Conf-950309--4

NREL/TP-441-7803 • UC Category: 1210 • DE95009203

Estimating the Economic Value of Wind Forecasting to Utilities

Michael R. Milligan
Alan H. Miller
Francis Chapman

Prepared for
Windpower '95
Washington, D.C.
March 27-30, 1995

National Renewable Energy Laboratory
1617 Cole Boulevard
Golden, Colorado 80401-3393
A national laboratory of the U.S. Department of Energy
Managed by Midwest Research Institute
for the U.S. Department of Energy
under contract No. DE-AC36-83CH10093

Prepared under Task No. WE517050

May 1995

Available to DOE and DOE contractors from:
Office of Scientific and Technical Information (OSTI)
P.O. Box 62
Oak Ridge, TN 37831
Prices available by calling (615) 576-8401

Available to the public from:
National Technical Information Service (NTIS)
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161
(703) 487-4650

Printed on paper containing at least 50% wastepaper and 10% postconsumer waste

DISCLAIMER

ESTIMATING THE ECONOMIC VALUE OF WIND FORECASTING TO UTILITIES

Michael R. Milligan
Alan H. Miller
National Renewable Energy Laboratory
1617 Cole Boulevard
Golden, CO 80401
U.S.A.

Francis Chapman
Environmental Defense Fund
5655 College Avenue
Oakland, CA 94618
U.S.A.

ABSTRACT

Utilities are sometimes reluctant to assign capacity value to wind plants because they are an intermittent resource. One of the potential difficulties is that the output of a wind plant may not be known in advance, thereby making it difficult for the utility to consider wind output as firm. In this paper, we examine the economics of an accurate wind forecast, and provide a range of estimates calculated by a production cost model and real utility data. We discuss how an accurate forecast will affect resource scheduling and the mechanism by which resource scheduling can benefit from an accurate wind forecast.

INTRODUCTION

Wind generated power has become a cost-effective utility scale generating resource. The costs per unit of energy for a wind facility have declined steadily over the last decade. As the electric utility of the future attempts to provide service while minimizing the costs of providing electricity, wind generated electricity promises to increase its share of the generation mix. For an electric utility seeking to integrate the output from a wind resource there remain, however, some issues. The intermittent nature of wind energy causes problems to utilities, because they must be able to serve loads with an acceptably low probability of failure. The economic, social, and political costs of failing to provide adequate capacity to meet demand are so high that utilities have traditionally been reluctant to rely on intermittent resources for capacity.

One of the factors in this intermittency issue is the uncertainty of the wind resource itself. Utilities often provide backup generation sources to compensate for the possibility of unanticipated low- or no-wind conditions causing unexpected outages at wind generation facilities. Knowing in advance how much wind power will be available could be valuable to an electric utility. With perfect foresight, the utility would not need to provide backup for the wind resource, and could thus achieve savings by avoiding excess generation.

In this paper we seek to quantify the potential value of accuracy in a wind forecast. The results are generated by using an electricity production cost model to measure the cost implications of varying degrees of accuracy in a wind forecast. We do not attempt to forecast wind or wind energy, merely to analyze the value associated with a particular degree of accuracy in a forecast. The role that the accuracy of a wind forecast can play is to reduce risk as the utility attempts to minimize the cost of service to its customers.

BACKGROUND

The simplest benefit of an accurate wind forecast is that wind generated energy can be planned for and used by the utility; the utility thus avoids the need to consume a fuel to produce electricity. Wind can be viewed as a free fuel. If all of the output of a wind turbine replaces the output of a generator which consumes a non-free fuel, then savings are maximized. This simplified point of view is complicated by constraints imposed by integration of the wind resource with the rest of the electricity supply system.

An electric utility is charged with two competing tasks, providing service with an acceptably small risk of failure, and providing service at the lowest cost. The risk, or reliability target, translates into an amount of reserve on-line generating capacity. Too little reserve capacity violates the utility's reliability obligations, too much reserve capacity violates the utility's cost minimization mandate. In calculating operating reserve margins, utilities often view wind plants as nonfirm resources, i.e., as having little or no capacity value. If the wind resource is viewed as strictly nonfirm, the utility must provide firm capacity to support the wind plant. If wind plant output can be forecast, surplus wind-supporting capacity can be minimized.

