

MODELLING AND FORECASTING DAILY SERIES OF ELECTRICITY DEMAND *

José Ramón CANCELO

Universidad de La Coruña

Antoni ESPASA

Universidad Carlos III de Madrid

We present a model to forecast daily electricity demand in Spain. Instead of following the usual practice of splitting the sample into homogeneous subsamples and building separate models, we advocate the use of a single model for the whole sample. Thus the paper focuses on how to achieve a single-equation model that encompasses highly heterogeneous data while remaining stable for a quite long period of time. This case study may also be useful as an example of how to proceed to model daily series of economic activity.

1. Introduction

Forecasting electricity demand in the short run is a very significant economic problem, because electricity cannot be stored and producers have to anticipate future demand in a very accurate way in order to meet it while avoiding overproduction. In this paper we present a model to forecast daily electricity demand in Spain; following Bunn and Farmer's (1985) classification, it is an off-line model, designed to improve the plan scheduling process from one day to one week ahead. The model is also used in the study mode (Gross and Galiana, 1987) to simulate the effect of different scenarios on the system.

Although the main purpose of the paper is to present a solution for a specific forecasting problem, we think that it is not restricted to electricity demand forecasting. The paper is an application of a general methodology to model daily series of economic activity (Espasa 1993, Espasa and Cancelo 1995, Espasa *et al.* 1996), and it may be useful for analysts working on related fields of economic time series analysis. The contribution of the paper to the literature of forecasting this kind of high-frequency economic data may be summed up in two main points:

* We thank Rosmarie Grafe and José Félix Izquierdo for helpful suggestions and efficient research. Thanks are also due to two referees whose comments have led to significant improvements over earlier versions. We acknowledge financial support from Spanish DGICYT, projects PB93-0236 (both authors) and PB93-0653 (Cancelo).

- 1) Forecasts are obtained from a single equation: contrary to the usual practice, we consider a general dynamic model that captures the basic underlying process of electricity demand into a single equation.
- 2) The way we deal with special days and weather variables generalizes previous proposals: special days are modelled by a detailed intervention analysis; the contribution of weather variables is taken into account by piecewise linear approximations to nonlinear transfer functions, which are allowed to vary according to the season, the day of the week, etc.

The paper is organized as follows: the main characteristics of electricity demand in Spain are reviewed in Section 2. Section 3 shows a general overview of the model. Section 4 centers on modelling special days. The effect of weather variables is the subject of Section 5. The dynamic structure of the disturbance and the validation analysis of the proposed model are discussed in Section 6, and the main conclusions are summarized in Section 7.

2. General characteristics of the system configuration

Electricity distribution in the peninsular part of the Spanish territory (excluding Canary and Balearic Islands and the african cities of Ceuta and Melilla) is centralized at a state-owned utility. One of its tasks is to set the production of the several producers, given their capacity and the expected demand.

Peninsular Spain covers an area of almost half a million squared kilometers, with more than 36 million inhabitants. Taken as a whole, the system is winter-peaking: in 1988 maximum demand in winter was 1.15 times maximum demand in summer. Cold days are more frequent than hot days: we built a maximum temperature index for the whole territory by computing a weighted average of the maximum temperature recorded in ten selected observatories; in 1988 the index was below 20C (68F) in 150 days, and above 24C (75.2F) in 119 days.

The variable to model is the net demand for electricity, defined as total production from all sources plus international interchanges balance less intermediate autoconsumption and pumping consumption. It results from aggregating very different regional systems; a spatially disaggregated analysis would certainly improve the results shown in this paper, but regional data were not available. The series shows a growing trend, related to economic and socio-demographic factors, and several seasonal oscillations, among which annual and weekly ones are the most remarkable. Special days (holidays, vacation periods and any anomaly that alters normal daily activity) and weather conditions also exert a major influence.

Figures 1 to 3 show some of these features. In Figure 1 the growing trend is easy to see, as well as the annual seasonal. In Figure 2 several intra-annual components may be observed: the weekly seasonal cycle; the extra consumption due to temperature effects in winter (both extremes of the

figure) and in summer (the relative maximum at mid-year); and the influence of vacation periods (see for instance August). Figure 3 shows the distortion induced by a special day (12 October), which breaks the usual pattern of behavior within the week.




