1985

Bootstrapping the Probability Distribution of Peak Electricity Demand

Michael R. Veall

8861

ISSN: 0318-725X
ISBN: 0-7714-0622-3

RESEARCH REPORT 8506

BOOTSTRAPPING THE PROBABILITY DISTRIBUTION OF PEAK ELECTRICITY DEMAND*

by

Michael R. Veall

April, 1985

ABSTRACT

For effective capacity planning, an electric utility requires an estimate of the probability distribution of future maximum demand, rather than simply a point prediction of expected peak. This paper proposes a method of obtaining this using the bootstrapping technique of Efron (1979) to forecast the peak demand of an actual utility, Ontario Hydro. While the technique is constructed from the standard procedure of forecasting a future variable using regression coefficients and known values for the right-hand side variables, it is modified to allow for uncertainty in these independent variable forecasts as well.

*Thanks are due to D. Fretz and G. Green for excellent research assistance and to A. Gregory, I. McLeod, F. Trimmell and A. Ullah for valuable advice. The assistance of the Centre for the Study of International Economic Relations at the University of Western Ontario, Ontario Hydro and the Social Sciences and Humanities Research Council of Canada is acknowledged.

I. Introduction

Most electric utilities make their capacity decisions based on a forecast of expected peak demand plus some relatively arbitrary allowance for a reserve margin. However, given the large costs of either too much or too little capacity and the high degree of uncertainty with respect to future usage, it would seem worthwhile to calculate formally an optimal probability of shortage (e.g. Telson, 1975) and to set target capacity accordingly. This appears to be seldom done, perhaps because of the lack of reliable estimates of the required probabilities. This paper attempts to solve this difficulty by estimating the entire probability distribution of peak demand, conditional on current information.

The standard method of estimating such a probability distribution would be to estimate a least-squares regression for some forecasting model and to calculate a forecast for peak as the product of the estimated coefficient vector with the predicted values for the right-hand side variables. The probability distribution of this forecast could then be calculated by repeated application of the standard formula for forecast confidence intervals (e.g. Johnston, 1972, pp. 154-155).

There are serious difficulties with this procedure. As Feldstein (1970) discusses, the conventional confidence interval formula does not allow for uncertainty in the predicted values of the exogenous variables, which may be extremely important. While Feldstein shows how to calculate correct standard errors in this situation, confidence intervals are much more difficult, as he notes. Even if the estimated coefficient vector and the exogenous variables are each normally distributed, their product the forecast will not be, as is assumed in the standard approach. Moreover, not only does the conventional method require normality of the forecast estimate, it requires normality of the variable to be forecast. In this case, the variable to be forecast is peak demand, which typically will have a complicated distribution as it is a maximal order statistic. There is also an added distributional complication in that the forecasts are based on an equation expressed in logarithms and the predicted values are obtained by taking exponentials. These issues are discussed in more detail below, with the conclusion favouring a nonparametric approach which can handle uncertainty in the exogenous variable forecasts.

Such a nonparametric technique is "bootstrapping", as proposed by Efron (1979) and also described in Efron and Gong (1983). It can be simply explained using the example of estimating the median of a sample of independent observations drawn from the same, perhaps unknown, underlying distribution. The probability distribution of the estimate is in general a complicated calculation except for a few cases, for example, if the observations are normally distributed. The bootstrap however, replaces this calculation simply by employing a large number of artificial samples that have been created by drawing randomly and with replacement from the original sample. The probability distribution of the sample median can then be approximated by the histogram of the medians of the artificial samples.

The bootstrap of regression forecasts is straightforward and is summarized in Section II. The artificial samples are now based on random draws from the residuals and the technique is directed toward the distribution of the forecast error. This extension has been applied elsewhere, for example in Freedman and Peters (1984a, 1984b), although their applications differ from that here in that they focus on the standard errors and not on the probability distribution, and they assume that the forecasts of the exogenous variables are certain.

Section III describes the actual application, starting with the estimation of the underlying model and a discussion of diagnostics. Initially the predicted values for the right-hand side variables are assumed to be certain and the standard forecast and confidence intervals are calculated. These results are then compared with those of the bootstrapping procedure. Then the basic bootstrap is extended to allow for uncertainty in the right-hand-side variables, both by formulating a two-equation recursive model and by allowing some future values to be generated randomly. Section IV presents the conclusions of the research and suggests areas for future study.

II. Bootstrapping

Bootstrapping is a technique which replaces theoretical assumptions and complex algebraic calculations with a large number of stochastic simulations. The heart of the idea is to use a computerized pseudo-random number generator in artificial resampling, and then to use these artificial samples to calculate an empirical probability distribution for the target variable. In the case of forecasting, begin by considering a standard regression model

$$Y = X\beta + \epsilon \qquad (1)$$

where Y is a TX1 vector of observations on the dependent variable to be forecast, X is a TXk matrix of observations on the exogenous variables, β is a kX1 vector and ϵ is a TX1 vector of independent disturbances. The β can be estimated by OLS on observations 1,...,T provided the rank of X is k. Denoting this estimate by $\hat{\beta}$, the standard method of forecasting y_{T*}, the dependent variable in period T*, is:

$$\hat{y}_{T*} = x'_{T*}\hat{\beta} \qquad (2)$$

where x'_{T*} is a row vector of exogenously given values for x at time T*.

To bootstrap, draw T times randomly with replacement from the elements of the least-square residuals $\hat{\epsilon}$ to create an artificial residual set $\hat{\epsilon}^1$. Repeat this N times to construct $\hat{\epsilon}^i$, i=1,...,N. Then each of N artificial samples can be constructed as

$$Y^i = X\hat{\beta} + \hat{\epsilon}^i, \qquad i=1,...,N \tag{3}$$

that is, by adding the sets of artificial residuals to the fitted values of the original regression. The distribution of $\hat{\beta}$ can then be estimated using the empirical distribution of $\hat{\beta}^i$'s, that is, the histogram of the coefficient estimates from the artificial samples.

An alternative in this case is to assemble T row vectors (y_t, x_t) and then draw randomly from these T times with replacement to create the N artificial samples (as discussed in Efron and Gong, 1983). It can be seen that this differs from the above in that there is no attempt to hold the design matrix X constant (which is shown most dramatically if one of the bootstrapped samples (Y^i, X^i) has by chance an X matrix of rank less than k so that $X'X$ is singular and the OLS $\hat{\beta}^i$ cannot be calculated). Therefore, while there is little formal reason to prefer either of the two bootstrap methods, the first is used exclusively here because it seems desirable that each bootstrap sample have the same time-series history except for the random disturbance ϵ. In addition, in the theoretical limit as $N \to \infty$ under the i.i.d. assumption, it is the first method which will yield the conventional OLS estimate of the variance-covariance matrix of $\hat{\beta}$ and not the second (Efron and Gong, 1983).

To bootstrap the forecast errors, randomly sample N more residuals from $\hat{\epsilon}$, call these $\hat{\epsilon}^{i}_{T*}$ and then calculate N "simulated actuals" by adding these to the original forecast:

$$y^{i}_{T*} = x'_{T*}\hat{\beta} + \hat{\epsilon}^{i}_{T*}, \qquad i=1,\ldots,N \tag{4}$$

For this step to be valid, it must be assumed that the distribution of the disturbances has not changed over time, implying they must be homoskedastic. The simulated forecast error is then the difference between (4) and the bootstrap prediction:

$$\text{Forecast error}^{i} = y^{i}_{T*} - x'_{T*}\hat{\beta}^{i} \tag{5}$$

and therefore the distribution of y_{T*} conditional on the regression forecast can be estimated as the empirical distribution of these forecast errors centred around $\hat{y}_{T*}$, the actual forecast. This method can be further modified, to deal with random x'_{T*} and other problems, as will be discussed in the context of the application in Section III.

The technical justification of the bootstrap is that the bootstrap distribution converges to the true distribution asymptotically as both N, the number of bootstraps, and T, the number of observations, grow large (Bickel and Freedman, 1981). In the case here, this occurs essentially because the empirical distribution of the residuals converges asymptotically (with T) to the distribution of the true disturbance vector.

In finite samples the case for the bootstrap is less clear.[1] One reason to use it in the context examined here is the lack of alternatives: conventional analytic approaches to estimating the probability distribution of a forecast variable all require normality assumptions or at least some known probability distribution for the disturbances. But there are several reasons in the current example that suggest such a parametric approach will be inappropriate.

