TECHNICAL ASSISTANCE IN THE DEVELOPMENT OF A JOINT NORTH AND SOUTH CAROLINA ELECTRICITY DEMAND FORECASTING MODEL

Prepared for

The National Regulatory Research Institute
2130 Neil Avenue
Columbus, Ohio 43210

CRA Report #413

December 1978

This report was prepared by Charles River Associates, Incorporated for The National Regulatory Research Institute (NRRI) under contract No. EC-77-C-01-8683 with the U. S. Department of Energy (DOE), Economic Regulatory Administration, Division of Regulatory Assistance. The opinions expressed herein are solely those of the contractor and do not necessarily reflect the opinions nor the policies of either the NRRI or DOE.

The NRRI is making this report available to those concerned with state utility regulatory issues since the subject matter presented here is believed to be of timely interest to regulatory agencies and to others concerned with utilities regulation.

The NRRI appreciates the cooperation of the North Carolina Utilities Commission and the South Carolina Public Service Commission with the contractor in preparing this study.

TABLE OF CONTENTS

Introduction

On June 15, 1978, the National Regulatory Research Institute (NRRI) contracted with Charles River Associates Incorporated (CRA) to provide on-site services in support of the Public Staff of the North Carolina Utilities Commission and the staff of the South Carolina Public Service Commission in developing long-range electric energy and peak load forecasts. This memorandum briefly reports on the conduct and achievements of the NRRI-sponsored work.

The project was designed and conducted with several considerations in mind. First, both agencies had a practical need for the improved load forecasting procedures developed in the course of the project. Second, the project emphasized the role of consultants as a teaching resource. Accordingly, we developed approaches to working with the North and South Carolina agency staffs that featured close interaction between CRA consultants and agency staff members in all phases of the work. Agency staff participated in key model development decisions, performed the main work of data assembly and preparation, and implemented and documented the load forecasting model with advisory assistance from the CRA consulting team. Third, the transfer value of the project to other regulatory agencies lies primarily in the lessons of this project regarding effective interaction between consultants and commission staffs to enhance the load forecasting capability of the commissions and to provide them with the ability to apply that capability to help develop forecasting procedures they can apply, maintain, and improve.

Background

The Public Staff of the North Carolina Utilities Commission was established by the North Carolina General

Assembly in 1975. Among its responsibilities, as established by statute in June 1977, is the annual filing of 15-year energy and peak demand forecasts for each of the state's utilities. The primary purpose of these forecasts is to serve as an independently derived check on the load forecasts that support expansion plans filed by the electric utilities. Building upon earlier work of the North Carolina Commission staff, the first such independent forecast was presented on December 15, 1977.

In its application to NRRI, the Public Staff requested assistance in developing long-term load forecasting methods. Assistance was sought in three areas: the development of durable forecasting procedures that could be updated and applied annually at low cost, the development of econometric forecasting methods, and a lesser degree of assistance in the development of noneconometric methods to serve as a check on the principally econometric procedures.

The request of the North Carolina Public Staff was joined by a similar request from the staff of South Carolina PSC. While the South Carolina PSC lacks a statutory obligation to provide independent load forecasts, such forecasts are routinely requested as an internal matter.

The proposed project had two features that made it an especially good candidate for NRRI support. First, the two staffs expressed an active interest in the joint estimation and use of forecasting models. Since the two largest utilities in North Carolina (Carolina Power and Light Company and Duke Power Company) are also the two largest South Carolina utilities, there were substantial payoffs to cooperative forecasting. Second, while neither staff had a great deal of experience in forecasting, both were able to make substantial commitments of high-quality staff for a three-month period during the summer. Hence, it was reasonable to expect that substantial progress could be made on model development and staff training during a brief but intense period of on-site consultation.

Contract Scope

The NRRI-CRA contract called for a series of five sessions of consultation in Raleigh over a period of two and a half months. During these sessions the consultant was to work with the commission staffs in developing techniques and models for forecasting residential and industrial load growth and for forecasting system peaks. The models were to be estimated for a utility to be chosen during the course of the project. If project resources allowed, some consultant time was to be reallocated to assist in developing noneconometric forecasting techniques.

As the project unfolded, scope was expanded to include assistance in commercial sector modeling. Less off-site consulting was required than had been anticipated. Instead, the staffs elected to have more on-site support. Because of the excellent progress made by the staffs in estimating the econometric models it proved possible to devote four consultant mandays to assisting the North Carolina Public Staff in noneconometric forecasting techniques.

