Analysis of the Value for Unit Commitment of Improved Load Forecasts

Benjamin F. Hobbs
Dept. of Geography &
Environmental Engineering
The Johns Hopkins University
Baltimore, MD 21218
bhobbs@jhu.edu

Suradet Jitprapaikulsarn,
Sreenivas Konda,
Vira Chankong,
Kenneth A. Loparo
Dept. of Systems, Control,
& Industrial Engineering
Case Western Reserve University
Cleveland, OH 44106

Dominic J. Maratukulam
Energy Delivery and
Utilization Division
Electric Power Research Institute
Palo Alto, CA 94303

Abstract: Load forecast errors can yield suboptimal unit commitment decisions. The economic cost of inaccurate forecasts is assessed by a combination of forecast simulation, unit commitment optimization, and economic dispatch modeling for several different generation/load systems. The forecast simulation preserves the error distributions and correlations actually experienced by users of a neural net-based forecasting system. Underforecasts result in purchases of expensive peaking or spot market power; overforecasts inflate start-up and fixed costs because too much capacity is committed. The value of improved accuracy is found to depend on load and generator characteristics; for the systems considered here, a reduction of 1% in mean absolute percentage error (MAPE) decreases variable generation costs by approximately 0.1%-0.3% when MAPE is in the range of 3%-5%. These values are broadly consistent with the results of a survey of 19 utilities, using estimates obtained by simpler methods. A conservative estimate is that a 1% reduction in forecasting error for a 10,000 MW utility can save up to $1.6 million annually.

Keywords: Load forecasting, power system economics, power generation dispatch

1. INTRODUCTION

Electric utilities make many short-term resource commitments that require forecasts of loads from a few minutes to one week ahead of time. Such decisions can include:

- commitment of generating units,
- short run hydropower scheduling,
- economic dispatch of committed units,
- predictive automatic generation control,
- spinning reserve,
- fuel allocation,
- short-term energy purchases and sales,
- real-time prices,
- load interruption,
- load control,
- generator and transmission line maintenance, and
- available transmission capability.

Forecast errors result in increased costs, or "regret." For instance, if loads turn out to be lower than forecast, then:

- units may have been unnecessarily committed, raising fuel costs and, perhaps, maintenance expenses,
- expensive power may have been purchased which wasn't needed, or a profitable opportunity to sell bulk power might have been spurned,
- hydropower may have been produced which would have been more valuable if generated at a later time,
- overly high real-time prices might have been quoted, depressing sales, or
- unnecessary interruptions or load controls might be invoked, annoying consumers and lowering revenue.

On the other hand, if loads are greater than anticipated, the following types of regret might result:

- Insufficient resources may be available for meeting security constraints, such as spinning reserve margins, thus endangering system reliability. (Zhai et al. [1] have analyzed the effect of load uncertainties on the probability of having insufficient committed capacity to compensate for unit failures and/or unanticipated load variation. Here, we examine not just these risks but also economic risks.)
- To meet the unanticipated load increase, uneconomic generation or purchases of spot power might be necessary. Alternatively, load interruptions or controls might be invoked that could have been avoided had the load been perfectly forecast.
- Commitments to sell power may have been made at a price less than the value of that power to the utility.
- Too low real-time prices might have been quoted, resulting in revenue falling short of the utility's cost.

The value of more accurate forecasts is the amount by which their use would reduce these various sources of regret.

The economic value of improved short-term forecasts is of particular interest now because of the recent development of new forecasting methods. These methods include artificial neural nets (ANNs), state-space approaches, stochastic models, and expert systems, in addition to refinements of traditional time series and regression methods. ANNs are a particularly promising approach because they do not require adoption of a particular functional relationship between inputs and outputs, and because of their ability to adapt as new data becomes available [2].

Proponents of these new methods argue that the more accurate forecasts those methods yield are valuable for the sorts of reasons just listed [e.g., 3]. But the worth of improved accuracy is rarely quantified, and so it has been difficult to compare the costs of new methods with their benefits.

To better understand the benefits of improved short-term load forecasts, we undertook two studies to estimate this

PE-353-PWRS-1-09-1998 A paper recommended and approved by the IEEE Power System Analysis, Computing and Economics Committee of the IEEE Power Engineering Society for publication in the IEEE Transactions on Power Systems. Manuscript submitted December 31, 1997; made available for printing October 7, 1998.

value for a particular neural net-based load forecasting system, ANNSTLF [4,5].