First, the estimated model will be in logs. If the true model is:

$$w_t = \beta_o z_t^{\beta_1} e^{\epsilon_t} \quad (6)$$

then

$$y_t = x_t \beta + \epsilon_t \quad (7)$$

where $y_t = \ln w_t$, $x_t = [1 \vdots \ln z_t]$, and $\beta = [\ln \beta_o \vdots \beta_1]$.

In general, an unbiased forecast of y_t cannot be translated into an unbiased forecast of w_t by simple exponentiation (see Goldberger, 1968). However, percentiles such as the median or the 95 per cent confidence intervals of y_{T*} can be so transformed in order to compute counterparts for w_{T*}. Bootstrapping will essentially do this for the entire estimated distribution of w_{T*} and hence an unbiased forecast can be calculated directly and the bias of the simple exponentiation procedure estimated.

The problem is further complicated as y_t and hence ϵ_t may not have the normal distribution. In Veall (1983), it is pointed out that peak demand is the maximal order statistic of all the momentary demands, so that even if the momentary demands are normal, the peak will not be. If there are many "peak-eligible" momentary demands and all of these "parents" are identically distributed, it is shown further that maximum demand will have the extreme value distribution. But because these assumptions seem unrealistic, it is possible that the peak probability distribution will be some unknown mixture, perhaps of the normal and the extreme value.[2]

Even if the distribution of y_t is known, there may be further problems. Note that:

$$\text{Forecast error} = y_{T*} - x'_{T*}\hat{\beta}$$
$$= x_{T*}\beta + \epsilon_{T*} - x'_{T*}(X'X)^{-1}X'Y$$
$$= \epsilon_{T*} - x'_{T*}(X'X)^{-1}X'\epsilon \qquad (8)$$

Initially assume that both x_{T*} and X are nonstochastic. In that case the true forecast error is the difference between the disturbance in period T* and a linear combination of past disturbances. In the conventional case, all disturbances are assumed to be normal so, as the set of normal variates is closed under addition, the forecast is normally distributed and confidence intervals are straightforward to compute. But if the distribution of the disturbances is non-normal as has been suggested is possible, the calculation of the distribution of (8) as a sum of random variates may be extremely complex and lead to no convenient distribution.

The final reason is that if x'_{T*} is also random, the second term of (8) becomes a product of random variables. Even if the strong assumptions are made that both ϵ and x'_{T*} are normal, their product would not be. When this point is considered along with the other points made above, it is clear that a nonparameteric approach may be desirable. Bootstrapping has the advantage both of being nonparametric and of allowing x'_{T*} to be random.

There is also considerable evidence suggesting that the bootstrap may perform well in finite samples. In terms of the estimation of the distribution of a statistic, Beran (1982) shows that, under regularity conditions, the bootstrap estimate will in general converge more quickly than the normal approximation, because the latter makes no allowance for bias. (The bootstrap will do exactly as well as the estimated first-order Edgeworth expansion.) Also with as few as 15 observations, the bootstrap has performed well in test situations when compared to an exact distribution which is either known or inferred from a much larger sample (Efron, 1979, 1982). In terms of forecasts specifically, Freedman and Peters (1984a) used Monte Carlo techniques to study the estimation of forecast standard errors and found that while bootstrapped standard errors were slight underestimates, they were far closer to the true standard errors than their analytic rivals.

III. Application to Ontario Hydro Peak

Most utilities forecast peak demand by simply dividing a forecast of average demand by an assumed "load factor". Here a slightly more general relationship is explored using a regression in logarithms[3]. The OLS results, using data for the period 1963-1982[4] for the East System of Ontario Hydro (which typically comprises about 90 per cent of the total provincial demand) are:

$$\log(PEAK_t) = \underset{(.1257)}{.8953} + \underset{(.0141)}{.9482} \log(AMW_t) + \hat{e}_t \qquad (9)$$

$R^2 = .9960$ Durbin-Watson statistic = 1.7709

where standard errors are in parentheses, $PEAK_t$ is peak demand in megawatts (mW) and AMW_t is average demand over the year, also in mW. The constant load factor assumption would imply that the log (AMW_t) coefficient be one, a hypothesis that can easily be rejected at the 5 per cent significance level.

The first step was to test the normality of the residuals of (9). Monte Carlo research by Huang and Bolch (1974) suggests that the preferred test is that of Shapiro and Wilk (1965) applied to the OLS residuals. This test statistic is .951, while a value below .905 would be required to reject the null hypothesis of normality at the 5 per cent level. The test based on moments proposed by Kiefer and Salmon (1983) yields a similar result[5]. Although the normality hypothesis cannot be rejected, the power of these tests in a sample of only 20 observations is questionable. In this sense a nonparametric approach is conservative, guarding against potential non-normality at this stage while also treating the other possible distributional problems discussed above, namely the prediction bias due to taking exponentials of a log forecast and the difficulty if x'_{T*} is also random.

It is also necessary to test for serial correlation and heteroskedasticity, either of which alone would invalidate the confidence intervals calculated both by standard methods and by bootstrapping. Five different diagnostic tests were employed, each using the 5 per cent level. As noted, the Durbin-Watson statistic is 1.77 (critical value: 1.41) and hense cannot reject the null hypothesis of no first-order autocorrelation. This result is also supported by a Godfrey (1978) t-statistic of -.28 (critical value: -2.11) for which the alternative hypothesis is either an AR(1) of MA(1) process. The chi-square test of White (1980) gives a value of 2.11 (critical value: 5.99) while that of Breusch and Pagan (1979) yields a statistic[6] of 3.37 (critical value: 3.84), so both cannot reject the null hypothesis of homoskedasticity. The Engle (1982) test against an alternative of autoregressive heteroskedasticity (ARCH) effects gives a chi-square statistic of .97 (critical value: 3.84) and hence cannot reject the null hypothesis[7].

Now that the necessary assumptions for bootstrapping have been tested, its application to forecasting can be discussed. A medium-term prediction for 1990 has been selected for illustrative purposes, and initially it is assumed that with certainty, 1990 average demand will be 14500.0 megawatt-hours, consistent with a 3 per cent growth rate from the 1982 level.[8] Using the coefficient estimates above, this implies a forecast of exp(9.9807) or 21605 mW for 1990 annual peak. Calculating the confidence intervals of this forecast using the method of Salkever (1976), the Student's t distribution 90 per cent confidence interval in logs is (9.9406, 10.02080) or (20756, 22490) in mW. In other words, there is an estimated 95 per cent probability that a capacity of 22490 would be sufficient in 1990.

When this estimate is bootstrapped the results are similar. This might be expected, given that there was little evidence of non-normality in the residuals and there has as yet been no allowance for a random x'_{T*}. As can be seen in Figure 1, with the number of bootstraps set at 1000,[9] the 95 percent probability point is now 22332mW, only slightly lower than that from the standard method. The other probability points listed in the bootstrap figure are related similarly to those generated by the Salkever procedure.

The other interesting point is that the estimated bias (i.e., the difference between the exponential of the log forecast and the empirical mean of the bootstrap distribution) is small, only 24mW or about .1 per cent of the forecast. This is still a little larger than the 6mW predicted by the Goldberger (1968, p. 469) adjustment.

The forecast method so far has assumed knowledge of x'_{T*}, the average demand for 1990. To allow random x'_{T*}, a model for average demand is specified and estimated using OLS over 1963-1982 as

$$\log(AMW_t) = \underset{(1.7291)}{1.2331} - \underset{(.0495)}{.3887}\log(P_t) + \underset{(.1038)}{.4336}\log(Y_t) + \underset{(.0046)}{.0293}\,TIME_t + \hat{\eta}_t$$

$$R^2 = .9985 \qquad \text{D-W statistic} = 1.7526 \qquad (10)$$

where standard errors are in parentheses, P_t is the real average price of electricity,[10] Y_t is total real income for the province of Ontario, and TIME is a linear time trend.[11] If weather variables (the log of cooling degree days or the log of heating degree days[12]) are added to (10) as well as the log of the real price of natural gas, all three coefficients are neither individually nor jointly significant at the 5 per cent level. Also, when the diagnostic tests previously applied to (9) are applied to (10), the null hypothesis of normality, homoskedasticity and no serial correlation cannot be rejected at the 5 per cent level.