Figure 1
Monthly Demand in 1983-1990




Figure 2
Daily Demand in 1989

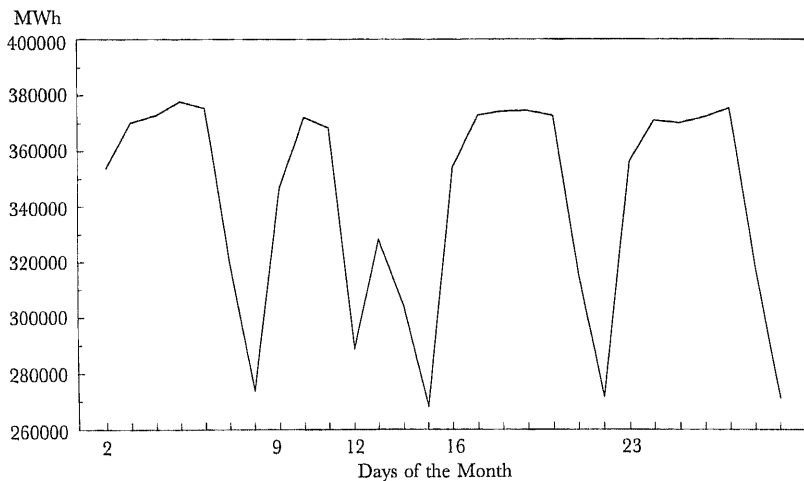


Figure 3
Distortion caused by a holiday
(Monday, Oct. 2 to Sunday, Oct. 29, 1989)

3. Overview of the model

Observed demand (D_t) may be expressed as the aggregation of the following unobserved components:

$$\ln D_t = SE_t + SW_t + CSD_t + CWEA_t + \varepsilon_t, \quad [1]$$

where $\ln(\cdot)$ denotes natural logarithm, SE_t a trend related to socioeconomic factors, SW_t a seasonal weekly component induced by the normal workdays/weekends pattern, CSD_t the contribution of special days that distort the normal pattern, $CWEA_t$ the contribution of weather variables and ε_t represents purely random transitory shocks not related to any of the previous factors.

A base demand BD_t can be defined from [1]:

$$\ln BD_t = \ln D_t - CSD_t - CWEA_t = SE_t + SW_t + \varepsilon_t. \quad [2]$$

As short-term oscillations caused by special days and sudden changes of the weather conditions have been eliminated, this new series displays a much more regular evolution. It may be modelled by the following ARIMA model:

$$\Delta \Delta_7 \ln BD_t = \eta_t, \quad [3]$$

stationarity being achieved by applying the variance reduction criterion as described in Espasa and Cancelo (1995) or Espasa *et al.* (1996) (η_t denotes a stationary disturbance that follows an ARMA process). From [2] and [3]:

$$\Delta \Delta_7 \ln D_t = \Delta \Delta_7 CSD_t + \Delta \Delta_7 CWEA_t + \eta_t, \quad [4]$$

Without loss of generality we may assume that

$$\begin{aligned}
 CSD_t &= f_1(L)' SD_t = \sum_{i=1}^m f_{1,i}(L) SD_{i,t} \\
 CWEA_t &= f_2(L)' WEA_t = \sum_{j=1}^{m'} f_{2,j}(L) WEA_{j,t}
 \end{aligned}
 \tag{5}$$

where $SD_t = (SD_{1,t} \dots SD_{m,t})'$ stands for a vector of dummy variables used to model special days and $f_1(L)$ for a vector of polynomials in the lag operator L , in the spirit of Box and Tiao's (1975) intervention analysis. In a similar way, $WEA_t = (WEA_{1,t} \dots WEA_{m',t})'$, refers to the set of weather variables and $f_2(L)$ represents a vector whose components are functions of L . No restrictions are imposed either on the variables or on the filters, so [5] entails no *a priori* simplification of the model. In particular, the orders of the polynomials $f_{1,i}(L)$ and $f_{2,i}(L)$ are free to adapt to the specific characteristics of the variables they are related to.