The first of these studies was a survey of 28 utilities that are members of the EPRI ANNSTLF users' group [6]. Of the 19 utilities that compared the accuracy of ANNSTLF to other methods, 18 have found that ANNSTLF has significantly improved forecast accuracy, in some cases by several percentage points. Nineteen utilities claimed significant economic benefits from using ANNSTLF. A conservative estimate of these benefits is an average of $800,000/utility, mostly stemming from improved unit commitment and power transactions. However, just four of those utilities actually quantified the commitment benefits. They concluded that a 1% reduction in mean absolute percentage error (MAPE) would translate into a $1.7, $28, $42, or $143 annual benefit per peak MW of demand (depending upon the utility). The highest value represents approximately 0.15% of that utility's variable generation costs; it results from an annual savings of $7.6M for a utility with a peak load of 35 GW for whom ANNSTLF lowered MAPE by 1.5%.

However, the above estimates were obtained by spreadsheet calculations making simple assumptions about how much would be saved by avoiding excessive power purchases (due to underforecasts), and how many unnecessary unit startups (due to overforecasts) would be avoided. To verify the reasonableness of those estimates, and to explore their sensitivity to various assumptions, a second study involving a detailed simulation analysis has been undertaken. The methodology and results of the latter study are the subject of this paper.

In the remainder of paper, we discuss the methodology used (Section 2), the case study utilities (Section 3), and the results (Section 4).

2. METHODOLOGY

The elements of the short-term decision problem can be structured as a decision tree (Fig. 1). A basic tree would start with a decision node representing "here-and-now" decisions that must be made before the future is perfectly known. Unit commitment would be an example of such a decision. Then in each hour (or some other appropriate period of time), a set of chance nodes would represent the realization of the actual load. A last set of decision nodes would then represent "recourse" decisions that can be postponed until after the true load is known. "Recourse" decisions would include real-time dispatch in which the exact level of operation of committed units can be altered in response to observed changes in load. Rescheduling of units that can be ramped up quickly is also a possibility [1]. A path through the tree represents a particular sequence of decisions and loads; at the end of each path is the cost associated with those choices and outcomes. After structuring the problem in this manner, the benefit of improved forecasts can be assessed by first examining how better information would alter here-and-now decisions and then calculating the expected cost savings.

We consider the use of short-term forecasts for unit commitment and evaluation of power sales/purchases, as they are

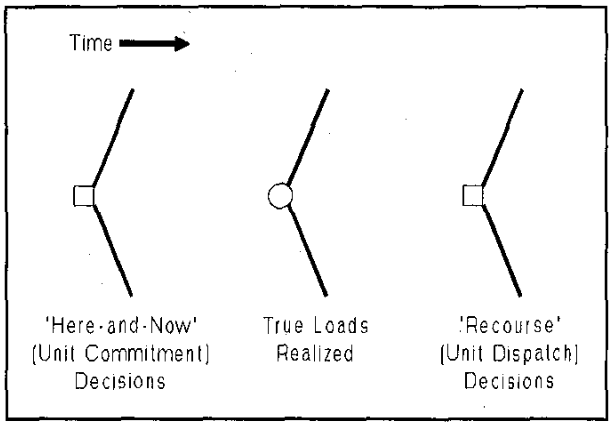

Fig. 1. Decision tree for valuing forecast accuracy

the most common application of such forecasts [7]. Our survey found that 13 of the responding utilities ranked unit commitment and dispatch as the most important use of forecasts, while 7 mentioned transactions as having the highest importance.

In this study, unit commitment is the here-and-now decision, and unit dispatch and spot market purchases are the recourse decisions. We also allow for recommitment of more costly combustion turbines as a recourse decision, since they can ramp up quickly relative to thermal units. (Alternatively, the cost of turbine power could be interpreted as a proxy for the cost of spot purchases to make up for capacity shortfalls.) There are assumed to be 25 decision points in the decision tree: the unit commitment decision (assumed to take place the previous day) and 24 hourly dispatch decisions. The information available at the time of the unit commitment decision is the 24 hour load forecast for the next day, while at the time of dispatch, the true load in that hour is assumed to be known. The only random variable is load itself.

We obtain decisions and outcomes via two models:

1. A unit commitment model [8], which uses Lagrangian relaxation to calculate a least cost unit commitment schedule based on the forecasted load. As is typical with Lagrangian algorithms, the model finds an optimal dual solution and then uses heuristics to construct a feasible primal schedule. The particular implementation in [8] produces multiple feasible primal solutions. To ensure that a good schedule is obtained for a given projected load, all the solutions for all load profiles (i.e., for all values of λ, defined below) for a given day are pooled. Then each solution is tested against the given projected load using the recourse model, described next. The schedule that performs best in the recourse model for the given load is then selected.
2. A recourse model, which in each hour optimizes the dispatch of the units and commitment of the combustion turbines, subject to the constraint that the commitment schedule for all other units is fixed at the values determined in the unit commitment model. This recourse program is formulated as a quadratic program (QP) with the following structure:
   - an *objective function* in which cost is expressed as a

convex quadratic function of each unit's generation (for simplicity, it is assumed that any costs for committing the combustion turbines aside from the variable generation cost are negligible);

decision variables representing generation from each unit, including combustion turbines; and

constraints, consisting of bounds upon each unit's generation based upon their capacity, minimum run levels, and maximum ramp rates, considering the previous hour's generation.