Regressions (9) and (10) can be bootstrapped recursively. This is done by randomly drawing year numbers uniformly and with replacement from 1963-1982 and taking the residuals from each of (9) and (10) for that year as a matched pair $(\hat{\epsilon}_t^i, \hat{\eta}_t^i)$. Then for each observation of artificial sample i, AMW_t^i is calculated using the coefficient estimates from (9), values of P_t, Y_t and $TIME_t$ plus a bootstrapped residual $\hat{\eta}_t^i$. $PEAK_t^i$ is then generated recursively using AMW_t^i in (10) plus the matching bootstrapped residual, $\hat{\epsilon}_t^i$.

To forecast average usage using the standard method, it is now necessary to obtain a forecast of 1990 real income. To use the trend annual rate of growth in our sample of 4.9 percent seems wildly optimistic; on the other hand, to assume the virtual zero growth of the last four years seems too pessimistic. In the end, it was decided to weight the recent experience more heavily and assume 2 per cent per year, which is about one per cent lower than suggested by the Economic Council of Canada (1983). This yields a 1990 forecast of approximately 41.8 billion 1971 dollars. Assuming a constant real price of electricity, the point prediction from (10) is 14848 megawatt-hours, about 350 mW.h higher than assumed above. This leads to a higher peak forecast of 22096 mW, an increase of about 500 megawatts. But more striking is that when the peak distribution is estimated by the described recursive bootstrapping, the 95 per cent point is 23399 mW, or more than 900 megawatts higher than with the basic nonrecursive bootstrap. The additional uncertainty modeled by the recursive bootstrap apparently makes a substantial difference in the tail. In contrast, the estimated bias of the forecast due to exponentiation is virtually zero (see Figures 2 and 3).

An alternative method of estimating the distribution of the forecast error is to employ the Feldstein estimate of the variance-covariance matrix. If the distribution of both the forecast from (10) and coefficient estimates from (9) are both normally distributed, then the forecast error cannot also be normally distributed but this nonetheless could be a valuable approximation. Under this Feldstein estimate based on normality, the 95 percent point can be calculated as 23622 mW, slightly higher than but reasonably close to the bootstrap estimate.

The final experiment was to allow for subjective uncertainty in the forecasts of the stochastic variables, price and income. The actual observations on these variables were not altered (so that all results are conditional on the same "history" of the right-hand side) but in calculating the simulated actuals as in (4), it was recognized that the future x_{T*} would not be known but would be random. Accordingly, x_{T*} was assumed to be normally distributed and an x^i_{T*} was calculated for each bootstrap run by adding a normal pseudo-random deviate to each of the price and income forecasts in log form. This new x^i_{T*} was then used to replace x_{T*} in (4) for the price forecast. For the price element, a standard deviation of 10 per cent was assumed while for the income forecast, a standard deviation of 7.1 per cent was assumed for one run, and 10 per cent for another.[13]

As might be expected, neither modification shifts the estimated probability distribution of peak very much but both flatten it. In each case, the estimated bias due to taking exponentials is about .1 per cent. But using the 7.1 per cent standard deviation on income, the 95 per cent point for peak climbs another 900 mW to 24287 mW. With a 10 per cent standard deviation, the 95 per cent point becomes 24608 mW (see Figure 4). Clearly this addition of uncertainty using subjective priors is somewhat *ad hoc*,[14] but it does illustrate the importance of relaxing the assumption that the forecasts of the independent variables are certain.[15]

Again the Feldstein method may be used as an approximation to yield alternative estimates. Using the 7.1 per cent standard deviation as the subjective forecast of income, the 95 per cent point is estimated as 24499 mW; with a 10 per cent standard deviation this estimate becomes 24782 mW.

While in both of these latter situations the Feldstein method again predicts a slightly higher 95 per cent point than the bootstrap, in all cases the results are similar (a similarity which extends to the entire probability distribution). This evidence helps substantiate the validity of the bootstrap conclusions and suggests for these cases, that the normality assumption using the Feldstein variance-covariance matrix may lead to slight overestimates of capacity requirements. However, another possibility has not yet been considered, and that is if the priors on the forecasts of the exogenous variables are also non-normal. In the electricity planning process for example, there might be the possibility of a large change in the pricing regime, such as using marginal cost as the basis for pricing rather than the lower average cost.

To give an example of this, assume (rather artificially) that there are two equally likely plans for electricity pricing in the future. Under the first plan prices would fall 14 per cent by 1990; under the other they would rise by the same percentage. These prices have been calculated to correspond to the assumptions of Figure 4--namely that the prior on percentage price change has a mean of zero and a standard deviation of 10. The bootstrap-estimated distribution is presented in Figure 5. Using the rest of the assumptions of Figure 4, the 95 per cent point is now 25054 mW, which is greater rather than less than the 24782 mW result of the Feldstein method. This difference apparently occurs because the empirical distribution is no longer approximately normal, with a Kiefer-Salmon statistic of 17.11 rejecting the null hypothesis of normality at the 5 per cent level (critical value: 5·99). The Kiefer-Salmon statistic for Figure 4 with the normal priors is only 2.15.

To conclude this section, it seems worthwhile to discuss possible qualifications. First, is 1000 sufficiently large for N, the number of bootstraps? To test this, examine Table 1, which gives various upper-tail probability points for $n = 100$, 1000 and 10000 for probably the most complex example, that of Figure 5, with the recursive bootstrap with subjective uncertainty on both incomes and prices, with the price forecast assumed to be a discrete binomial. The agreement is reasonably close, except for the 99 percent point for N = 100, which is somewhat smaller than its counterparts. The conclusion is that N = 1000 appears to be adequate.

Second, a more refined bootstrapping approach would include bias-corrected confidence intervals (Efron, 1981, p. 146), which allow for the possibility that the empirical distribution of estimated forecast errors might be biased and not have mean zero. However, as all such biases are in the order of .1 per cent of the forecast, there seems to be no need to pursue the bias corrected technique.

IV. Summary and Conclusions

In summary, there are two main points. First, while these results are only illustrative, they suggest estimation of the probability distribution of peak will be valuable. A 95 per cent probability of capacity being sufficient in 1990 does not seem excessive, but that requires under some assumptions 2500 mW or more additional capacity than would be indicated by the regression forecast. This is a substantially greater margin than standard confidence intervals would indicate, yet it still should be regarded as a conservative estimate because it does not take into account possible shifts in model structure.

Second, as might be expected, it appears to be very important in estimating the probability distribution of peak demand to allow for uncertainty in the forecasts of the independent variables. This is illustrated by the recursive bootstrap results and also by the results of a bootstrapping procedure which allows for subjective uncertainty in these forecasts.

Finally, the work suggests a possible goal for future research. Ultimately, it would be desirable to be able to estimate the probability distribution not only of peak demand but of the entire load duration curve (which is drawn by taking any number y and mapping it against the length of time that demand exceeds y). As this is the principal tool for the optimal planning of capacity mix (see for example Crew and Kleindorfer, 1979 and Rowse, 1980), it would then enable these calculations to be done in a stochastic as opposed to a deterministic context. The result would permit empirical calculation of the effect of uncertainty on the choice of different types of capacity with different fixed and marginal costs (as in Ellis, 1980) and hence potentially contribute to a framework that would improve the entire range of capacity decisions.

Footnotes

[1]The probability distribution estimated by the bootstrap does turn out to be closely related to the Bayesian posterior associated with an uninformative symmetric Dirichlet prior (Efron, 1981). This similarity has led Rubin (1981) to suggest that other, more reasonable priors be used to construct a Bayesian bootstrap.

[2]As mentioned in the application, peak demand is measured in terms of its logarithm, which in no way invalidates this discussion. Note also that the assumption of a large number of identically distributed parent densities is probably more realistic for a factory with distinct shifts (as in Veall, 1983) than for aggregate data as here. If the parents are identically distributed but there are few peak-eligible periods, the standard order statistic approach could be used, but this requires knowing both the form of the parent density and the exact number of peak-eligible demands. Galambos (1978) discusses the difficulties non-identical parents may cause. Gumbel (1958, pp. 184-187) describes Barricelli's generalization. This treats a particular departure from the identical parent population which states formally how the resulting peak distribution may be between the normal distribution and the extreme value distribution.