Wind power has value even if forecasts are not accurate. Clearly, the amount of wind energy that is available at a given time is independent of the forecast. Unexpected wind power can still be used to displace energy that is provided by a load-following unit. However, as wind plants continue to become more economically competitive with conventional energy sources, it is important to correctly assess the capacity value of wind. A wind resource that cannot be counted on for capacity has minimal contribution to system reliability. An accurate forecast allows the utility to count wind capacity, reducing costs without violating reliability constraints.

We address two processes, the commitment process and the spinning reserve process, both of which must be considered by utilities when scheduling resources to meet load, and which affect the value of a wind forecast. Understanding of these two processes is helped by a quick outline of some issues affecting utility resource planning decisions.

An electric utility meets instantaneous demand by coordinating output from a variety of generation sources. The output of multiple independent generators must be provided in such a way as to provide electricity to the distribution grid with the correct frequency and voltage. If a generator fails suddenly, and there is no compensating action to make up for the lost capacity, the whole grid can fall out of balance, potentially requiring a system-wide shutdown and restart. To avoid imbalance, a utility will keep some capacity spinning in reserve, i.e., a running generator capable of producing electricity but not providing current to the grid. If an on-line generator fails, the spinning capacity contributes to the grid and restores balance.

Wind turbines are intermittent generators. If wind turbines are a non-trivial component of system capacity, their intermittent output can unbalance the grid. One utility strategy to smooth output from wind turbines is to have spinning reserve capacity assigned to compensate for fluctuating wind output. We refer to this backing-up of the output as firming, or making firm, the wind resource. Wind generation is less cost-effective if a utility has to compensate for intermittency by running fossil-fired generators as firming resources. With perfect foresight of wind generator output, wind energy can replace fossil-fueled energy on a one-to-one basis. The spinning reserve benefit of an accurate wind forecast would be to allow the utility to reduce the amount of spinning capacity it requires to firm up wind generation.

Commitment issues involve a longer time scale than spinning reserve issues. First we will explain what commitment refers to, then how commitment is affected by a wind forecast. The amount of lead time required for a generator to go from a nonoperating state to an electricity producing state varies among technologies. An example of a "quick-start" unit is a combustion turbine, whose lead time can be measured in minutes. An example of a "slow-start" unit is a nuclear plant, which requires hours or days of lead time to bring to an operating state. Once a slow-start unit is operating it must continue to run or go through a lengthy shutdown and restart. We refer to a slow-start unit's generation threshold, the level below which the unit must be shutdown and restarted, as its minimum output level. Slow-start units are ramped down during low load periods, possibly as far as their minimum operating level, and then ramped up to assist in meeting peaks. Intermediate and peaking units are also ramped up or started, or both, in order to meet the peak. This method of operation is termed load-following.

An electric utility meets loads through a combination of slow- and quick-start capacity. Typically, the utility must plan at least one day ahead in order to have sufficient slow-start capacity on-line to be able to meet the next daily peak. The process of deciding which slow-start units to operate is termed the unit commitment process.

In order to minimize production costs, an electric utility will favor generation from least-cost sources. Plants will be dispatched in a merit order, with decreasing capacity factors for higher priced plants. Slow-start capacity tends to have high fixed costs (including capital costs) with low operating costs, whereas quick-start capacity tends to have low fixed costs and high operating costs. The utility attempts to minimize total costs, by generating with the least-cost mix of slow- and quick-start units which satisfies the obligation to serve and the reliability criteria.

The utility must commit sufficient slow-start capacity to be able to meet as much demand as possible with low-cost generators. However, committing too much slow-start capacity has a cost. When the system reaches minimum load conditions, the capacity across the minimum operating levels of committed slow-start plants may exceed demand, causing the utility to dump or sell excess energy. Utilities are constrained by the lead times required to cool down and restart slow-start units; they cannot simply turn off slow-start units that are surplus at night and restart them for the next day's peak. Typically, a utility avoids the problem of having too much slow-start capacity during non-peak hours by relying on quick-start units to serve peak loads for a few hours each weekday.

In contrast to overcommitment problems, if a utility undercommits, i.e., starts up too few slow-start units, it will rely excessively on higher cost quick-start energy or purchases to meet peaks. In the extreme case, the utility may even be short of capacity at the time of system peak.