Conduct of the Contractual Effort

The first step in the contract was to meet with the North and South Carolina staffs and an NRRI representative to delineate the roles of the consultant, NRRI, and the various members of the Carolina staffs. The functional requirements of the Carolina staffs (in terms of model geographic scope, forecast period, methodological constraints, and so forth) were also determined. We next reviewed the existing forecast methodology employed by the staffs to determine whether it would be appropriate to concentrate efforts on only some demand sectors.[1] It was determined that we should go forward initially with the residential and industrial sectors, then add commercial energy and peak demand as time permitted. It was further decided to concentrate modeling efforts on Duke Power with the expectation that the Carolina staffs would subsequently extend the methodology to other major utilities.

The principal focus of the first session was on data. Various modeling strategies and associated data requirements were discussed and compared with data already collected or available from Duke, state agencies, and outside sources. From this discussion evolved a general modeling strategy and a detailed data collection and database compilation strategy to be pursued by the staffs. In broad outline, the strategy was to develop a highly structured, end-use model of residential energy sales and a two-sector model of industrial sales that would give separate treatment to electricity use by textiles manufacturers. These models would be estimated using annual data. If resources permitted, an annual model of electricity sales to Duke's commercial customers would be developed. Lastly, a structural approach to peak load forecasting would be mapped out toward the end of the project. This approach would have the objective of forecasting Duke's peak demand in a manner fully consistent with the forecasting of Duke's energy sales.

The specific strategy for residential energy use was to develop a model capable of forecasting the annual consumption of electricity by customers belonging to each of Duke's residential rate classes: standard service, standard plus water heat service, and all-electric service. Although lacking

---

[1] The methodology employed by the staffs prior to this project is described in Appendix A.

full appliance detail, this model would nonetheless account for differences in consumption between rate classes due to differential rates of ownership of electric space heat and electric water heat. Distinct per-customer average annual utilization equations would be developed for air conditioning, water heating, and space heating; all other uses of electricity would be lumped into a single "baseload" consumption equation. In general, per-customer utilization rates would be forecast to depend on average annual household income, electricity prices and, where appropriate, weather. Moreover, the model would be designed to forecast customer selection of one or another rate class on the basis of the relative prices of electricity and fossil fuels as well as the availability of new natural gas hookups to Duke's residential customers. To effect this modeling strategy, historic estimates of average annual per-customer consumptions of electricity by end use and rate class would have to be developed from monthly utility sales data. This would be a major task. Historic data on residential electricity use determinants would also have to be developed.

It was decided that the forecasting of Duke's industrial energy sales would benefit most from two structural model improvements. One would be the separate forecasting of electricity sales to textile establishments and all other electricity sales. This would be beneficial both because textile production is highly electricity-intensive and because the textiles industry in the Carolinas is much more nationally oriented than other Carolina industries. The second improvement would be to relate electricity use by Duke's industrial customers to average annual industrial employment in the counties comprising the Duke service territory. Previously Duke's industrial sales had been related to measures of statewide industrial activity, although Duke's service area is very different in its industrial makeup from the Carolinas as a whole. The development of an activity database specific to the service territory would enable the more accurate forecasting of Duke's industrial sales. The major prerequisite to carrying out the industrial strategy was the collection of time series data on county-level employment by industry in North and South Carolina and its assembly into a Duke-specific database.

The strategy for commercial model development was quite modest. If time permitted (as it eventually did), a simple aggregate model of commercial energy sales would be estimated. The basic difference between this model and ones previously used by agency staff would be its use of territory-specific commercial employment data as a measure of the size of the commercial sector as opposed to statewide retail sales. The model would also rely on annual, not monthly, data.

Two essentially different peak demand modeling approaches exist. The simpler is to relate peak load directly to activity variables, electricity prices, and perhaps peak weather. This approach tends to be highly unsatisfactory in that it ignores the fundamental relationship between peak demand and energy consumption, namely, that the peak is a function of the distribution of electricity consumption over time. The consultant advised the development of an alternative, preferred approach. This approach, which itself has a number of variants, is to derive the system peak demand forecast from forecasted system energy sales. The particular strategy to be pursued if time permitted would be to develop a derived peak demand model that was sensitive to major changes in the composition of system energy sales. Such a model's forecasts would thereby reflect the radically different load factors of, for example, the industrial and residential sectors.