The recourse model is solved using the commercial optimization package IMSL; for the few cases in which IMSL fails to find a feasible solution, we instead apply Lemke's algorithm.

By subtracting the first of the below quantities from the second, the cost of inaccurate forecasts is obtained:

1. the cost *if* the true load were known at the time that unit commitments were made, which is usually less than:
2. the actual cost of dispatching the system (from the recourse model based upon the true load) plus the start-up and other fixed costs of committing the units (from the Baldick model, based on the forecast load).

The procedure is described more specifically in Section 4. It is repeated for all days in the load forecast data base (in the case of the utilities below, about 440 days). It is then repeated for each of several different levels of the mean forecast error. This allows us to determine how the expected economic cost of inaccurate forecasts will change as the accuracy of load forecasts change.

In particular, we adopted the following procedure to simulate alternative levels of forecast accuracy. A revised forecast $L^{F\prime}$ is obtained as a weighted combination of the true load L^T and the utility's forecast L^F as follows:

$$L^{F\prime} = L^T + \lambda(L^F - L^T). \qquad (1)$$

The mean error (either root mean square or mean absolute percent) of the new forecast $L^{F\prime}$ is $|\lambda|100\%$ of the original forecast L^F. Thus, an increase in $|\lambda|$ simulates a worsening of accuracy, while a smaller $|\lambda|$ represents an improvement. In the analyses of this paper, λ is varied from -2 to +2 to simulate different degrees and directions of error. If $|\lambda|>1$, this results in forecasts with greater error than the original forecasts, while $|\lambda|<1$ implies less error (with the extreme case $\lambda=0$ implying a perfect forecast). A $\lambda<0$ changes the sign of the error; e.g., if the original forecast exceeded the true load, the revised forecast would understate it.

Our methodology is similar to that of Ranaweera et al. [9], with two important differences. First, their study assumed that forecast errors were distributed randomly and independently from hour to hour; our study is instead based on the actual distribution of errors, which for the two utilities we studied shows a high autocorrelation (0.96 in one case). Second, we consider a wider range of conditions, including several generation systems and two different utilities' loads.

3. CASE STUDY UTILITIES

Two utilities provided ANNSTLF forecasts and true loads. Table 1 presents summary data concerning those loads and forecasts. Several different generation systems were considered. All are based on the systems defined by Bard [10] and Shaw [11], with some modifications. Table 2, for instance, shows cost and capacity data for our modified Bard system. A minimum spinning reserve margin of 3%-5%, depending on the system, is required. The Shaw system has 13 units with 6425 MW of capacity. It differs from the Bard system in that marginal costs vary more among the units (*a* varies from 6.05 to 14.62 \$/MWh, several times its range in Table 2). Additional generation systems were defined by replicating the Shaw system (26 rather than 13 units), by modifying start-up and fixed costs of the Bard system (doubling and quadrupling them), and by altering the cost of combustion turbine energy. Our purpose in testing several systems is to determine how sensitive the value of forecast accuracy is to particular characteristics of the generation system.

Table 1. Load Data

System:	Northeastern utility; Jan. 1 - March 30, 1996	Southern utility; Jan. 1 - Dec. 10, 1995
Load factor:	0.70 (entire period) 0.88 (mean daily)	0.56 (annual) 0.84 (mean daily)
24 hour forecast error:	6.8% (RMSE) 5.4% (MAPE)	5.6% (RMSE) 3.9% (MAPE)

4. APPLICATION RESULTS

The economic cost of inaccurate forecasts was obtained by applying the models, load data, and generator data sets described above in 3 steps:

1. For each day in the load data base and for each value of λ considered (-2 to +2, in increments of 0.25), the Baldick [8] unit commitment model was used to create a commitment schedule based on the forecast loads.
2. For each hour of each day and for each λ, the scheduled units are dispatched against the actual load using the QP recourse model, yielding the actual dispatch cost. When summed over the 24 hour day, and then added to the start-up costs from the unit commitment model, this yields the actual daily production cost.
3. For each day and each $\lambda\neq0$, actual daily production costs are compared to the costs for $\lambda=0$ (zero forecast error), giving the cost increase due to forecast error. This cost is usually positive because the unit commitment has been optimized for the incorrect loads. (However, because unit commitment ignores uncertainty in loads, it may fail to identify the unit commitment that minimizes expected cost. Ideally, stochastic optimization should be used [12]; however, we instead simulate the prevailing utility practice of solving deterministic models. Because expected costs are not minimized by such models, it is possible for more accurate forecasts to yield *higher* costs; yet this occurs only occasionally in our simulations.)