[3]While a serious attempt has been made to produce reasonable forecasts, the following analysis tends to use somewhat simplified models. This is to minimize the institutional detail so that the bootstrap results are illustrated as clearly as possible and also to ensure that only publicly available data are required. All data used in the following are available upon request.

[4]When this equation was estimated using either 1941, 1946 or 1950 as starting points, stability testing using the Chow F-test (Johnston, 1972, p.199) suggested a break in 1963. Further Chow testing uncovered no further break in the 1963-1982 period. The sample ends in 1982 because provincial income estimates, which prove to be important in subsequent analysis, are only available to 1982 at the time of writing.

[5]As might be expected, the residuals between the actual values and exponents of the corresponding fitted values appear to have a positively skewed distribution and tests reject normality. This suggests that there will be at least some bias in the point prediction calculated by taking the exponential of the prediction in logs.

[6]The Breusch and Pagan test is conducted against an alternative that the variance is a linear combination of the right-hand side variables of (9), plus a random component.

[7]Nonparametric analogues to the serial correlation and heteroskedasticity tests were also performed. The Geary (1970) residual sign change test could not reject the no autocorrelation hypothesis nor could an obvious modification of it using the absolute value of residuals reject the hypothesis of no ARCH effects. Homoskedasticity could not be rejected by a nonparametric test based on the Spearman rank coefficient and described by Johnston (1972, p. 219). All these tests were also at the 5 per cent level.

[8]This is consistent with published forecasts conditional on 1982 information but is slightly high compared to more recent ones (Ontario Hydro, 1983).

[9]The uniform pseudo-random number generator used in this research is that of Wichmann and Hill (1982), Applied Statistics algorithm 183. When later in the paper a normal random generator is needed, this is combined with Applied Statistics algorithm 111 (Beasley and Springer, 1977).

[10]This is calculated as average revenue, which is clearly a major simplification given the variety of Hydro rates. However these rates tend to move together closely and Berndt (1978) has pointed out for a single rate system, the difference between the marginal rate and the average rate does not usually affect estimates greatly. The price is deflated using the annual Gross Provincial Product deflator for Ontario.

[11]The formulation of an equation for $\log(AMW_t)$ naturally suggests re-estimating equation (9) with Two-Stage Least Squares to test whether in that equation $\log(AMW_t)$ is correlated with the disturbance. Using the instruments implied by (10), the 2SLS results are:

$$\widehat{\log(PEAK_t)} = \underset{(.1258)}{.8953} + \underset{(.0141)}{.9482} \log(AMW_t) \qquad (9')$$

virtually identical to estimates (9). Following Hausman (1978), a specification test consists of taking $\widehat{\log(AMW_t)}$ from (10) and adding it as an additional variable in (9). When this is done the coefficient is .0086 with a t-statistic of .0023, which is clearly unable to reject the null hypothesis that $\log(AMW_t)$ is uncorrelated with the disturbance in (9).

[12]Cooling degree days are a well-known measure of the need for air conditioning and are defined as $\sum_{i=1}^{365} (TEMP_i - 18^\circ C)$, $TEMP_i \geq 18^\circ C$, where $TEMP_i$ is the average daily temperature in degrees Celsius for day i. Heating degree days are a measure of heating requirements and are defined as $\sum_{i=1}^{365} (18^\circ C - TEMP_i)$, $TEMP_i \leq 18^\circ C$.

[13]It was felt that some idea of the uncertainty could be obtained from the forecast uncertainty using the Salkever method on a regression of income on trend. Even though the 4.9 per cent growth rate was rejected as too high a forecast, the estimate of 7.1 per cent as the standard error of the forecast was used here. The 10 per cent estimate is completely arbitrary.

[14]The uncertainty of forecasts of independent variables can be treated less arbitrarily if the distribution of the independent variable can be assumed constant over time, as in the case of weather variables. For example, in related work, models were estimated to forecast monthly peaks, particularly for July which usually contains maximum summer usage. In the July model, temperature variables proved to be important and these were bootstrapped by drawing randomly from past observations. In principle, there is no reason to restrict this draw to the sample but the entire historical record could be used.

[15]The 90 per cent confidence intervals for Figure 4 correspond to an average growth rate confidence interval of [1.2,4.1] over the 1982-1990 period. This corresponds to a range of about 5000 mW. Capacity of that amount would be currently valued at more than 1 billion Canadian dollars, based on the average accounting value of currently operating Hydro facilities.

The estimates from Baughman, Joskow and Kamat (1979, pp. 242-243) suggest this cost estimate is conservative, as the 1985 estimate of capital costs for the cheapest type of peaking component, gas turbines, is 250 U.S. dollars per kW or about 1.6 billion Canadian dollars for 5000 mW. Any other type of capacity has a capital cost of at least twice that amount.

References

Baughman, M. L., Joskow, P. L. and D. P. Kamat (1979), Electric Power in the United States: Models and Policy Analysis, M.I.T. Press, Cambridge.

Beasley, J. D. and S. G. Springer (1977), "The Percentage Points of the Normal Distribution," Applied Statistics, 26, 118-121.

Berndt, E. R. (1978), "The Demand for Electricity: Comment and Further Results," Department of Economics Resources Paper No. 38, The University of British Columbia, August.

Beran, R. (1982), "Estimated Sampling Distributions: The Bootstrap and Competitors," Annals of Statistics, V. 10, No. 1, 212-225.

Bickel, P. J. and D. A. Freedman (1981), "Some Asymptotic Theory for the Bootstrap," Annals of Statistics, 9, 1196-1217.

Breusch, T. S. and A. R. Pagan (1979), "A Simple Test for Heteroskedasticity and Random Coefficient Variation," Econometrica, 47, 1287-1294.

Crew, M. A. and P. R. Kleindorfer (1979), Public Utility Economics, St. Martin's, New York.

Economic Council of Canada (1983), The Bottom Line: Technology, Trade and Income Growth, Ottawa.

Efron, B. (1979), "Bootstrapping Methods: Another Look at the Jackknife," Annals of Statistics, 7, 1-26.

Efron, B. (1981), "Nonparametric Standard Errors and Confidence Intervals," Canadian Journal of Statistics, 9, 139-172.

Efron, B. (1982), The Jackknife, The Bootstrap and Other Resampling Plans, SIAM, Philadelphia.

Efron, B. and G. Gong (1983), "A Leisurely Look at the Bootstrap, the Jackknife and Cross-Validation," The American Statistician, 37, 36-48.

Ellis, R. (1980), "Optimal Pricing and Investment Decisions by Electric Utilities Under Demand Uncertainty," M.I.T. Energy Lab Working Paper No. MIT-EL 80-048 WP, November.

Engle, R. F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflations," Econometrica, 50, 987-1008.

Feldstein, M. S. (1971), "The Error of Forecast in Econometric Models when the Forecast-Period Exogenous Variables are Stochastic," Econometrica, 39, 55-60.

Freedman, D. A. and S. C. Peters (1984a), "Using the Bootstrap to Evaluate Forecasting Equations," Technical Report No. 20, Department of Statistics, University of California, Berkeley, January.

Freedman, D. A. and S. C. Peters (1984b), "Bootstrapping an Econometric Model: Some Empirical Results," Journal of Business and Economic Statistics, 2, 150-158.

Galambos, J. (1978), The Asymptotic Theory of Extreme Order Statistics, Wiley, New York.

Geary, R. C. (1970), "Relative Efficiency of Sign Changes for Assessing Residual Autoregression in Least Squares Regression," Biometrika, 57, 123-127.

Godfrey, L. G. (1978), "Testing Against General Autoregressive and Moving Average Error Models When the Regressors Include Lagged Dependent Variables," Econometrica, 46, 1293-1302.

Goldberger, A. S. (1968), "The Interpretation and Estimation of Cobb-Douglas Functions," Econometrica, 36, 464-472.

Gumbel, E. J. (1958), Statistics of Extremes, New York, Columbia University Press.

Hausman, J. A. (1978), "Specification Tests in Econometrics," Econometrica, 46, 1251-1271.

Huang, C. J. and B. W. Bolch (1974), "On the Testing of Regression Disturbances for Normality," Journal of the American Statistical Association, 69, 330-335.

Johnston, J. (1972), Econometric Methods, McGraw-Hill, New York.

Judge, G. G., W. E. Griffiths, R. C. Hall and T. C. Lee (1980), The Theory and Practice of Econometrics, Wiley, New York.