The Carolina staffs had excellent success in quickly developing data. As a consequence, in the second session, which was scheduled to involve only a review of collection efforts and further work on model specification, we began estimation of some of the algorithms of the residential model. The residential model specification called for separate forecasts of water heating, space heating, air conditioning, and baseload power uses. However, historic data on the amounts of the end-use specific consumption of energy did not exist, and thus they had to be developed using the semirigorous load separation procedure described below.

The residential load separation task had the objective of best identifying historic air conditioning, water heating, space heating, and baseload energy sales for Duke's three residential rate classes given limited information. In Duke's case the basic available information was monthly data on average electricity sales per customer by rate class. The load separation procedure therefore relied on monthly rate class sales data, various reasonable assumptions about the underlying composition of those sales, and secondary information. Assumptions, where they were imposed, were of three kinds. One category specified months of each year in which electricity is used for air conditioning and space heating. A second category dealt with the usual monthly distribution of annual electricity consumption for certain end uses. The third type of assumption related the magnitude of average end-use consumption in one rate class to that in others. Secondary information included rate class level appliance stocks survey data. By its nature, the load separation exercise was in part a search for a reasonable set of assumptions that fit the monthly sales data well. The procedure could therefore be complicated by identifiable factors seriously affecting historic sales for which no quantifiable information exists. Alterations in data series were identified and had to be diagnosed. As a result, the load separation task required more time and effort than was originally anticipated. Major unquantified changes over time in the Duke territory housing stock, and in the characteristics of residential electric water heaters, appeared to be responsible for much of the difficulty experienced in deriving sensible time series estimates of historic end-use consumption patterns. Because of the dependence of the residential modeling strategy on historic end-use consumption data, successful completion of the load separation task became the critical pacing element for the project as a whole. Ultimately, data of sufficient accuracy and stability for econometric estimation were constructed.

Apart from delaying residential model development, the slow progress of the residential load separation task caused another problem because of the unexpectedly large amount of data processing that it required. The problem was simply a severe lack of data processing capacity to devote to industrial, commercial, and peak load model development. This problem, which was apparent as early as the second work session, was never fully overcome. This problem was reduced by obtaining additional data processing capacity in South Carolina for industrial and peak model estimation and by using a local interactive time-sharing system for estimation of the commercial model. As a response to this problem and in order to more efficiently use the time of the two consultants and several Carolina staff members involved, we abandoned the sequential model development strategy that had been originally adopted in anticipation of lags in industrial and commercial data preparation. Beginning with the third session, somewhat specialized staff/consultant teams worked independently on the various energy and demand models.

The third and all subsequent sessions were devoted to model estimation. The general pattern of these sessions was that the Carolina staffs would present results that had been obtained since the previous session. These would be reviewed and the consultants would make suggestions concerning respecifications that might improve the results. In some cases, these suggestions would imply redoing previously acceptable portions of the model. In other cases, they required collection of additional data. The new specifications would then be tried. This process was iterated until satisfactory results were obtained. The number of iterations varied, depending on the number of session days, computer turnaround, and how quickly we converged on acceptable results. In the context of the project, an "acceptable result" was a usable forecasting model judged better than available alternatives in terms of several broad criteria. Acceptance criteria included consistency of the model's parameter estimates with the prior knowledge and reasonable expectations of CRA consultants and agency staffs, and the quality of test statistics for statistically derived model parameters.

Until well into August most of the analysis was done with North Carolina data. It took considerably longer to assemble the South Carolina data since little preparatory work had been done on it. It became clear that the project could not be completed by its original completion date (September 1) without abandonment of estimation of the South Carolina portion of the Duke model. The decision to stretch the project schedule allowed time for staff to complete the data collection for the South Carolina estimates and estimate the South Carolina equations while preserving the availability of the CRA consulting team to provide assistance and critical review.

Noneconometric Models

The North Carolina Public Staff had requested assistance on noneconometric forecasting methods. Noneconometric forecasts are produced by a separate group within the Public Staff and serve as an information check on the econometric forecasts. This is not a primary activity for this group, and they can make little time available to produce the forecast. Historically, trend analysis has been used as a forecasting methodology.

CRA provided an expert in noneconometric methods to work with this group. He reviewed the existing methodology, gave a seminar on more advanced forecasting methods, and made practical suggestions on how the group's methodology could be improved within existing resource restraints. The content of this course is briefly outlined in Appendix B.