In the subsections below, several groups of results are summarized. First, costs for three different commitment schedules for a single day are presented to illustrate why forecast inaccuracies inflict economic penalties. Then average results for four systems (Shaw and Bard systems for the

Table 2. Generator Data Set I (based on Bard [10])

Unit	Energy cost coefficients[a]			Start-up coefficients[b]			Max. ramp rate (up or down) [MW/hr]	Generation limits		Minimum up and down times	
	Linear term a [\$/MWh]	Quadratic term b [\$/(MWh)^2]	Fixed cost F [\$/hr]	Constant term c [\$]	Exponential term d [\$]	Time constant e [hr]		Min. run [MW]	Capacity [MW]	Up time [MW]	Down time [MW]
1	9.02	0.00226	820	2050	825	4	450	300	1000	5	4
2	7.65	0.00320	400	1460	650	3	220	130	400	3	2
3	8.75	0.00294	600	2100	950	4	215	165	600	2	4
4	8.43	0.00300	420	1480	650	4	185	130	420	1	3
5	9.22	0.00468	540	2100	900	3	310	225	700	4	5
6	7.05	0.0103	175	1360	750	2	75	50	200	2	2
7	9.12	0.00262	600	2300	950	4	350	250	750	3	4
8	7.76	0.00342	400	1370	550	3	190	110	375	1	3
9	8.16	0.00256	725	2200	950	4	400	275	850	4	3
10	8.15	0.00904	200	1180	625	2	100	75	250	2	1
11	8.33	0.0044	450	1760	780	3	250	175	575	3	3
CT	30.00	0	0	0	0	n.a.	none	0	4000	0	0

[a]Energy generation cost [\$/hour] = $aMW + bMW^2 + F$, where F is incurred only if the generator is committed and MW is the real power production.
[b]Start-up cost [\$] = $c + d*exp[-t/e]$, where t is time since last shutdown.

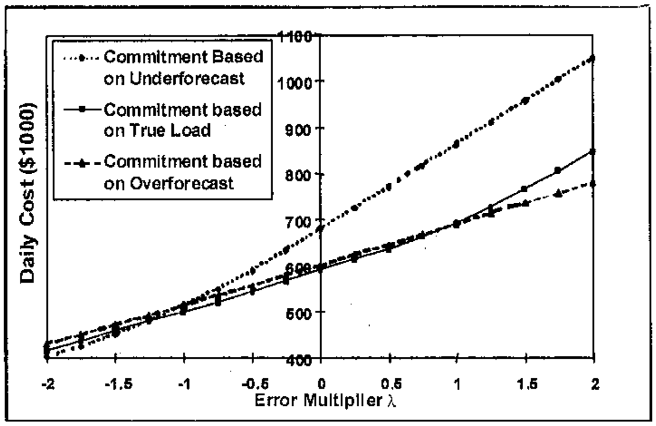

Fig. 2. Generation cost of three different unit commitment schedules under varying loads

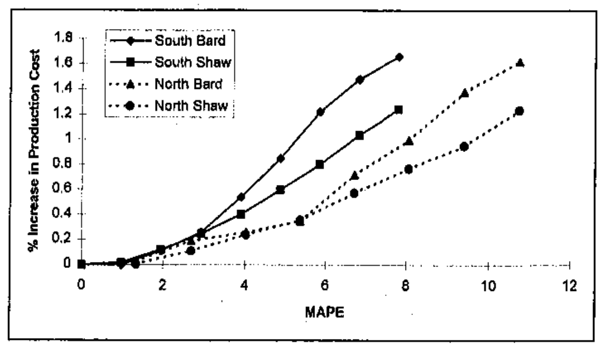

Fig. 3. Increase in variable production costs due to inaccurate load forecasts, averaged over year, for four systems

southern and northeastern utilities) are presented. The distribution of errors is subsequently discussed, including an analysis of the effects of underforecast versus overforecast errors. Finally, we present sensitivity analyses in which generator characteristics are varied, including system size, backup power costs, and start-up and other commitment costs.

4.1. Example of Inaccuracy Penalty

In this subsection, an illustration of the cost penalties that can occur as a result of forecast inaccuracy is presented. The southern Bard system is used as an example. On Jan. 2, 1995, that utility's load was overforecast by 15.5%, on average. As a result, unit commitments based on the forecast would not have minimized cost. Fig. 2 illustrates the costs resulting from three distinct unit commitment schedules as a function of the load. Alternative load profiles are represented by the scaling factor λ, where $\lambda = 0$ stands for the true load, λ=-2 represents a lower load (31% below the true load), and λ=+2 is a higher load (31% above the true load). For the high load, the schedule represented by the dashed line is best; for the low load, the dotted line is superior; and, finally, for the true load, the solid line schedule has the lowest cost. The schedules differ in the number and timing of start-ups and shut-downs.