Kiefer, N. and M. Salmon (1983), "Testing Normality in Econometric Models," Economics Letters, V. 11, No. 1-2, 123-127.

Ontario Hydro (1983), Load Forecast Report 831212, Ontario Hydro, Toronto.

Rowse, J. (1980), "Intertemporal Pricing and Investment for Electric Power Supply," Bell Journal of Economics, 11, 143-165.

Rubin, D. B. (1981), "The Bayesian Bootstrap," Annals of Statistics, 9, 130-134.

Salkever, D. S. (1976), "The Use of Dummy Variables to Compute Predictions, Prediction Errors and Confidence Intervals," Journal of Econometrics, 4, 393-397.

Shapiro, S. S. and M. B. Wilk (1965), "An Analysis of Variance Test for Normality (Complete Samples)," Biometrika, 52, 591-611.

Telson, M. L. (1975), "The Economics of Alternative Levels of Reliability for Electric Power Generation Systems," Bell Journal of Economics, 6, 679-694.

Veall, M. R. (1983), "Industrial Electricity Demand and the Hopkinson Rate: An Application of the Extreme Value Distribution," Bell Journal of Economics, 14, 427-440.

White, H. (1980), "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," Econometrica, 48, 817-838.

Wickman, B. A. and I. D. Hill (1982), "An Efficient and Portable Pseudo-random Number Generation," Applied Statistics, 31, 188-190.

Table 1

A Comparison of Probability Points for Different Numbers of Bootstraps Using the Recursive Bootstrap Model with Uncertain Forecasts of Independent Variables[1]

Probability Point	N = 100	N = 1000	N = 10000
		(in megawatts)	
.5	21934	22006	22030
.75	23431	23446	23358
.90	24564	24514	24360
.95	25100	25054	24936
.99	25645	25779	25994

[1] All runs are based on (9) and (10) and are for the year 1990. The forecasts of the real electricity price and of the level of real income in 1990 are consistent with 0 and 2 per cent growth respectively. However in each bootstrapping run, a pseudo-random normal deviate with standard deviation .1 has been added to the 1990 log levels of real income. Real price is forecast as a binomial with .5 probability of a price decline of 14 per cent and .5 probability of a price increase of 14 per cent. The results for N = 1000 correspond to Figure 5.

Figure 1

Basic Bootstrap of Annual Peak Demand, 1990[1]

BOOTSTRAP FORECAST, ANNUAL, 1990
****

THE FORECAST FOR 1990 IS 21604.54536

THERE IS A PROBABILITY OF	.50	THAT PEAK WILL NOT EXCEED	21618.64	MEGAWATTS
THERE IS A PROBABILITY OF	.75	THAT PEAK WILL NOT EXCEED	21865.46	MEGAWATTS
THERE IS A PROBABILITY OF	.90	THAT PEAK WILL NOT EXCEED	22108.80	MEGAWATTS
THERE IS A PROBABILITY OF	.95	THAT PEAK WILL NOT EXCEED	22331.80	MEGAWATTS
THERE IS A PROBABILITY OF	.99	THAT PEAK WILL NOT EXCEED	22800.45	MEGAWATTS

FREQUENCY DISTRIBUTION OF SIMULATED FORECAST ERRORS AROUND THE FORECAST PEAK LOAD DEMAND

PEAK LOAD	FREQ
19845.03518	1
19915.11341	0
19985.19163	1
20055.26986	0
20125.34808	2
20195.42631	0
20265.50453	0
20335.58276	3
20405.66098	6
20475.73921	4
20545.81743	10
20615.89566	4
20685.97388	12
20756.05211	1
20826.13033	3
20896.20856	17
20966.28678	17
21036.36501	25
21106.44323	26
21176.52146	32
21246.59968	41
21316.67791	33
21386.75613	62
21456.83436	66
21526.91258	67
21596.99081	69
21667.06903	61
21737.14726	31
21807.22548	72
21877.30371	70
21947.38193	43
22017.46016	40
22087.53838	28
22157.61661	25
22227.69483	11
22297.77306	7
22367.85128	2
22437.92951	4
22508.00774	6
22578.08596	3
22648.16419	3
22718.24241	7
22788.32064	11
22858.39886	3
22928.47709	2
22998.55531	1
23068.63354	2
23138.71176	0
23208.78999	0
23279.36921	1

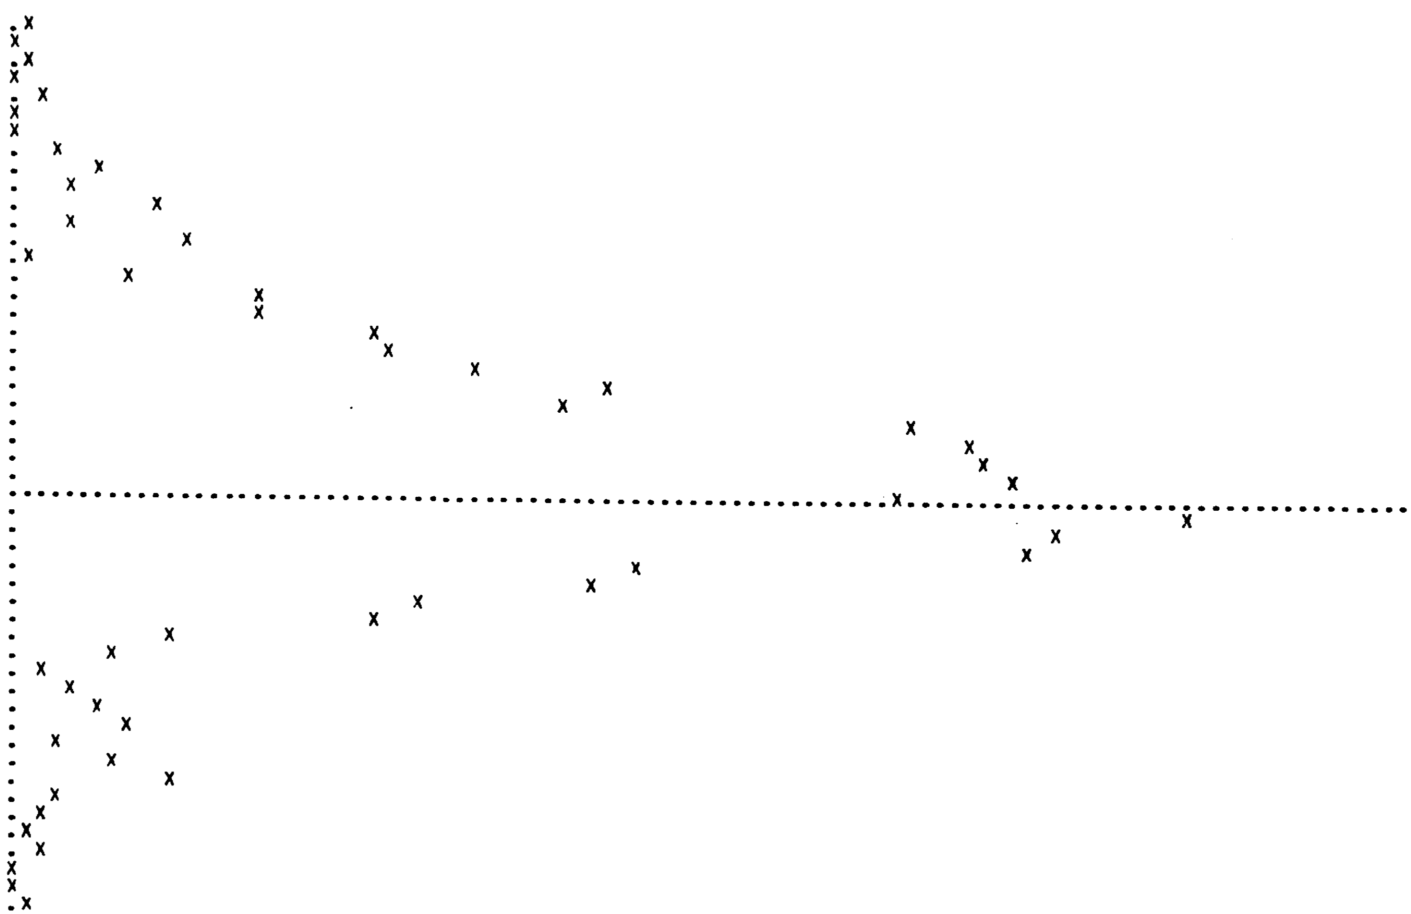

---

[1] Based on (9) and a forecast for average megawatts of 14,500.