Optimized commitment is made more complicated by intermittent output from a wind resource. If wind generation represents a significant portion of total capacity, the utility must balance the risk of the wind resource being a noncontributor against the risk of committing excess slow-start capacity. A commitment benefit of an accurate wind forecast would be to allow the utility to optimize the amount of slow-start capacity it commits, minimizing the risk of producing either too much or too little low-cost energy.

An accurate forecast could enable a utility to plan its generation schedule by using the wind power forecast as the "schedule" of wind output. As long as the forecast is accurate, the utility finds itself in its optimal position, able to integrate wind output without backup or support for the wind resource. An inaccurate forecast, however, causes problems. If the forecast is too optimistic (i.e., forecasted wind exceeds actual wind) then the utility finds its reserve margin too low. A capacity deficit could violate minimum reliability constraints imposed on utilities by regulators, and may cause system costs to be higher than optimal.

Conversely, the reliability target can be exceeded when the forecast is too pessimistic (the forecast wind is less than the actual wind). In this case the utility has completed its commitment list, assuming the forecast wind power level. The actual wind power output is greater than forecast. From a reliability perspective, units are committed that need not be, resulting in excess spinning capacity. While the excess reliability is not a problem itself, it does result in a nonoptimal schedule of resources, increasing energy costs needlessly. It is interesting to note that this type of scheduling error reduces the likelihood that wind is accredited with the appropriate capacity value.

METHODS

We used the Elfin probabilistic production cost model, version 2.05[1]. Given forecasted loads and data on the existing resource base as inputs, Elfin calculates expected generation from each resource. We modeled differing degrees of accuracy in a wind forecast using techniques described below, and measured the total cost of producing electricity in each of these scenarios. The differences in total cost for each scenario were measured and used to project the benefits of accuracy in the forecast. All production simulations were carried out for a single year, and forecast benefits are measured as annual values. There are long-term implications that wind forecasting accuracy would have on optimal generation expansion, but they are not explored here.

We use the Wind Power Simulator (WIPS) to calculate wind plant output. WIPS uses wind turbine parameters and wind strength data to calculate the net power output of a wind resource. WIPS takes factors into account such as the hub-height of the turbine, losses due to mechanical, electrical, and wake effects. WIPS is further described in Milligan and Miller (1993).

We developed two methods to analyze the benefit of a wind forecast. The first one uses hourly wind data to model the wind power forecast and hourly data to represent the deviation from forecast. This is called the hourly method. A second method uses the forced outage rate of the wind plant as an approximation for the forecast error. This is called the forced outage rate (FOR) method.

The first method, which uses hourly wind power data, is described here. The scenario begins with the meteorological forecast, which we assume results in an hourly wind speed forecast for the time period in question (on the order of 24–48 hours). This data is used to calculate the hourly wind power output, which is then given to the utility dispatcher. The power system schedule is then based on the wind forecast, and decisions are made on which of the units should be committed. Actual wind power may vary from its predicted value. The extent of the wind forecast error determines the financial and reliability cost to the utility. If the forecast error is zero, then the utility has planned in an optimal way. Penalties occur when errors in forecasting occur, and these penalties increase as a function of the forecasting error. Our modeling attempts to capture these scenarios in as much detail as possible. We provide Elfin with the hourly wind power forecast, so that optimal resource scheduling can be completed by the model. The wind forecast is treated as any other scheduled resource and is considered firm. We then provide Elfin with hourly deviations from the forecast. In order to capture the best time resolution possible, we modeled the wind resource as a load modifier. While it would be useful to use the probabilistic modeling technique developed in our earlier paper, it was not possible to model forecast errors with that method.

In order to account for the effects of wind power deviating from the forecast, we assume the utility relies on the forecast in planning its commitment schedule. Wind power above or below the forecast is not counted towards the commitment target. In other words, the forecast wind power counts towards commitment, whereas the deviation from the forecast does not. Among all scenarios with a constant amount of wind energy, differences in total production cost are a function of the portion of wind output that was counted on during the commitment process.

We believe that this method is extremely flexible and is a reasonable way to approximate the forecasting error. The method allows us to separately examine the question of the effects of overforecasting versus underforecasting. As examples, consider these cases. First, suppose that the forecast is completely accurate.

---

[1] Elfin uses the Baleriaux-Booth method to account for the probabilities of full or partial outages occurring at a generator. For a more detailed discussion, see the Elfin Algorithms Guide, Environmental Defense Fund, (1992).