The figure shows that choosing the wrong schedule, which would result from having the wrong forecast, can impose large cost penalties. In particular, forecasting a low load (λ=-2) but realizing a middle load (λ=0) results in costs being more than 15% higher than necessary (\$680,000 rather than the optimal \$591,000). Meanwhile, forecasting a high load (λ=+2) but realizing the middle load incurs a 1.9% penalty (a cost of \$602,000 rather than \$591,000).

4.2. Average Costs of Forecast Inaccuracy

In Fig. 3, the production cost increase due to forecast inaccuracy, averaged across all days, is shown as a function of MAPE for each of the four possible combinations of loads (southern vs. northeastern) and generation systems (Bard vs. Shaw). Each point is obtained from a distinct value of $|\lambda|$. (Because the results for λ and -λ are similar, they are combined, and only the average results for $|\lambda|$ are presented.)

The curves show, as anticipated, that production costs increase as accuracy worsens. At a MAPE of 3%, for instance, inaccurate forecasts inflate production costs by 0.1% to 0.25%. Increasing MAPE to 5% causes this cost penalty to rise to between 0.35% and 0.85%. All values differ significantly from zero ($p<0.01$). The southern system's penalties are double those of the northeastern system. This may be due in part to its lower daily load factor, which implies more start ups and shut downs over the day. Meanwhile, the Bard system, whose costs are more uniform across generating units, has higher penalties than the Shaw system for both the northeast and the south. Thus, these few examples show that penalties for forecast inaccuracy can vary substantially among utilities, just as reported in our survey [6].

To make these values more concrete, consider a 5000 MW (peak) system with a $20 mean variable production cost, a 0.6 load factor, and a 5% MAPE for its forecasts. The above figures imply that the economic loss due to inaccuracy is between $1.8M and $4.5M annually.

The curves of Fig. 3 are convex in most places, implying that a 1% worsening in accuracy imposes more of a penalty for a system that is already inaccurate. If MAPE is between 3% and 5%, a 1% change in MAPE changes variable production costs by 0.12% to 0.3%. For the system of the previous paragraph, this implies that improved accuracy is worth about $0.6M to $1.6M annually per 1% improvement, or $125 to $315/peak MW/year.

These estimates of the worth of better forecasts can be compared to values reported elsewhere. As mentioned in Section 1, four utilities we surveyed have estimated values of between $1.7 and $143/peak MW/yr for a 1% improvement in MAPE [6]. Using a similarly simple methodology, Wisconsin Electric [13] calculated a value of over $200/peak MW/yr. The lowest of those values are too conservative, because they did not consider the full range of benefits arising from avoiding overcommitments and purchases of backup power. The higher values are more comprehensive, but did not explicitly model how commitment schedules would be altered as accuracy improved, and so may be inaccurate.

In another study, Erwin et al. [14] report that a 200 MW improvement in forecast accuracy for the Southern Companies resulted in savings of $3.6M/year. At a load factor of 0.6 and an average variable cost of $20/MWh, this is equivalent to 0.17% reduction in cost (over $150 per peak MW) for each 1% improvement in forecast accuracy. Meanwhile, an estimate that cost will fall by roughly 0.2% per 1% improvement can be inferred from the statement that £10M/yr would be saved if accuracy was improved by 1% for the UK system in 1984 [15]. The latter estimate was based solely on fixed cost savings resulting from avoiding overcommitment.

Ranaweera et al. [9] report a slightly lower value of accuracy: decreasing MAPE from 5% to 3% for their 20 generator system lowers the penalty of inaccuracy from 0.44% to 0.27%. This translates to a 0.08% decrease in cost per 1% increase in MAPE, which falls below our range. A possible reason for their lower estimate is that they assume that forecast errors are uncorrelated from hour to hour; as a result, if too much capacity is committed in one hour because load was underforecast, that excess capacity might be useful in the next hour when an overforecast might have occurred. Thus, the schedule based on the forecast load might still be optimal,

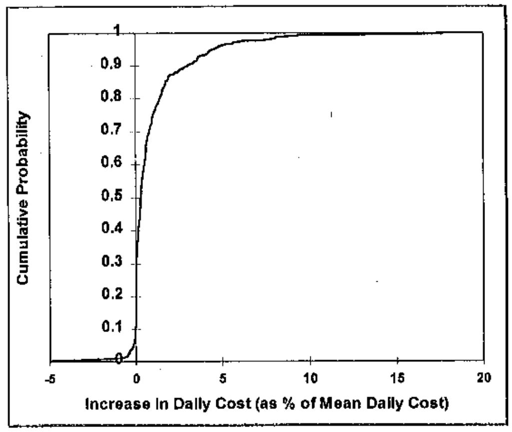

Fig. 4. Distribution of daily cost penalties due to inaccurate forecasts (Southern Shaw system)

or near optimal. However, actual errors are highly correlated in real systems; for instance, the hourly autocorrelation for the southern utility is 0.96. (That is, the correlation between the forecast error in hour t with the error in $t+1$ is nearly perfect.) Thus, for most days, hours will either be all overforecast or all underforecast, or nearly so. As a result, schedules based on forecasted loads are less likely to be optimal than assumed in [9], and the penalties should be larger than they calculated.