Figure 2

Basic Bootstrap of Annual Average Demand, 1990[1]

BOOTSTRAP FCRECAST, ANNUAL, 199C
****

THE FORECAST FOR 1990 IS 14847.95153
THERE IS A PROBABILITY OF .50 THAT ENERGY WILL NCT EXCEED 14839.21 MEGAWATT-HOURS
THERE IS A PROBABILITY OF .75 THAT ENERGY WILL NCT EXCEED 15153.16 MEGAWATT-HOURS
THERE IS A PROBABILITY OF .90 THAT ENERGY WILL NCT EXCEED 15397.77 MEGAWATT-HOURS
THERE IS A PROBABILITY OF .95 THAT ENERGY WILL NCT EXCEED 15598.57 MEGAWATT-HOURS
THERE IS A PROBABILITY OF .99 THAT ENERGY WILL NCT EXCEED 15903.25 MEGAWATT-HOURS

FREQUENCY DISTRIBUTION OF SIMULATED FORECAST ERRORS AROUND THE FORECAST OF ENERGY

ENERGY	F?EQ
13436.86943	1
13495.50244	0
13554.13546	0
13612.76847	2
13671.40143	1
13730.03450	0
13788.66751	2
13847.30052	5
13905.93354	13
13964.56655	8
14023.19956	10
14081.83258	14
14140.46559	11
14199.09860	22
14257.73162	23
14316.36463	23
14374.99765	19
14433.63066	40
14492.26367	35
14550.89669	39
14609.52970	53
14668.16271	41
14726.79573	43
14785.42874	71
14844.06175	52
14902.69477	40
14961.32778	52
15019.96079	50
15078.59381	44
15137.22682	49
15195.85983	42
15254.49285	43
15313.12586	29
15371.75887	23
15430.39159	19
15489.02490	14
15547.65791	10
15606.29093	17
15664.92394	10
15723.55695	7
15782.18997	7
15840.82298	2
15899.45600	6
15958.08901	0
16016.72202	4
16075.35504	1
16133.98805	0
16192.62106	1
16251.25408	0
16310.88709	2

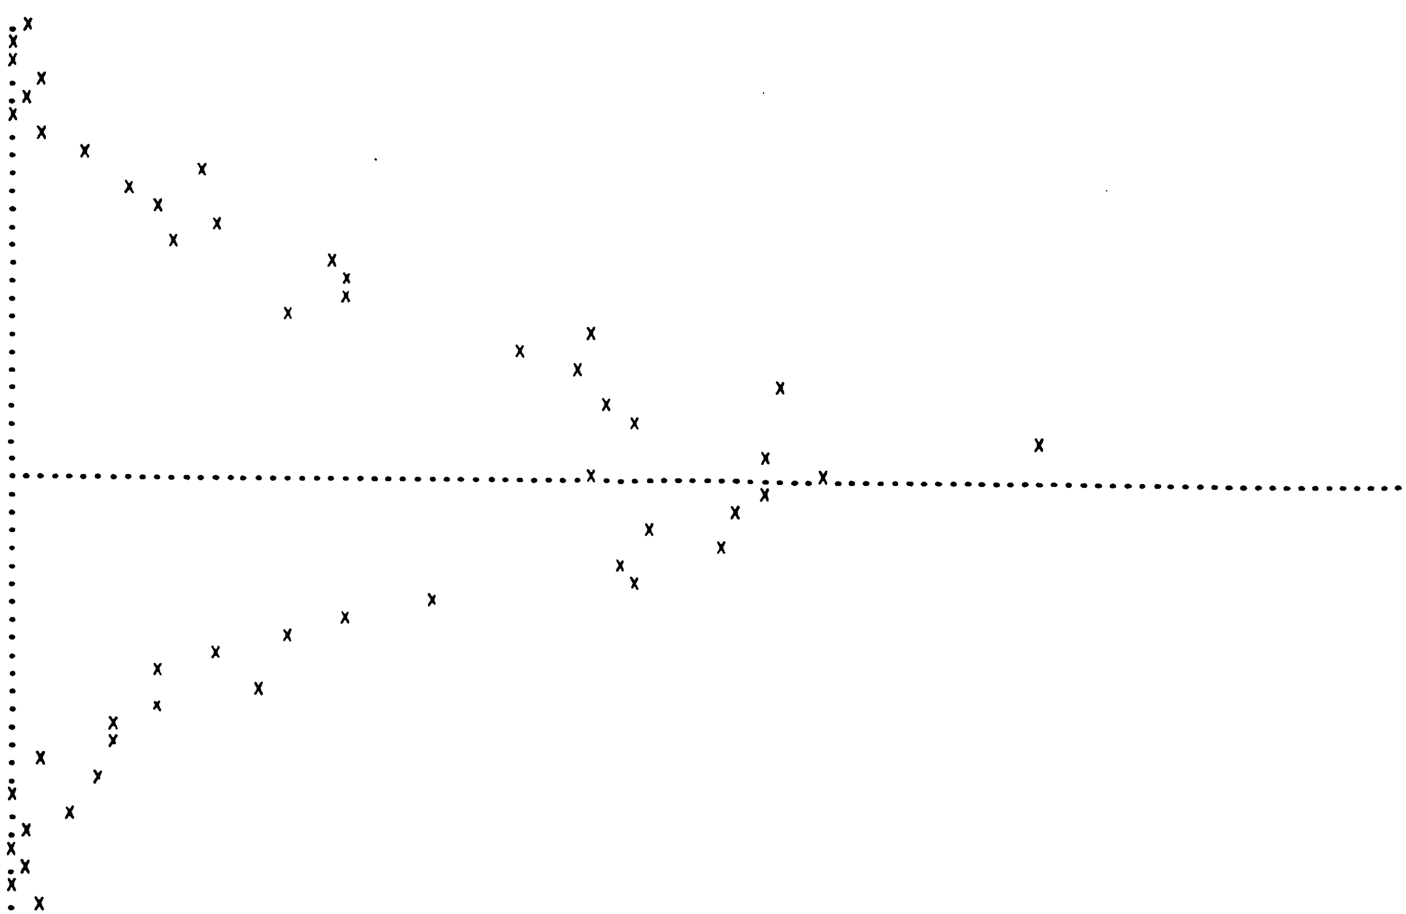

[1] Based on (10) and forecasts of a constant real electricity price and real income growth of 2 percent.

Figure 3

Recursive Bootstrap of Annual Peak Demand, 1990[1]

```
BOOTSTRAP FORECAST, ANNUAL,  1990
*******************

THE FORECAST FOR 1990 IS          22095.81430
THERE IS A PROBABILITY OF   .50   THAT PEAK WILL NOT EXCEED    22121.02   MEGAWATTS
THERE IS A PROBABILITY OF   .75   THAT PEAK WILL NOT EXCEED    22604.42   MEGAWATTS
THERE IS A PROBABILITY OF   .90   THAT PEAK WILL NOT EXCEED    23062.96   MEGAWATTS
THERE IS A PROBABILITY OF   .95   THAT PEAK WILL NOT EXCEED    23398.81   MEGAWATTS
THERE IS A PROBABILITY OF   .99   THAT PEAK WILL NOT EXCEED    23992.98   MEGAWATTS
```

FREQUENCY DISTRIBUTION OF SIMULATED FORECAST ERRORS AROUND THE FORECAST PEAK LOAD DEMAND

PEAK LOAD	FREQ
18990.64141	1
19102.73591	0
19214.83041	1
19326.92491	0
19439.01941	0
19551.11391	0
19663.20841	0
19775.30291	0
19887.39741	2
19999.49191	2
20111.58642	3
20223.68092	3
20335.77542	5
20447.86992	7
20559.96442	9
20672.05892	18
20784.15342	15
20896.24792	21
21008.34242	20
21120.43692	24
21232.53142	23
21344.62592	30
21456.72042	41
21568.81492	43
21680.90942	41
21793.00392	53
21905.09842	56
22017.19292	53
22129.28742	56
22241.38192	56
22353.47642	53
22465.57092	68
22577.66542	49
22689.75992	38
22801.85442	38
22913.94892	42
23026.04342	27
23138.13792	16
23250.23242	19
23362.32692	14
23474.42142	11
23586.51592	8
23698.61042	6
23810.70492	6
23922.79942	5
24034.89392	4
24146.98842	5
24259.08292	1
24371.17742	0
24483.77192	2