The final model is obtained by combining [2], [4] and [5]:

$$\Delta \Delta_7 \ln D_t = \Delta \Delta_7 f_1(L)' SD_t + \Delta \Delta_7 f_2(L)' WEA_t + \eta_t.
 \tag{6}$$

This expression can be reformulated as a transfer function model on $\ln D_t$ with non-stationary noise:

$$\ln D_t = f_1(L)' SD_t + f_2(L)' WEA_t + \frac{\eta_t}{\Delta \Delta_7}.
 \tag{6'}$$

Leaving aside the contribution of special days and weather conditions, which will be discussed separately in the next sections, its main characteristics are:

3.1. The logarithmic transformation

The usual practice in short-term electricity forecasting is to work with non-transformed data, although it is well known that ARIMA models for monthly or quarterly economic series are usually based on logged data due to heteroskedasticity. A similar result holds for daily data: to work with non-transformed data may not induce serious distortions in a sample of three or four months, but it is not a good choice when dealing with several years of data at a time; Bogard *et al.* (1982) provide some valuable reflections on this issue.

3.2. Trend and weekly seasonal modelling with difference operators

Economic conditions, electricity prices, population levels and the determinants of the regular weekly pattern do not change on a day-to-day basis, and they may be modelled implicitly by slowly evolving trend and seasonal components. The difference filters proposed in [6]-[6'] provide the minimum variance of the transformed series and they allow for a simple, parsimonious way of explaining non-stationary characteristics of the variable. Their usefulness in modelling daily or hourly energy series is well

documented in the literature: see for instance Ernoult and Meslier (1982), Goh *et al.* (1986), Pigott (1985) or Bogard *et al.* (1982). O'Connor and Kapoor (1984) compare the performance of a purely stochastic ARIMA-type model with a deterministic plus stochastic model, where sums of sinusoids at the fundamental frequency and its harmonics are explicitly estimated. Not surprisingly –from our point of view–, the latter shows a better fit to the sample data, but the ARIMA model is better for forecasting; the authors point out that this result seems to be the consequence of the superior flexibility of the ARIMA specification to adapt to changes in the structure of the data.

Specification, estimation and validation are based on data from 1983 to 1988; all results shown in this paper correspond to this sample. The final specification has 126 parameters, which have been estimated from a sample of 2174 observations; 87 parameters are for special days, 30 for weather conditions and 9 for the random disturbance. Joint estimation was performed with the SCA software (Scientific Computing Associates, 1986), which provides maximum likelihood estimators based on the exact likelihood function. Data for 1989 were used to test post-sample stability. Since then, the model has been operated by the utility and used for real forecasting. The parameters of the model are reestimated every six months; sample size is kept constant, i.e., each time a new semester of data is introduced the first six months at the beginning of the sample are discarded.

4. Modelling special days

4.1. Outline of the procedure

The distortion on the demand caused by special days has been treated in very different ways: compare for instance Hesterberg and Papalexopoulos (1988), Ernoult and Mattatia (1984), and Bolzern and Fronza (1986). Our procedure for estimating their influence is based on:

a) Some authors advocate to begin by erasing special days from the database and modelling solely «normal» data; see for instance Pigott (1985). Quite the reverse, we do not pursue blind elimination of anomalies, but try to detect stable patterns of the demand behavior in special days, so that feasible rules to forecast the influence of similar events in the future might be found.

b) We analyze the characteristics of each type of special day by considering the information about it contained in several years of data. We seek to identify stable patterns of behavior that may lead to group homogeneous special days, and to assign them the same effect on the demand, in order to get more efficient estimators: this analysis allows to impose general restrictions with respect to models that estimate one coefficient for each anomaly. It was checked, by using likelihood ratio tests, that the restrictions were not rejected by the data.

Special days are classified into three groups: single bank holidays, vacation periods (Easter, August and Christmas) and other anomalies (strikes, elections, etc.). In this paper we make a brief review of the treatment of bank holidays, as an illustration of the type of intervention analysis carried out in the modelling process; a more detailed exposition of the influence of special days in this series is available in Cancelo and Espasa (1991) or from the authors upon request.