Project Results

Primary results achieved were as follows:

- A complete econometric model estimation database specific to the Duke Power Company service territory was completed.
- Strategies for forecasting the model's explanatory variables were identified.
- A residential model was developed that permits separate forecasts of the saturations of electric water heating and space heating, of per-customer utilizations of these heating systems, and of per-customer power consumptions for air conditioning and lighting and miscellaneous purposes. Subsequent to the project, satisfactory preliminary results have been obtained by agency staff in the development of a similar model for the Carolina Power and Light Company service area.
- A commercial sales model was satisfactorily estimated. The model is more highly aggregated than the residential and industrial model. The parameter estimates are good, and it should forecast rather well.
- A two-sector industrial model with substantial power was developed. This model treats textiles separately from other industries.

- A fully operational system peak load model was not developed by the project's conclusion. However, the project had developed a preliminary model, and agency staffs were advised on several model development strategies to follow. They have subsequently estimated a peak load model that reflects those strategies and is consistent with the above energy sales models.
- Most important, the project experience materially enhanced the agency staffs' knowledge and capability in load forecasting. Moreover, the commissions now possess a much more highly structured and accurate set of econometric energy sales and peak demand forecasting procedures than was previously the case. These procedures should be fairly durable, so that in the future agency staffs can concentrate more effort on developing sound forecast assumptions as opposed to reliable forecast procedures.

Suggestions Bearing on Future Projects

There are a number of lessons that we learned from this project that bear on the smooth functioning of other projects of a similar nature. These are:

1. The project was cost effective. Consultant development of a similar load forecasting model would have been somewhat more expensive on a real resource cost basis (that is, including the cost of Carolina staff personnel and state-provided computer services) and materially more expensive on a budgetary basis.
2. A very substantial amount of learning by agency staffs took place. Model building is more an art than a science. Hence, this type of learning by the state staffs is best achieved by close and continuous integration of consultant and staff personnel.

3. Coordination of the two states was not a problem in this case. Key reasons for this success were the strong common purpose of modeling the same utilities, the mutual dependence implied by the need for data from both states, professional respect among individuals on the staffs, and the lack of conflicting model output or methodology requirements.
4. Substantially full-time dedication of the staffs and the relative absence of conflicting activities were instrumental to maintaining project momentum. While the time schedule of the project turned out to be unrealistically tight, this tightness imposed a useful discipline.
5. It turned out that frequent two- and three-day working sessions were more effective than longer but fewer sessions. While this was substantially due to computer turnaround problems, it would probably have been preferable even had a fully interactive time-sharing system been used.
6. Good early planning of data needs is very important due to the time lags involved in getting data from utilities and from other state agencies.
7. Problems with computer turnaround are a very serious impediment to successful on-site support of model development activities.
8. Given the high technical quality of personnel assigned to this project by the states, the level of consultant support provided by this contract was sufficient. However, considerably more support would have been required if state staffs had been less qualified to take initiatives on their own.
9. It would probably be useful to maintain consultant involvement through the first exercising of the model

in a forecasting mode. This would make the consultant available for advice on compilation of the forecast database and on any model modifications required to correct deficiencies that become apparent in the forecast simulation.

10. The project affirmed the value of tight scheduling combined with schedule flexibility in the face of unanticipated problems in the execution of certain project tasks. Despite the delays that were encountered, the project was successfully completed by September 20, 1978, within 13 weeks of project start and within three weeks of the original deadline.

Appendix A
FORECAST METHODS PREVIOUSLY EMPLOYED BY THE CAROLINA STAFFS

Prior to this project, the Carolina staffs had made some progress in the development of econometric techniques for forecasting electric energy sales and peak demands. In particular, they had previously forecast utility sales separately for each major class of service: residential, commercial, and industrial. Class-of-service energy sales were typically forecast using single equation models in which aggregate monthly sales were econometrically related to a number of explanatory variables. Explanatory variables included, for example, average electricity prices, population, total retail sales, the Federal Reserve Board index of industrial activity, and heating and cooling degree days.

There were three major drawbacks to these energy models. First, the population and economic data on which they were based were not well matched to the utility territory in question.