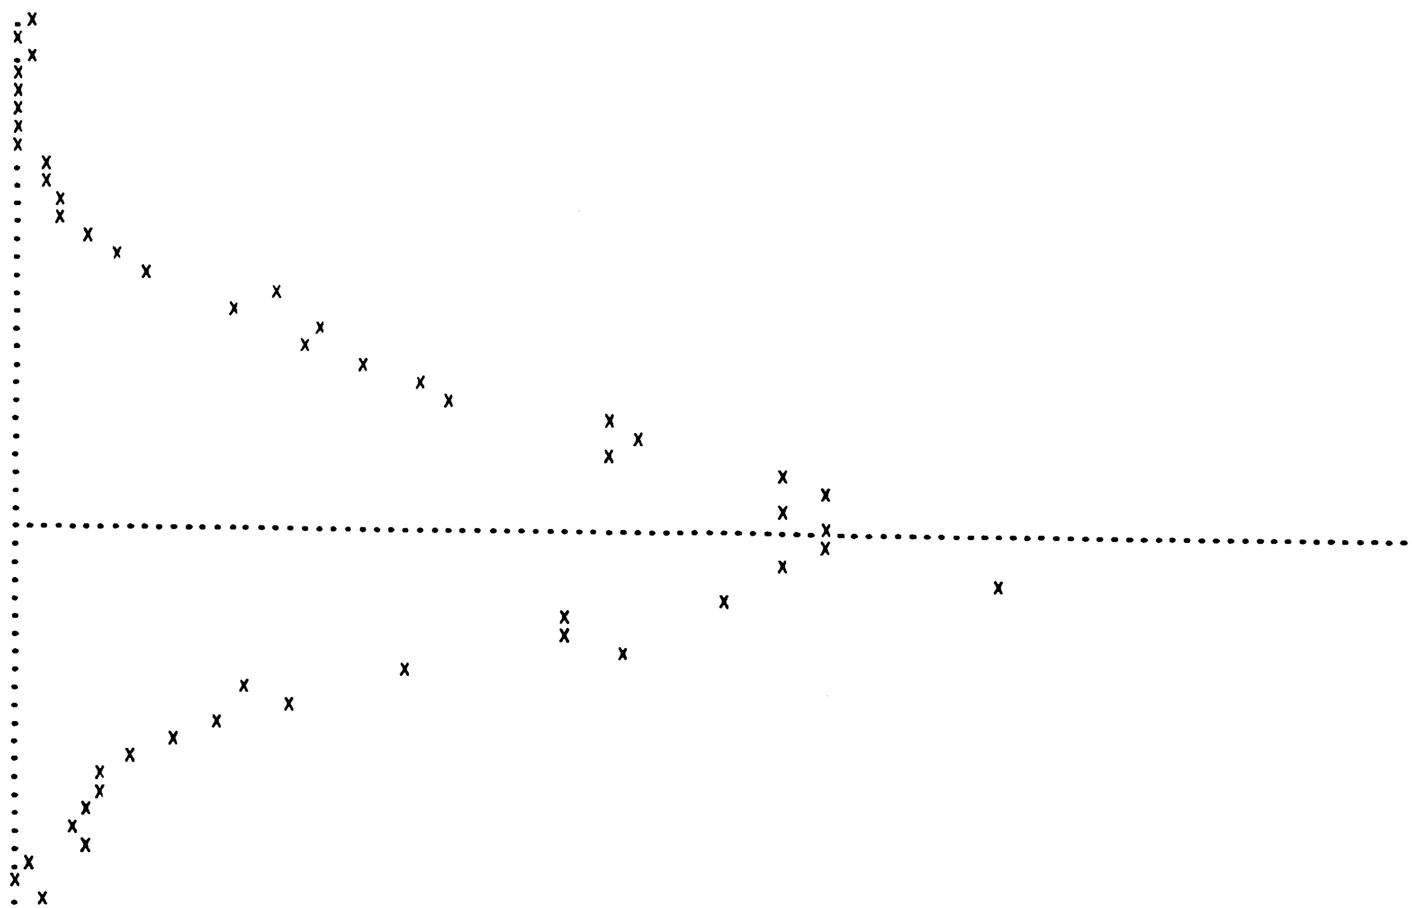

[1] Based on (9) and (10) and forecasts of a constant real electricity price and real income growth of 2 percent.

Figure 4

Recursive Bootstrap of Annual Peak Demand with Uncertain Forecasts of Independent Variables, 1990[1]

```
BOOTSTRAP FORECAST, ANNUAL,  1990
*******************

THE FORECAST FOR 1990 IS              22095.81430
THERE IS A PROBABILITY OF  .50  THAT PEAK WILL NOT EXCEED    22065.62  MEGAWATTS
THERE IS A PROBABILITY OF  .75  THAT PEAK WILL NOT EXCEED    23116.82  MEGAWATTS
THERE IS A PROBABILITY OF  .90  THAT PEAK WILL NOT EXCEED    23928.28  MEGAWATTS
THERE IS A PROBABILITY OF  .95  THAT PEAK WILL NOT EXCEED    24607.83  MEGAWATTS
THERE IS A PROBABILITY OF  .99  THAT PEAK WILL NOT EXCEED    25708.23  MEGAWATTS
```

FREQUENCY DISTRIBUTION OF SIMULATED FORECAST ERRORS AROUND THE FORECAST PEAK LOAD DEMAND

PEAK LOAD	FREQ
18119.81940	4
18300.07670	2
18480.33401	1
18660.59131	5
18840.84861	7
19021.10592	5
19201.36322	1
19381.62052	6
19561.87783	11
19742.13513	15
19922.39243	25
20102.64974	20
20282.90704	28
20463.16434	25
20643.42165	32
20823.67895	39
21003.93625	38
21184.19356	44
21364.45086	44
21544.70816	38
21724.96546	42
21905.22277	44
22085.48007	45
22265.73737	37
22445.99468	46
22626.25198	39
22806.50928	48
22986.76659	46
23167.02389	45
23347.28119	37
23527.53850	32
23707.79580	32
23888.05310	22
24068.31041	21
24248.56771	11
24428.82501	9
24609.08232	10
24789.33962	7
24969.59692	8
25149.85423	7
25330.11153	5
25510.36883	3
25690.62614	7
25870.88344	1
26051.14074	2
26231.39804	2
26411.65535	0
26591.91265	0
26772.16995	1
26952.92726	1