That implies that the nonfirm component of wind is zero, and the hydro-thermal system can be scheduled in an optimal way (ignoring other sources of uncertainty). If the forecast is too high, we reduce actual wind power from the forecast level, but do not allow commitment to change. Similarly, if the forecast is too low, we increase available wind power, but do not count this increase towards commitment.

This method could also be expanded so that other forecast error patterns can be introduced. Here, we assume the forecast error is consistently higher or lower than actual wind output. In reality, forecast errors are likely to be too high in some hours and too low in others. It is also important to note that, unless forecast error is explicitly taken into account, it is *implicitly* assumed to be zero. Our earlier paper, for example, assumes that no forecast errors occur and that the hydro-thermal system can be optimally dispatched.

We used 2 large, interconnected utilities, Utility A and Utility B, with a broad range of resources, including hydro, oil and gas, and nuclear. For this study, we did not allow Elfin to price any resources at system lambda, because altering the forecast error tends to change system marginal costs and thus artificially alters total cost. We also held the load profile constant, so that differences in correlation between windpower and load would not play a role in influencing the results.

We used a single year of wind data from the High Plains wind site, calculating equivalent power output, and running the Elfin model. We ran 112 different cases per utility, each of which represents a different combination of actual and forecast wind power. The actual wind power was simulated at 10% increments of the historical wind levels. If, for example, the historical wind data resulted in wind output of 100 MW for the hour, we ran at 10 MW, increasing by 10 MW, up to 100 MW. For each of these wind power levels, there are several ways of combining 10% increments of forecast wind and forecast error values. For example, 50 MW of actual output can be realized, by several combinations, such as a 60 MW forecast and -10 MW forecast error; 70 MW forecast and -20 MW forecast error, or 30 MW forecast and 20 MW forecast error. For the site maximum wind speed, which translates to rated wind plant output, forecast error can only be positive. Likewise, for 0 MW actual wind power, the forecast error cannot exceed 0. However, zero wind output can be achieved with offsetting values of forecast and forecast error: 0 forecast and 0 error; 10 forecast and -10 error, for example.

The method allows us to choose any combination of wind power forecast and forecast accuracy. For a given hour, for example, we can forecast 50 MW of wind, and actually have 80 MW of wind – the forecast error is 30 MW. In an extreme case of a missed forecast, we can simulate a 100 MW forecast with 0 MW wind (forecast error = 100 MW). In this case, the utility has planned on 100 MW of wind, but must dip into reserves to meet load. Conversely, the utility can plan on 0 MW of wind and receive 100 MW of wind. However, spinning reserve and commitment levels are set in advance, so the benefit of 100 MW of unplanned wind power is diminished. It is important to note that we hold wind energy constant in order to assess the value of the forecast. Therefore, the primary impact of the forecast error enters via the unit commitment logic. It is also important to note that we made sure that spinning reserve targets were also held constant in the various scenarios. However, the units that are called on to meet the fixed spinning reserve target may in fact differ from the perfect forecast case. This is due to differences in unit commitment that arise from differing forecast accuracies.

We now describe the forced outage method. Because of the large number of exogenous calculations required by this method, we were unable to obtain annual forecast values. We did, however, obtain results from this method that support the results of the hourly method described above, by applying the FOR method to the utility's peak week. We study the effects of an accurate forecast on the commitment process by changing the reliability, or variance, of the wind capacity while holding everything else constant.

A utility would have to forecast wind perfectly to maximize a wind resource's capacity benefits during the commitment process. With perfect foresight, the utility could rely on wind capacity as being equally as firm as a thermal resource. However, even thermal resources are not all equally firm. The reliability of a unit is expressed as a forced outage rate, or a probability of a full or partial outage. Coal units and older gas-fired steam boiler forced outage rates can range from 10% to 40%, while newer gas units are projected to achieve forced outage rates in the range from 3% to 10%. In choosing a combination of units to have on-line to meet peak load, a dispatcher must account for the reliability of each unit expected to contribute to load. The Elfin model accounts for reliability effects on commitment by derating each unit's capacity by the unit's forced outage rate.