4.3. Distribution of Forecast Errors

The above annual averages mask large day-to-day variations in penalties. For the Southern Shaw system (MAPE = 3.9%), daily values range from zero or less (for 35% of the days) to a few values over 5% (Fig. 4), with a standard deviation of 1.1%. Meanwhile, their mean (Fig. 3) is 0.35%. As the distribution is highly skewed, less than 5% of the days account for half of the aggregate annual penalty.

(Note that negative penalties can occur, as in Fig. 4, because Lagrangian relaxation cannot guarantee a global optimum. Thus, the commitment obtained by the algorithm under an *erroneous* forecast might be better under the true load than the solution obtained assuming instead the *true* load. This cannot occur if global optimality is assured.)

Further analysis reveals that the bulk of the costs occur when loads are underforecast, when it becomes necessary to dispatch combustion turbines or buy spot power because insufficient thermal capacity has been committed. To demonstrate this, we divide the daily results into three groups:

1. Days in which 20 or more hours were underforecast (26% of days),
2. Days in which 20 or more hours were overforecast (35% of days), and
3. Days in which forecasts were a mix of over- and underforecasts (39% of days).

Fig. 5 shows for the Southern Shaw system the average cost penalty as a function of MAPE for each of the three subsets of days. For smaller MAPEs, under- and overforecasts impose similar penalties, but for medium and large errors,

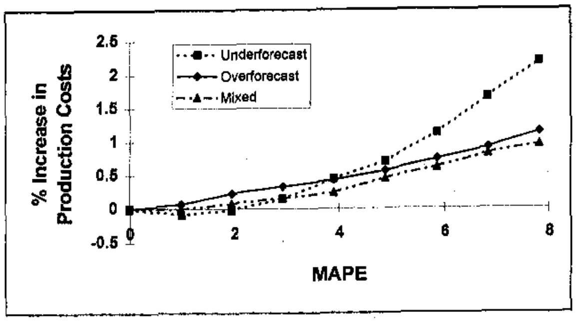

Fig. 5. Mean cost of inaccurate forecasts for three types of forecast errors (Southern Shaw system)

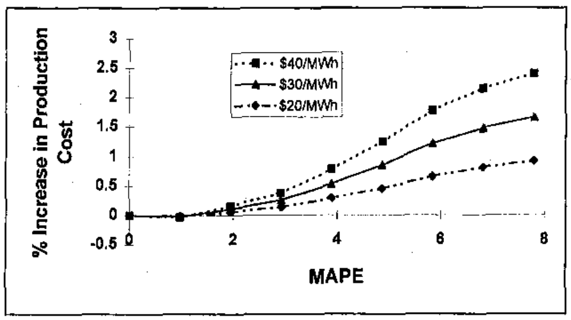

Fig. 6. Effect of back-up (turbine or spot market) costs upon the cost of inaccurate forecasts (Southern Bard system)

underforecast costs are more severe.

For other systems, underforecast costs tend to dominate for all MAPEs; for instance, for the Northeastern Bard system, for a MAPE equal to 5.4%, 65% of the annual penalty is incurred during underforecast days. In this situation, the cost of back-up power when underforecasts occur imposes more of a penalty than does incurring unnecessary start-up costs due to overforecasts. This result is consistent with the simpler analysis performed by the Southern Companies [summarized in 6], which estimated that underforecasts were responsible for two-thirds of the annual economic penalty.

Statistical analysis of the daily cost penalties for the $\lambda=1$ case reinforces the above conclusions. Various regression models including MAPE, level of loads, load factors, and number of errors of each type (over- and under-forecast) as independent variables were tested. The best models (by an adjusted R^2 criterion) accounted for over- and underforecasts separately and omitted the other variables. For the Southern Shaw system, the percent cost increase $\%C$ for a given day as a function of the MAPE for that day is:

$$\%C = 0.0954\, MAPE^{+} + 0.0200(MAPE^{-})^2,\ R^2 = 0.59 \qquad (2)$$
$$(0.0116) \qquad (0.0009)$$

Standard errors for the coefficients are given in parentheses. $MAPE^{+}$ is the MAPE that occurs if that day's loads are overforecast, on average, and is zero otherwise. $MAPE^{-}$ is the MAPE when loads are underforecast. For the Northeastern Shaw system,

