40% (since a 10% price increase would reduce load by 5%), implying that load per household would actually decline by about 25% over the 30 years due to rising electricity rates.[6] Combined with the load growth due to rising population of 45%, this implies only a 20% increase in residential load or about one-third of the 60% projected by Hydro.

Hydro projects total load growth of about 7,899 Gwh, from 24,367 Gwh in 2011/12 to 32,266 Gwh in 2031/32 (NFAT, ch.12, 2-3). It appears that residential load growth is more rapid than other growth, but assume that only one-third of this growth is residential, or 2,633 Gwh. If actual growth is only one-third of that figure because of reduced household usage due to rising electricity prices, as suggested above, then load growth would be reduced by more than 1,755 Gwh. The NFAT (ch.12, p.2) suggests that one year of load growth constitutes 420 Gwh, so this amounts to a reduction in load growth of 4.2 years. By comparision, the revisions to the load forecast for 2013 amount to a reduction in load growth of 3 years by 2031/32, which defers the need for new resources by one year. This suggests that electricity conservation in the residential customer base alone, arising from the rate increases projected by Hydro, would defer the need for new resources by at least another year. Since the commercial sector would also be sensitive to increases in the price of electricity, reductions in load growth in the General Service Mass Market and Top Customer sectors might be expected to defer load growth correspondingly by as much as three years. While this is only illustrative, these are quite significant numbers that would substantially affect planning.

Indiana has seen an increase in prices in real terms since 2005. The SUFG forecasting models indicate real price elasticities of -0.4 for the residential sector, -0.26 for the commercial sector, and -0.48 for the industrial sector. Thus a 2% real price increase in the residential sector would result in 0.8% less electricity use. While one would expect the actual price elasticities to be different in Manitoba than

---

[6] These are estimates of the average effect of electricity price increases on consumers. There may be differences across households. Lower income households, for example, may be less able to reduce electricity consumption than higher income households, since their electricity use is already devoted primarily to necessities. Harvey Stevens and Wayne Simpson explore this issue in a separate submission to this hearing.

they are in Indiana, there still should be a dampening of electricity demand as real prices rise. Yet Manitoba Hydro is projecting electricity usage to grow at the same rate as it has historically, with an average annual increase of 0.4% in utilization per customer.

It should be noted that The Brattle Group uses a price elasticity of -0.4 in the export price forecast model that was used as an input to Manitoba Hydro's export price forecast. Furthermore, they label that value as conservative and low (NFAT, Appendix 3.1, slide 51).

Assessment of Forecast Reliability

The reliability of the forecast depends on two components: (i) the within-sample reliability of the forecast instrument and (ii) the beyond-sample accuracy of projections of the inputs (e.g., economic or demographic projections) to the forecasting model. In a 30 year forecast based on annual data, it is safe to say that (ii) likely matters far more than (i) because the variability arising from (i) will lead to short-run forecast errors (over a few years) that will largely cancel over the longer run (30 years) while the projection errors from (ii) are far more likely to accumulate over time, e.g. a prolonged economic or population/immigration slowdown.

Hydro's forecasting methodology makes it difficult to assess within-sample forecasting reliability. The extensive discussion of the impact of weather on forecast reliability does not make up for the absence of reliability estimates based on other forecast inputs such as population and economic growth. A consistent econometric approach to forecasting would solve this problem, but other statistical solutions to assess the within-sample reliability of the present forecasting method (Monte Carlo or bootstrapping approaches, for example) are likely feasible as well.[7]

It is also very difficult to assess the reliability of the beyond-sample accuracy of projections to the inputs to the forecasting model. Elenchus emphasizes the need for alternative population and economic growth projections to assess the sensitivity of the load forecasts, and subsequent plan evaluations to these two components of their forecasting model. In addition, there are important inputs to the load forecast, such as prices, that are not considered at all by Hydro or extensively by Elenchus.

The most disturbing omission from the Hydro forecasting methodology must be the impact of rising electricity prices because all the evidence implies that the bias introduced by this omission is upward; that is, the omission of price effects leads to inflated load forecasts and requirements for new system capacity. Indeed, our illustrative results with fairly conservative estimates of the responsiveness of U.S. consumers to electricity price increases imply requirements for new system capacity may be overestimated by several years.

[7] For any forecasting methodology where the data can be measured and characterized in terms of one or a series of empirical probability distributions, repeated draws from the distribution(s) can be used to measure the difference between the forecast and actual outcomes to assess forecasting error.