In analyzing the value of a wind forecast we address two components, capacity and energy. In order to isolate the capacity value of accuracy, we model the wind resource as a thermal resource and keep the energy contribution from the wind resource constant across cases. We correlate actual wind generator output to forecasted wind output by using the unit's forced outage rate. A zero forced outage rate models a perfectly reliable forecast, while a forced outage rate of 40% implies a forecast error of 40% (accuracy of 60%), and so on. All scenarios contribute an equal amount of effective capacity, i.e., capacity after derating for forced outages. We kept the effective firm capacity contribution of the wind resource constant by modifying the nominal rating in proportion to the forced outage rate. For example, to achieve an effective contribution of 50 MW, we could have a 50-MW nominal resource with a zero forced outage rate, or a 100-MW nominal resource with a 50% forced outage rate, or a 200-MW nominal resource with a 75% forced outage rate. In order to isolate the effects of the accuracy of the forecast from other factors, we recorded total costs at each scenario's optimal commitment level[2]. This technique therefore solves for the optimal (i.e., least-cost) commitment target at each forecast accuracy, measuring the cost differences to account for the economic benefit of the accurate forecast.

To capture commitment effects, the wind resource had to be sizable enough to cause changes in total production cost because of changes in firmness, or reliability, of the wind plant. Therefore, the capacity portion was modeled to yield a constant 4% of effective system peak-coincident firm capacity, across all cases, i.e., 4% of peak load would be expected to be served by the wind plant[3]. By modeling varying degrees of forecast accuracy and recording total cost at the optimal commitment level, while holding the wind energy contribution and all other factors constant, we calculate the change in system variable costs due solely to the change in accuracy of the wind forecast.

RESULTS

We begin with results from the hourly method. Some results from the hourly method are presented in Table 1, below. The results of all of the hourly cases show that an accurate wind forecast does indeed result in the highest benefit from the wind resource – and for inaccurate forecasts, wind benefits are reduced. The direction of the forecast error does influence the benefit reduction for a given utility, but is not constant between the utilities we analyzed. The relative costs of over- or underforecasting will depend on the tradeoffs faced by each utility – what is the penalty imposed by the power pool or regulators for undercommitting, as balanced against the cost of overcommitting.

Our modeling shows that an accurate wind forecast does have value. Several conclusions can be drawn, although the specific results depend heavily on the utility's generator mix, load swings, correlation of wind to load, and other factors. First, accurate forecasts always have economic value, even at lower levels of wind output. Second, even if there is no wind power available, it is in the utility's best interest to forecast this event accurately. Third, our results indicate that a forecast that is too high is more costly than one that is too low for Utility A, but the

---

2 Optimal commitment was calculated to an accuracy of one-tenth of a percent of peak load.

3 To put the 4% figure in perspective, it represents a larger proportion of firm wind capacity to total peak capacity than any major U.S. utility could currently achieve.

opposite is true for Utility B. These results are highly sensitive to the utility's contractual and power pool arrangements and should be studied more closely in that context.

TABLE 1 - FORECAST BENEFITS, UTILITIES A, B

50% Wind Output			100% Wind Output		
Forecast Error (MW)	Util. A ($1,000,000)	Util. B ($1,000,000)	Forecast Error (MW)	Util. A ($1,000,000)	Util. B ($1,000,000)
-175.5	50.0	26.0	0.0	106.9	96.0
-140.4	50.9	29.0	35.1	105.2	92.0
-105.3	51.7	33.0	70.2	103.3	87.0
-70.2	52.4	37.0	105.3	101.3	83.0
-35.1	53.0	42.0	140.4	99.6	79.0
0.0	53.8	47.0	175.5	99.0	74.0
35.1	52.3	42.0	210.6	97.8	70.0
70.2	50.8	40.0	245.7	96.2	64.0
105.3	49.5	34.0	280.8	95.3	60.0
140.4	48.6	29.0	315.9	95.0	56.0

Figures 1 and 2 illustrate the results of some of the simulation runs. Figure 1 assumes that actual wind power output is at 50% of rated capacity for the wind plant. The benefit, as measured on the vertical axis, shows that with a perfectly accurate forecast (forecast error of zero) there is about a $54 million benefit provided by the wind plant. This benefit is measured as fuel saving from conventional units. The diagram shows how this benefit is reduced as the absolute value of the forecast error increases. In the case of a 50% forecast error, the benefit of the wind plant is reduced to about $50 million. Therefore, the forecast provides approximately $4 million in benefits. Conversely, a negative 50% forecast error results in a benefit reduction of about $2 million. The graph in Figure 1 has a typical shape for the various percentages of wind capacity we looked at.