4.2. The influence of single Bank Holidays

The starting hypothesis is that the effect of a given bank holiday only depends on the day of the week on which it falls. While the data seem to support this assumption for most holidays, it was detected that imposing this restriction entailed bad forecasts for very specific holidays, and we had to distinguish general holidays from special holidays.

General holidays: Most Spanish holidays belong to this category, and a uniform treatment suits them. Six dummy variables were used, one for each day of week (no effects were detected for general holidays falling on Sundays); they take values in the $[0,1]$ interval according to the weight in total consumption of the territory affected by the holiday. All variables are assigned a dynamic filter in order to capture possible pre and post effects.

Special holidays: There are three single bank holidays –6 January, 1 May and 15 August– which do not fit into the general scheme and are modelled separately. It was carefully checked that they needed special consideration; to base estimates on a single observation is not a good practice as a general rule, but for special holidays it seems worse to impose a common effect because they influence electricity demand in a very different way. Even so, some restrictions, such as some symmetry in the responses for holidays falling on Tuesdays and Thursdays, were imposed in order to increase the efficiency of the estimators.

The final estimates are shown in Table 1. They reflect the change in the logged demand caused by each holiday assuming that it has effects on the whole territory; for regional and local holidays these coefficients should be multiplied by the weight of the affected area on total consumption.

5. Modelling weather conditions: the influence of temperature

Given their expected influence on demand and the availability of data, the following meteorological indicators were selected: maximum daily temperature (which was preferred to average temperature as this was the usual practice at the utility), a temperature dispersion indicator (defined as the difference between maximum and minimum temperature) and solar light hours. Following the usual practice (see for instance Ernoult and Meslier (1982) for France, Pigott (1985) for the UK, or Bolzern and Fronza (1986) for Italy), all meteorological data used in the model are weighted averages of ten selected observatories that cover the whole territory.

Table 1
Estimated coefficients for holidays

	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
	-						
Sunday	-0.0503 (7.1)						
		-0.2566 (45.8)	-0.0409 (7.4)				
Monday		-0.2782 (33.2)	-0.0254 (3.2)				
		[2]	[2]				
		-0.1809 (27.2)	-0.0412 (6.4)				
		-0.1017 (17.3)	-0.2962 (44.4)	-0.0333 (5.8)			
Tuesday		-0.1003 (12.4)	-0.2847 (32.6)	-0.0379 (3.5)			
		-0.1142 (10.7)	-0.3357 (37.5)	-0.0539 (4.9)			
		[2]	[2]	[2]			
			-	-0.2604 (38.0)	-0.0398 (5.8)		
Wednesday			-	-0.2782 (33.2)	-0.0254 (3.2)		
			-0.0538 (6.5)	-0.3537 (37.5)	-0.0538 (6.5)		
			-	-0.1548 (19.6)	-0.0416 (5.2)		
					-0.2520 (57.0)	-0.1083 (21.6)	-0.0255 (5.8)
Thursday					-0.2847 (32.6)	-0.1003 (12.4)	-
					-0.3080 (38.3)	-0.1025 (12.8)	-
					-0.1548 (19.6)	-0.0416 (5.2)	-
						-0.2544 (35.9)	-0.0839 (11.9)
Friday						-0.2847 (32.6)	-0.1003 (12.4)
						-0.3080 (38.3)	-0.1025 (12.8)
						-0.1548 (19.6)	-0.0416 (5.2)
							-0.0795 (18.9)
Saturday							[2]
							[2]
							-0.0693 (6.8)

Notes: [1] The figures refer to the estimated effect on the logged demand. In each cell the first figure corresponds to a general holiday, the second to 6 January, the third to 1 May and the fourth to 15 August; t-ratios in absolute value are shown in parentheses. [2] No information in the sample 1983-1988 to estimate this effect.

In this paper we focus on the relationship between the level indicator of temperature and the demand for electricity, because it is by far the most important for forecasting purposes.