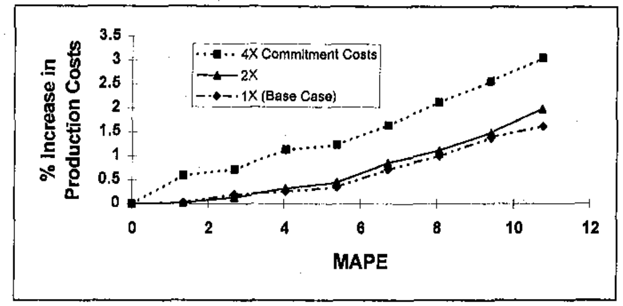

Fig. 7. Effect of alternative commitment costs upon the cost of inaccurate forecasts (Northeastern Bard system)

$$\%C = 0.00714(MAPE^{+})^2 + 0.0262(MAPE^{-})^2,\ R^2 = 0.53 \qquad (3)$$
$$(0.00234) \qquad (0.0025)$$

These equations show that a change from 3% to 5% in MAPE has a 1.5- to 4-fold larger impact for underforecasts than for overforecasts. This reinforces the result of Fig. 5 that underforecasts incur a larger penalty for these utilities. Also, confirming the results of Fig. 2, a 1% improvement in forecast accuracy when MAPE is between 3% and 5% drops costs by 0.06% to 0.42% (depending on whether over- or underforecasts are involved, and which system is considered).

4.4. Sensitivity to Generation System Characteristics

In this subsection, the effects of alternative assumptions concerning the generation system are examined.

The first sensitivity analysis concerns the size of the generating system. For the Southern Shaw system, this was tested by doubling the number of generating units and loads, and redoing the analysis. As a result, the economic penalty, as a fraction of annual production costs, increases by as much as 50%. However, for the Northeastern Shaw system, a similar doubling usually, but does not always yield an increase in the percentage cost penalty. Indeed, in another set of simulations, we found that for a simple system consisting of identical generating units, the size of the system (MW peak and number of units) does not significantly affect the percentage penalty. Therefore, we conclude that the use of relatively small systems (12 or 13 units) does not cause an upwards bias in our estimate of the economic penalty, expressed as a percentage of the total production cost.

The other sensitivity analyses examine two cost assumptions that directly contribute to the penalty: the expense of back-up (combustion turbine or spot market) power, in the case of underforecasts; and fixed costs associated with generator commitment, for overforecasts. The simple analysis performed by Southern Companies [summarized in 6] assumes that penalties are proportional to these costs; the results below show that the relationships are not necessarily that simple, but that they are still strong.

Fig. 6 shows the impact of varying the cost of back-up power from \$20 to \$40/MWh for the Southern Bard system. It turns out that penalties incurred during underforecast days are almost directly proportional to this cost; as underforecast costs dominate for that system, the results turn out to be very sensitive to this assumption.

The effect of varying fixed commitment costs is portrayed in Fig. 7 for the Northeastern Bard system. These costs include the fixed hourly costs F and start up cost terms c and d (Table 2). Three curves are shown, one representing the base case costs, and the other two modeling the situation in which commitment costs are doubled and quadrupled. The 2X cost case is only slightly greater than the 1X (base) case; for instance, for a MAPE of 5.4%, the base case penalty is 0.35%, which rises to 0.45% if commitment costs are doubled. This relatively modest effect is reasonable, since, as pointed out earlier, the costs of overforecasts (and, thus, overcommitment) are generally smaller than underforecast costs. In the 1X case, 14% of the annual penalty occurs during overforecast days, while in the 2X case that percentage rises to 41%. However, if commitment costs are quadrupled (case 4X), Fig. 7 reveals that the annual economic penalty is sharply higher; thus the penalty appears to be a nonlinear function of commitment costs.

5. CONCLUSION

Better information should lead to better decisions; for the case of more accurate short-term electric load forecasts, this paper has quantified the dollar value of improved unit commitment decisions. For a typical utility whose annual fuel costs amount to several hundred million dollars, a 1% reduction in the average forecast error can save hundreds of thousands or even millions of dollars.

Future research should consider benefits that flow from improving other types of short-term commitments that are becoming increasingly important in restructured power markets. Examples include:

- real-time prices,
- available transmission capability, and
- short-term spot sales.

Acknowledgement. This research was sponsored by the Electric Power Research Institute, with partial support by the National Science Foundation, grant ECS 9696014. We thank EPRI research managers D. Sobajic and M. Wildberger for their encouragement and comments, the members of the EPRI ANNSTLF Users' Group for their feedback and cooperation in providing data, R. Baldick of the University of Texas for providing the unit commitment model, and anonymous reviewers for their comments. The usual caveats apply.

References

[1] D. Zhai, A.M. Breipohl, F.N. Lee, and R. Adapa, "The Effect of Load Uncertainty on Unit Commitment Risk," IEEE Trans. Power Sys., 9(1), 1994, 510-517.

[2] D. Sobajic and Y-H. Pao, "Artificial Neural Net Based Dynamic Security Assessment for Electric Power Systems," IEEE Trans. Power Sys., 4(1), 1989, 220-227.