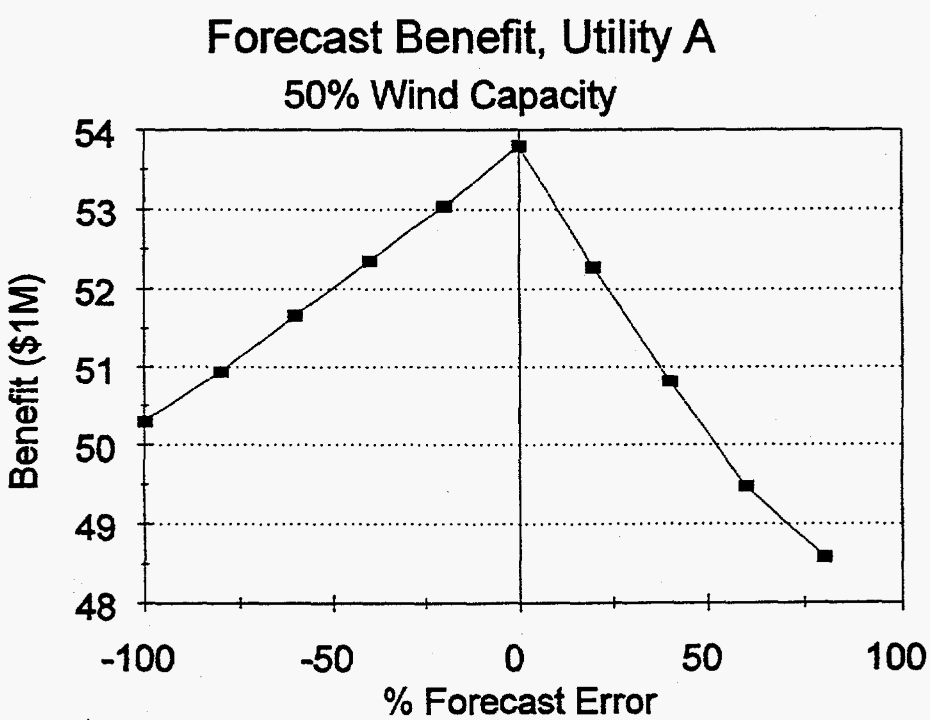

FIGURE 1 - EFFECT OF FORECAST ERROR ON BENEFIT, 50% OUTPUT

Figure 2 assumes full wind plant output. With full output, it is not possible for the forecast to exceed the actual wind. Therefore, the graph only shows positive forecast errors (i.e. in all cases, actual output exceeds forecast output). This graph shows that a forecast error of 90% costs the utility about $12 million. The consequences of less severe errors can be read from the graph – for example, a 40% forecast error results in approximately $7.5 million in lost benefits.

Our test wind plant has a capacity of 1,250 MW. After accounting for various losses the net output capability is 1,006 MW. For such a wind plant our results indicate that an accurate forecast can be worth up to about $12 million/year for utility A, and up to $40 million/year for utility B. The difference in results for the two utilities reflects differences in costs of the marginal units displaced by wind generation. A utility that has a relatively large slow-start unit on the margin during low-load periods is likely to see a larger cost impact of an accurate forecast. Of course it is extremely unlikely that forecasting with 100% accuracy can be achieved on average, but this does suggest that it would be worthwhile for a utility to pursue some reasonable level of forecasting accuracy.

Results from the FOR method support those from the hourly method. We calculated optimal commitment for each of several utilities and measured the savings that would result in a new optimum based on an accurate forecast. Figure 3 graphs total production cost versus forced outage rate for a wind resource, across three distinct