Second, they were basically nonstructural. For example, residential energy's primary dependence on the number of residential customers and its secondary dependence on per-customer behavior were not explicitly modeled. The modeling of Duke Power's industrial sales, as a second example, suffered from the fact that sales to textiles manufacturers were lumped together with all other industrial sales. This was a problem for several reasons. First, Duke's sales to the textiles sector are a very large fraction of its total industrial sales. Further, electricity use patterns are very different for textiles firms than for other industrial firms, and the Carolina textiles industry is growing at a different rate from the rest of the region's economy.

The third drawback of these energy models was their reliance on monthly, as opposed to annual, data. The use of monthly data to estimate long-range forecasting models imposed burdensome data and estimation requirements on the Carolina staffs. Development of models using monthly data requires, in particular, that purely seasonal phenomena be dealt with statistically or otherwise, even though these seasonal phenomena are of no particular interest to the long-range energy forecaster. Early in the project the consultant advocated model development primarily using annual data, since annual data provide adequate information for most long-range forecasting purposes.

The procedures previously used to forecast peak demands involved the direct modeling of peak as a function of independent explanatory factors including electricity prices, economic activity, and weather. The major limitation of this approach is that it ignores the fundamental relationship between peak demand and energy consumption, namely, that peak is a function of the distribution of electricity consumption over time.

Appendix B
NONECONOMETRIC ENERGY DEMAND FORECASTING METHODS

A short course on noneconometric methods of forecasting electricity sales was presented to this group by a member of the CRA staff. Introductory material included a comparison of econometric and noneconometric methods for forecasting electricity sales.

Noneconometric models rely only partially, if at all, on formal econometric procedures. The simplest example of a noneconometric forecast is a sales trend extrapolation. Often, trend extrapolations of other variables such as appliance saturations are used to adjust the extrapolations of the variables of interest. These procedures suffer from the disadvantages of certain econometric time series models in that resulting forecasts are largely a result of historical trends. They do, however, offer the advantage that most of the data are readily available. In situations where time and staff constraints are problems, the trend extrapolations represent a reasonable first step at implementing a noneconometric forecast.

The preferred noneconometric approach is often characterized as an inventory and/or employment based approach. This approach includes the forecasting of the basic determinants of energy demand, namely, the stock of appliances in residential and commercial sectors and a measure of output for industrial sectors.

For example, an inventory based model of the residential sector can be developed by using Bureau of Census forecasts of population, and by employing estimates of future household formation rates to determine a forecast of the future stock of dwelling units. Dwelling units forecasts automatically provide forecasts of space heating equipment for most states. These forecasts along with Census of Housing forecasts on the stock of dwelling units for the base year, and some assumed depreciation rate for the existing dwelling unit stock, provide estimates of the stock of dwelling units added during any year in the forecast period. Using information from engineering studies on average energy required for each electric space heating system (perhaps by housing type) or information from utilities, where it exists, along with an estimated saturation rate allows one to forecast electricity required for electric space heating. Some uncertainties arise in this forecasting approach for reasons similar to those of econometric models. That is, the number of households, electric space heating saturation, and choice of housing type must be estimated. The advantage of the approach, however, is that energy use requirements of electric space heating systems is based on much more detailed information concerning requirements of individual appliances. This consideration is especially important when one is considering the impacts of appliance standards that will become effective in 1980. Tracking the stock of space heating in appliances in an accounting framework that keeps information by vintage of the stock allows easy implementation of such policy changes. Another advantage is the ability to incorporate estimated voluntary changes in efficiency of appliances that arise in response to increasing fuel prices. This inventory-based approach requires forecasts of the stocks of individual appliances for residential and commercial sectors. This approach has been applied at Oak Ridge National Laboratory for both of these sectors; thus, documentation exists to guide other organizations in implementing this modeling approach.

The output-based modeling of the industrial sector consists of relating energy use per unit of output by industry or by employment within the industry and using forecasts of industrial output or employment and electricity intensity to drive forecasts of future electricity sales to industrial clients. The effects of conservation policies, which have been analyzed on a sectorial basis for the U.S. Department of Energy, can then be used to adjust the energy use per unit of output or energy use per employee to incorporate future changes in production efficiency in the industrial sector.

In summary, while the noneconometric approach outlined above suffers from some of the same uncertainties concerning forecasts of relevant exogenous variables such as number of households, the focus of the model approach on appliance- or output-related energy requirements allows long-run policy forecasts that are well grounded on the results of a voluminous engineering literature.