5.1. *A priori information on the relationship*

The main characteristics of the relationship between temperature and electricity demand are:

- a) The relation is highly nonlinear. There are two critical values of temperature, $T^* < T^{**}$, that define a neutral zone in which temperature does not influence demand. Below T^* we enter into the cold zone, and above T^{**} in the hot zone. The response function may be nonlinear in any of the cold and hot zones.
- b) With daily data there is a dynamic response, as the demand for day t will be influenced by temperatures observed in $t, t-1, \dots, t-h$. EPRI (1983) provides arguments explaining this feature; see also Gross and Galiana (1987).
- c) There may exist a temperature value low (high) enough as to oblige every heating (cooling) system to operate at full capacity; and additional decreases (increases) of temperature will have no additional effects on demand.
- d) The influence of a given temperature may be different according to the time of the year: Bogard *et al.* (1982) and Ernout and Meslier (1982) remark this point.
- e) The response of the demand to a given temperature may be different for working and non-working days: see for instance Hesterberg and Papalexopoulos (1988).
- f) For a small period of time the stock of appliances may be assumed to be constant, but when several years are considered this assumption may not be valid: see Fischler and Nelson (1986) for a further discussion.

The two first characteristics are the most important for forecasting, and together they define what we shall call from now on the basic effect of temperature; a procedure to model this basic effect was presented in Cancelo and Espasa (1995). The remainder are enlarged effects, and we shall refer to them as exhaustion effect (characteristic c), time-of-the-year effect (d), type-of-day effect (e) and stock-of-appliances effect (f). Enlarged effects contribute to get a more reliable approximation to the relation and more precise forecasts for specific dates.

5.2. *A methodology for searching for specification*

We propose the following sequential strategy:

Stage 1. To avoid the distortion caused by special days an ARIMA model with intervention analysis is built:

$$\ln D_t = f^*_1(L)'SD_t + \frac{\theta^*(L)}{\Delta \Delta_7 \Phi^*(L)} e_t, \tag{7}$$

and the logged demand corrected for special days is computed:

$$\ln DCSD_t = \ln D_t - f^*_1(L)'SD_t. \tag{8}$$

This will be the dependent variable in Stages 2 to 6.

Stage 2. We begin with the basic effect, the most important and the most difficult to model. The procedure described in Cancelo and Espasa (1995) was used: in short, it consists of approximating a nonlinear dynamic relationship by linear transfer functions defined on threshold variables, so that a piecewise linear approximation results. A first estimation of the relationship is obtained:

$$\ln DCSD_t = \sum_{j=1}^n \beta_c^{(j)}(L) C_t^{(j)} + \sum_{j=1}^{n'} \beta_h^{(j)}(L) H_t^{(j)} + \frac{\theta(L)}{\Delta \Delta_7 \Phi(L)} a_t \quad [9]$$

Cold threshold variables take non-zero values only for temperatures below a given knot j : $C_t^{(j)} = \max(0, j - T_t)$, T_t being the maximum temperature of day t . A similar definition applies to the hot thresholds $H_t^{(j)} = \max(0, T_t - j)$. Call expression [9] model A.

Model A was the final result of the application in Cancelo and Espasa (1995). However, for the present problem of forecasting it is not a final solution because it imposes a priori restrictions on all enlarged effects. In the next stages we carry out a sequential procedure to test whether these constraints are supported by the data, and extend the model if necessary.

Stage 3. Exhaustion effects are tested by considering several candidates for exhaustion thresholds, estimating a different model for each candidate, choosing the best one, and comparing it with model A by using likelihood ratio tests. Both cold and hot zones are assigned exhaustion thresholds. Call this model B.