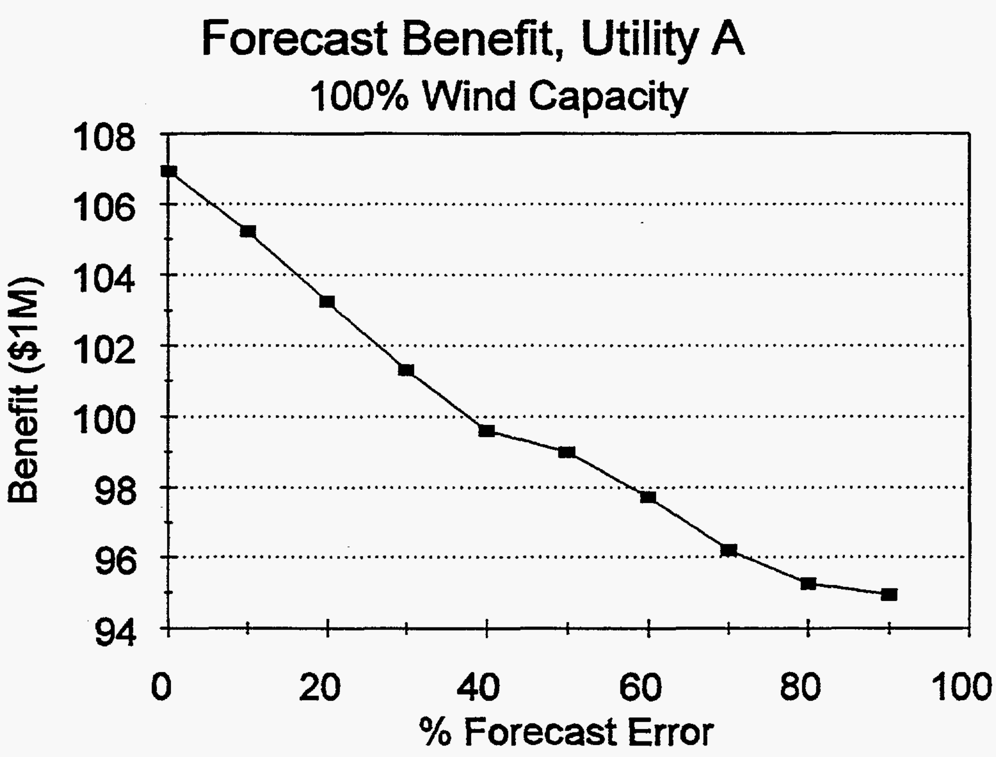

FIGURE 2 - EFFECT OF FORECAST ERROR ON BENEFIT, FULL OUTPUT

utilities. Utility A could not be solved for its optimal commitment level, so was not used in the FOR cases. In addition to Utility B, we also used Utility C. a synthetic data set, with a 2,905 MW peak and 4,028 MW of firm slow-start capacity, and Utility D, which is based on a small, non-interconnected utility with a 1,250 MW peak and 1,435 MW[4] of firm slow-start capacity. Modifications were made to Utility D to increase the proportion of slow-start to quick-start capacity. Total production cost penalties caused by inaccurate forecasts were normalized amongst the three utilities. The shape of the curve indicates declining benefits as the accuracy of the forecast increases. If one assumes that the cost of a wind forecast increases with the degree of accuracy, then the benefit to cost ratio of an investment in improving the accuracy of a wind forecast declines as accuracy increases.

The shape of the curve is a product of merit order dispatch. To achieve least-cost generation, the first unit dispatched should be the least-cost unit, the second should be the second lowest cost unit, and so on until the most expensive unit needed is dispatched last. Unit capacity factors should decrease as production costs go up. The marginal units, those that would increase production to serve an extra unit of demand, start with the highest priced resource and will exclude the least cost, or baseload, generators. Compared to a non-wind case, the capacity scenario with a zero forced outage rate displaces the top, or highest-priced, 4% of marginal generating capacity. As the forced outage rate increases, the probability of high priced marginal generation being required to contribute to load during wind resource outages increases, driving up expected costs. This modeling result should hold true for a utility seeking to integrate a significant wind resource. If some portion of wind capacity was firm because of an accurate wind forecast, displaced firm resources would likely be idled in descending cost order.