[3] T.M. Peng, N.F. Hubele, and G.G. Karady, "An Adaptive Neural Network Approach to One-Week Ahead Load Forecasting," IEEE Trans. Power Sys., 8(3), 1993, 1195-1203.

[4] A. Khotanzad, R.-C. Hwang, A. Abaye, and D. Maratukulam, "An Adaptive Modular Artificial Neural Network Hourly Load Forecaster and its Implementation at Electric Utilities," IEEE Trans. Power Sys., 10(3), 1995, 1716-1722.

[5] A. Khotanzad, R. Afkhami-Rohani, and D. Maratukulam, "ANNSTLF-Artificial Neural Network Short-Term Load Forecaster-Generation Three," IEEE Trans. Power Sys., in press.

[6] B.F. Hobbs, "Identification and Quantification of Benefits of Using ANNSTLF," Final Report, submitted to EPRI, October 1997.

[7] A. Debs, "Demand Forecasting for Operation Planning," Notes for Short Course on Modern Operations Planning, Decision Systems, Inc., Atlanta, GA, April 3, 1995.

[8] R. Baldick, "A Generalized Unit Commitment Model," IEEE Trans. Power Sys., 10(1), 465-475, 1995.

[9] D.K. Ranaweera, G.G. Karady, and R.G. Farmer, "Economic Impact Analysis of Load Forecasting," IEEE Trans. Power Sys., 12(3), 1388-1392, 1997.

[10] J.F. Bard, "Short Term Scheduling of Thermal-Electric Generation Using Lagrangian Relaxation, Oper. Res., 36(5), 756-766, 1988.

[11] J.J. Shaw, "A Direct Method for Security-Constrained Unit Commitment," IEEE Trans. Power Sys., 10(3), 1329-1342, 1995.

[12] A. Carpentier et al., "Stochastic Optimization of Unit Commitment: A New Decomposition Framework," IEEE Trans. Power Sys., 11(2), 1067-1073, 1996.

[13] T. Bilke, "Short Term Load Forecasting at Wisconsin Electric," Unpublished manuscript, WEPCO, Milwaukee, WI, Nov. 13, 1997.

[14] S.R. Erwin, J.T. Wood, and J.S. Griffith, "Operational-Planning Software Realizes Production Cost Savings," *Comput. Appl. Power*, 4(3), 55, July 1991.

[15] D.W. Bunn and E.D. Farmer, "Economic and Operational Context of Electric Load Prediction," in D.W. Bunn and E.D. Farmer, Comparative Models for Electrical Load Forecasting, Wiley, London, 1985.

Benjamin F. Hobbs (Member, IEEE) is Professor of Geography and Environmental Engineering at the Johns Hopkins University. Previously, he was Professor of Systems Engineering and Civil Engineering at Case Western Reserve University (CWRU). Between 1982 and 1984, he was a Wigner Fellow at Oak Ridge National Laboratory, and during 1977-1979 he was an Economics Associate at Brookhaven National Laboratory. He is a consultant to the FERC Office of Economic Policy, EPRI, and GRI.

Suradet Jitprapaikulsarn is a Ph.D. student in the Dept. of Systems, Control, & Industrial Engineering at CWRU. Mr. Jitprapaikulsarn obtained a B.S. in Mathematics from Chulalongkorn University, Bangkok, Thailand in 1991. Between 1991 and 1994, he worked for Computer Systems Consulting Co., Ltd. and International Software Factory Co., Ltd. in Bangkok.

Sreenivas Konda is a Ph.D. student in the CWRU Dept. of Systems, Control, and Industrial Engineering. Mr. Konda earned a B.S. degree in Mechanical Engr. from Osmania University, India in 1993, and a M.S. in Industrial Engr. from the University of Miami in 1995.

Vira Chankong (Senior Member, IEEE) is Associate Professor of Systems Engineering at CWRU. His research interests include large scale and multiobjective optimization, decision analysis and decision support systems, and systems methodology. He is senior author of "Multiobjective Decision Making: Theory and Methodology" (Elsevier-North Holland, 1983).

Kenneth A. Loparo is Professor of Systems Engineering and Mechanical and Aerospace Engineering at CWRU. His research interests include nonlinear and stochastic stability, nonlinear filtering, and control of stochastic systems, with applications to electric power systems.

Dominic J. Maratukulam is Manager, Power System Design at EPRI. He received the B.Tech. degree from the Indian Institute of Technology, Madras, India, a M.S. (material science) from the Univ. of Washington (1972), and M.Eng. (Electrical Engr.) from the Univ. of British Columbia (1974). Before joining EPRI in 1987, he was at Systems Control Inc., Palo Alto, California; B.C. Hydro, Vancouver, Canada; and Bharat Heavy Electricals, Trichy, India.