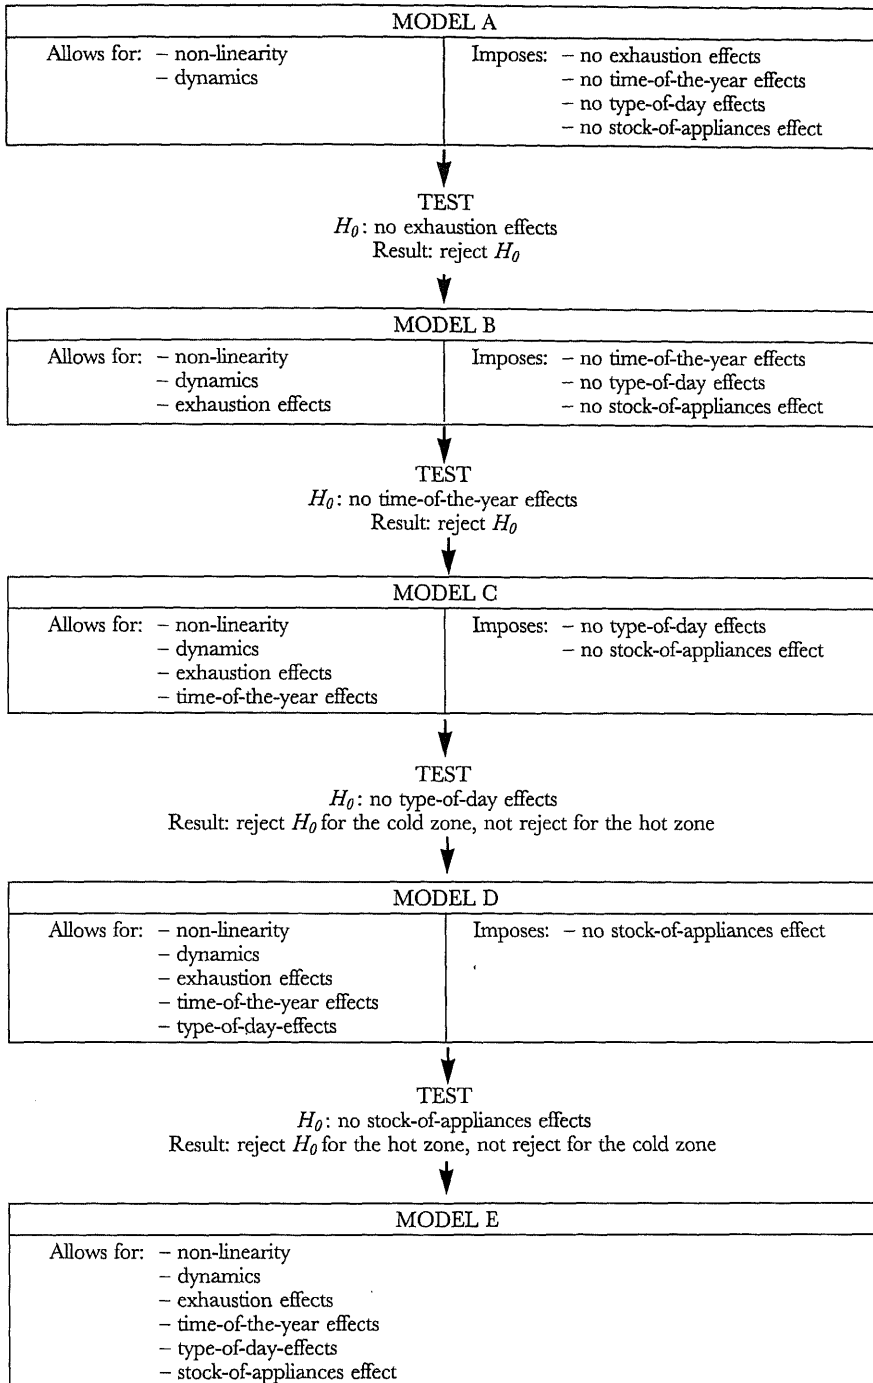
Stage 4. The year was split into four seasons which do not match the standard division as the latter is not adequate for our purposes; Bogard *et al.* (1982) use almost the same partition as ours, although they reached it by a purely empirical analysis and we made extensive use of extrasample information from meteorological experts. Afterwards, the corresponding temperature variables were generated, a generalisation of model B that allowed for time-of-the-year effects was specified and estimated, non-significant coefficients were deleted, and likelihood ratio tests were used to determine which between-seasons constraints could be imposed. Finally a model that includes time-of-the-year effects, call it model C, was obtained.

Stage 5. The sample was divided into working and non-working days and the same procedure as in Stage 4 was applied. The results of the whole process point to a response function for the cold zone that varies according to the type of day. This led to model D.