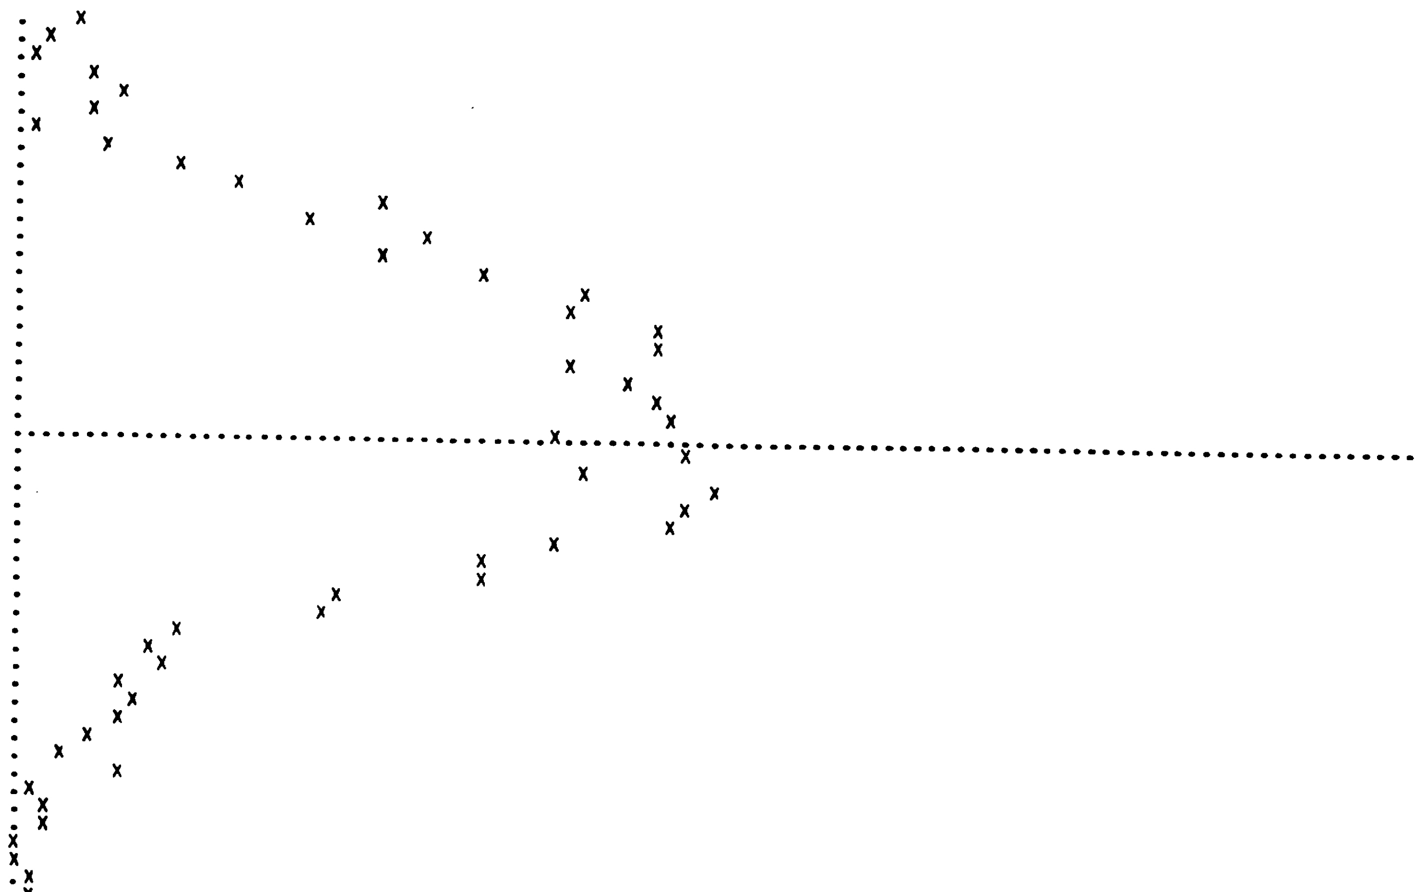

[1] Based on (9) and (10). The forecasts of the real electricity price and of the level of real income in 1990 are consistent with 0 and 2 percent growth respectively. However, in each bootstrapping run, a pseudo random normal deviate with standard deviation .1 has been added to the 1990 log levels of real price and real income.

Figure 5

Recursive Bootstrap of Annual Peak Demand with Uncertain Binomial Forecasts of Independent Variables, 1990[1]

BOOTSTRAP FORECAST, ANNUAL, 1990

THE FORECAST FOR 1990 IS 22055.81430
THERE IS A PROBABILITY OF .50 THAT PEAK WILL NOT EXCEED 22006.09 MEGAWATTS
THERE IS A PROBABILITY OF .75 THAT PEAK WILL NOT EXCEED 23446.91 MEGAWATTS
THERE IS A PROBABILITY OF .90 THAT PEAK WILL NOT EXCEED 24514.03 MEGAWATTS
THERE IS A PROBABILITY OF .95 THAT PEAK WILL NOT EXCEED 25054.26 MEGAWATTS
THERE IS A PROBABILITY OF .99 THAT PEAK WILL NOT EXCEED 25779.39 MEGAWATTS

FREQUENCY DISTRIBUTION OF SIMULATED FORECAST ERRORS AROUND THE FORECAST PEAK LOAD DEMAND

PEAK LOAD	FREQ
17532.47333	2
[illegible]	[illegible]
17754.47444	[illegible]
[illegible]	[illegible]
18170.51555	5
18363.52611	5
18556.53667	4
18749.54722	6
18942.55778	14
19135.56834	11
19328.57889	16
19521.58945	16
19714.60000	24
19907.61056	25
20100.62112	24
20293.63167	31
20486.64223	[illegible]
20679.65279	33
20872.66334	43
21065.67390	[illegible]
21258.68446	[illegible]
21451.69501	38
21644.70557	45
21837.71613	37
22030.72668	26
22223.73724	[illegible]
22416.74779	35
22609.75835	[illegible]
22802.76891	[illegible]
22995.77946	[illegible]
23188.79002	[illegible]
23381.80058	41
23574.81113	[illegible]
23767.82169	[illegible]
23960.83225	24
24153.84280	25
24346.85336	[illegible]
24539.86392	22
24732.87447	14
24925.88503	22
25118.89559	11
25311.90614	14
25504.91670	11
25697.92725	6
25890.93781	4
26083.94837	2
26276.95892	3
26469.96948	1
26662.98004	1
26855.99059	2

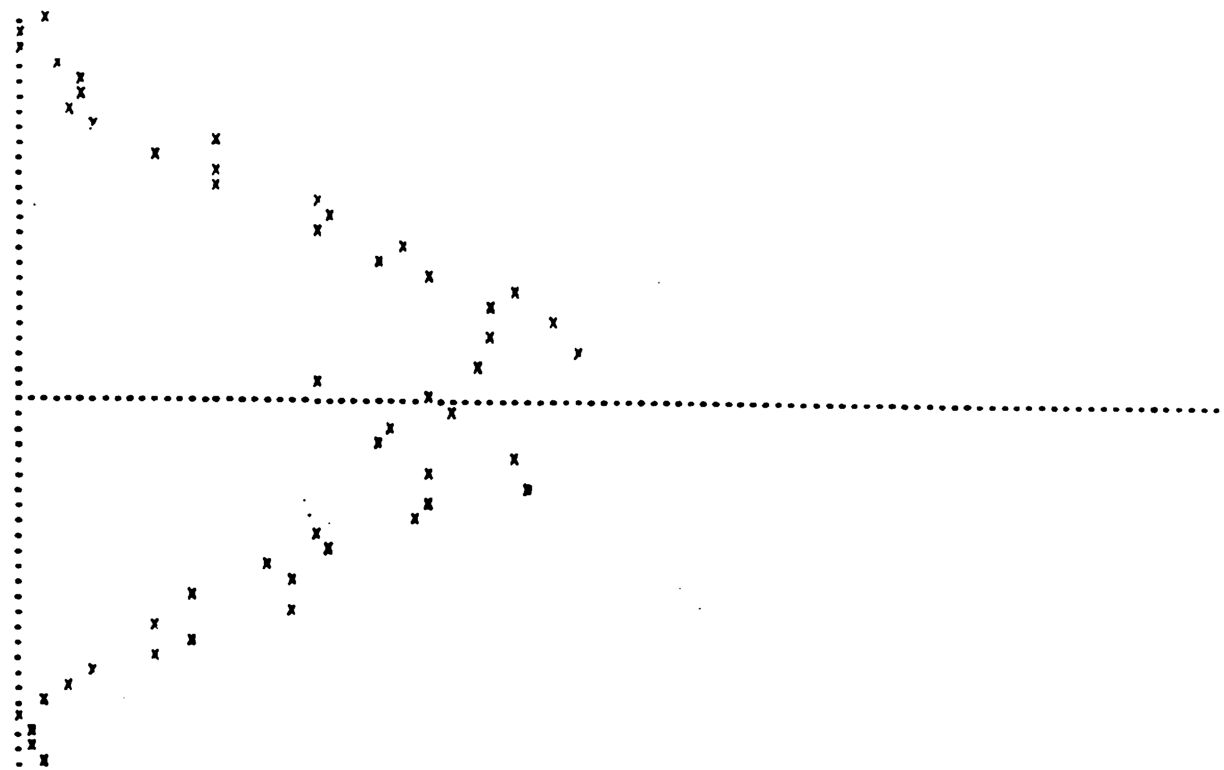

[1] Based on (9) and (10). The forecasts of the real electricity price and of the level of real income in 1990 are consistent with 0 and 2 per cent annual growth respectively. However, in each bootstrapping run, a pseudo random variate has been added to the 1990 log levels of real price and real income. For price, this variate has .5 probability of being -.14 and .5 probability of being .14, implying a standard deviation of .1. For income, the variate is normal with standard deviation .1.