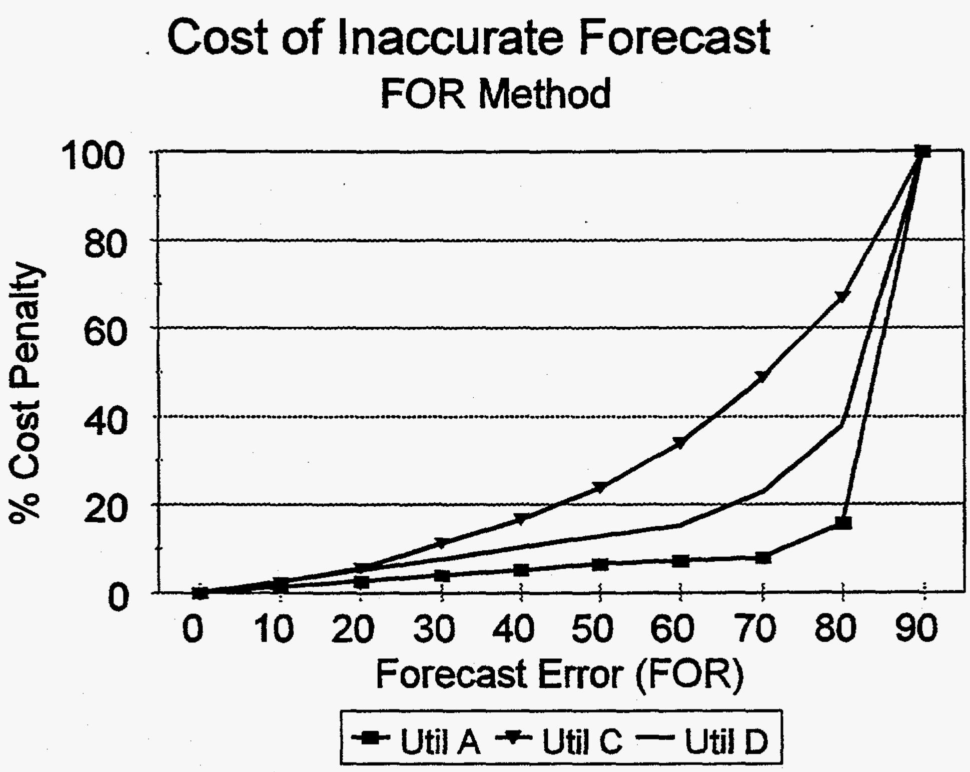

FIGURE 3 - WIND FORECAST ERROR EFFECT ON COSTS, FOR METHOD

---

[4] This total excludes 260 MW of capacity firm for dispatch but nonfirm for commitment.

CONCLUSIONS

Forecast benefits will vary according to the utility's load pattern, generation mix, fuel costs, correlation of load with wind power, and other factors. However, in all cases we analyzed, an accurate wind forecast has economic benefits. The declining marginal benefit curve from the forced-outage method suggests that the goal of a wind forecasting project should not be 100% accuracy, but should balance the degree of accuracy against the degree of benefit derived by that level of accuracy. Further work in this area, however, is clearly needed.

The hourly load-modifier method allows us to distinguish between the impacts of underforecasting or overforecasting wind power. As outlined in the discussion above, the reliability and economic impacts of under- or overforecasting may not be symmetric, owing to differences in the cost structure and consequences of overcommiting and undercommiting. The benefits of accurate wind forecasting that we calculate are significant, but should be interpretted as a first approximation of an iterative process. Further improvements on our techniques could improve the accuracy of these results. It is also important to note that, given an inaccurate forecast, we assume that the only effect is via the utility's commitment decisions. All wind energy, therefore, is assumed to be used by the utility. This means that the economic benefits we measure here are probably understated. However, given these caveats, a couple of points can be made. First, a utility that treats wind generators purely as economy energy resources with no firm capacity would be the best candidate to invest in some degree of accuracy in a wind forecast. Second, concentration of wind generators in regions where wind forecasts were maximized for accuracy could bring economies of scale, resulting in more cost-effective forecasting.

Third, utilities are already managing uncertainty and intermittency in forecasting load, even without the use of wind or other intermittent power generators. Weather forecasting plays a role in load forecasting, especially in areas which experience extreme weather conditions. Many utilities have already invested in weather forecasting and may be able to leverage existing weather forecasting capabilities in developing wind forecasting expertise. Other intermittent, weather-dependent renewable resources, such as solar-powered generators, could also benefit from accurate forecasting techniques.

To the extent that an improved wind forecast reduces the use of fossil fuels, pollution would also be reduced. If this cost is internalized by the utility, the economic benefit of the accurate wind forecast would be increased accordingly.

REFERENCES

Elfin Algorithms Guide. (1992). Version 1.92. Oakland: Environmental Defense Fund.

Milligan, M.R; Miller, A.H. (1993). "The Value of Windpower: An Investigation Using a Qualified Production Cost Model." *Windpower '93 Proceedings; July 12-16, 1993; San Francisco, California.* Washington, DC: American Wind Energy Association; pp. 74-81.