Stage 6. It was checked whether there exists an increase-of-the-stock-of-appliances effect. We centered the study on the hot zone, as it was known that sales of cooling systems in Spain experienced a boom during the second half of the eighties. The data confirmed that there is a shift in the response function for the hot zone, and model E includes it. We remark that with aggregate data the only way to detect this type of effect is by looking for shifts in the response function; Fischler and Nelson (1986) propose an alternative procedure based on a saturation measure of the existing stock, but their approach is not feasible in most applications because no reliable information on the saturation degree is usually available (EPRI, 1983).

Table 2

Sequential procedure for approximating the relation between temperature and demand



Stage 7. Although all temperature effects already included in the model and the parameters of the disturbance process are estimated in each stage, the coefficients of special days remain fixed on the estimates of Stage 1. So the procedure ends with a full efficient estimation of all the coefficients of the final specification.

The whole procedure is summarized in Table 2. It may be argued that the final specification depends on the way enlarged effects are ordered; although we can not rule out this possibility, trying to consider all effects at once is unfeasible and some type of sequential analysis has to be done. We asked the experts of the utility to rank enlarged effects on a more-to-less-important basis and used their ordering in the search; as a consequence the procedure in Table 2 may be termed a simple-to-general specification search guided by extrasample information.

5.3. Main results of the estimated relationship

Neutral zone is defined as the interval of temperatures between 20C (68F) and 24C (75.2F). Below 20C we enter into the cold zone, its main results being summarised in Table 3. Temperatures above 24C constitute the hot zone, and their effect on demand is summarised in Table 4.

Table 3
Summary of the estimated coefficients for the cold zone

Season	Non-linearity: knots at	Dynamics: lags	Type-of-day effects:	
			working	non-working both
winter/spring	9C / 48.2F	1		1.03 (3.1)
		2		.57 (2.8)
		3		.57 (2.8)
	11C / 51.8F	1		-.57 (2.6)
		1		.26 (3.3)
	20C / 68F	0	.54 (18.4)	.62 (15.8)
		1	.21 (5.3)	.27 (5.4)
		2	.23 (7.4)	.30 (7.3)
		3	.13 (4.2)	.00 (0.4)
		4	.12 (4.1)	.15 (3.7)
		5	.09 (3.0)	.12 (2.6)
		6	-	.08 (1.7)
	7	-	.08 (2.1)	
	summer/fall	20C / 68F	0	

Notes:

1. The coefficients are semielasticities expressed in %: a coefficient equal to 0.1 implies that for each Celsius degree below the given knot the demand for electricity increases an additional 0.10%.
2. Exhaustion effect at 9C / 48.2F.
3. Type-of-day effect only detected for the threshold variable with knot at 20C / 68F.
4. No stock-of-appliances effect detected.
5. No temperatures below 11C / 51.8F observed in spring.
6. No temperatures below 14C / 57.2F observed in summer or fall.
7. *t*-ratios in parenthesis.

Table 4
Summary of the estimated coefficients for the hot zone

Season	Non-linearity: knots at	Dynamics: lags	Stock-of-appliance effects:	
			1983-1987	1988 onwards
Summer/Fall	24C / 75.2F	0	.11 (2.9)	.29 (3.5)
		1	.31 (8.4)	.29 (3.5)
		2	.07 (2.0)	.15 (1.8)

Notes:

1. The coefficients are semielasticities expressed in %: a coefficient equal to 0.1 implies that for each Celsius degree above the given knot the demand for electricity increases an additional 0.10%.
2. Exhaustion effect at 33C / 91.4F.
3. No type-of-day effect detected.
4. No temperatures above 33C / 91.4F observed in fall.
5. Temperatures on the hot zone observed in spring do not affect demand. No temperatures on the hot zone were observed in winter.
6. *t*-ratios in parenthesis.

6. Stochastic properties of the disturbance term and model evaluation

Simple and partial correlograms of the disturbance η_t in [6] suggest that it is generated by the multiplicative moving average process (1, 2) (7, 14) (357, 364, 728, 731, 735). The estimates of its parameters are shown in Table 5. Although the annual part has not a straightforward interpretation, it has been checked that it is not caused by outliers. Moreover, the coefficients are such that if an annual AR (2) process with positive coefficients –which has a much more appealing interpretation– is used instead, the residual standard deviation and correlograms would be almost the same, with no remarkable changes on the forecasting performance of the model.

Table 5
Estimates of the parameters of the moving average process of the random disturbance

Parameter	Estimate	<i>t</i> -ratio
θ_1	.21	9.6
θ_2	.15	6.7
θ_7	.83	37.6
θ_{14}	.09	3.7
θ_{357}	-.06	2.6
θ_{364}	-.14	5.9
θ_{728}	-.09	3.4
θ_{731}	-.09	3.5
θ_{735}	-.08	3.2

The final estimate of the residual standard deviation equals 0.0130: the standard deviation of the one-period forecast error, given the true values of the weather variables, is 1.30% times the point forecast. This was considered

a pretty good performance by the experts of the utility given previous forecasting experiences for this variable. The actual forecasting error is higher because some values of meteorological variables are unknown and must be replaced by forecasts. Weather forecasts are obtained from meteorological experts and are neutral, in the sense that all competing procedures use this type of meteorological forecasts.

To validate the model, the following analyses were carried out:

1) The usual misspecification tests for the residuals as a whole:

a) Zero-mean: the t -ratio to test this hypothesis equals 0.002.

b) Autocorrelation: $Q(7) = 11$, $Q(14) = 14.2$, $Q(30) = 42.3$ and $Q(365) = 406$, where $Q(k)$ denotes the Box-Pierce-Ljung portmanteau statistic computed from the first k autocorrelations.

c) Skewness and kurtosis: the tests based on the asymptotic distributions of the sample coefficients of skewness and kurtosis (Spanos 1986, p. 454) reject the null hypotheses of a symmetric (t -ratio -3.7), mesokurtic (t -ratio 13.3) distribution. But if the fourteen residuals greater than $3\hat{\sigma}_\epsilon$ in absolute value are deleted then the hypothesis of symmetry is no longer rejected (t -ratio -1.1) and the t -ratio to test kurtosis falls to 2.7. These results indicate that although normality does not hold the true underlying distribution is not far from it.

2) The residuals were split by days of the week, months of the year and years, in order to analyze separately each group. Splitting by two of these factors at a time was not advisable because too few observations resulted.

a) Zero-mean: in all cases except one (residuals of July, t -ratio 2.8) the null hypothesis of a non-significant mean is not rejected.

b) Homoskedasticity: Table 6 shows the standard deviation for all groups, and at a first sight it looks like the residuals are not homoskedastic: if they are classified according to the day of the week, volatility is higher around weekends; if they are grouped by the month of the year, December, April and August are the most difficult months to forecast, while November and October are the easiest; and the variance is not the same by years, although no regular pattern of change is detected. To confirm formally what the sample figures seemed to indicate we computed a test on the equality of several variances (Mood *et al.* 1974, pp. 439-440), and the null was rejected for all criteria of classification; the statistic takes the value 15.0 (asymptotically distributed as a chi-square with 6 degrees of freedom under the null) if the sample is split by days of the week, 52.8 (chi-square with 11) by months and 54.8 (chi-square with 5) by years.

This seems to be a serious shortcoming; but we think that the interesting result is not that residual standard deviation still varies, but that its variation is much smaller than the variation in the original data. Table 6 also shows the standard deviations computed from the stationary transformation ($\Delta \Delta_7 \ln$) of the original series for all the groups; a quick comparison provides strong

evidence that the model was successful in explaining most of the original heteroskedasticity.

Table 6
Standard deviations for different subsamples:
differenced series and residuals of the model

Day-of-the-week			Months-of-the-year			Years		
days	$\Delta \Delta \ln D_t$	residual	month	$\Delta \Delta \ln D_t$	residual	year	$\Delta \Delta \ln D_t$	residual
Sun.	.0421	.0140	Jan.	.1065	.0130	1983	.0788	.0152
Mon.	.0689	.0134	Feb.	.0250	.0125	1984	.0802	.0130
Tues.	.0776	.0120	Mar.	.0666	.0136	1985	.0729	.0139
Wed.	.0914	.0120	Apr.	.0974	.0142	1986	.0680	.0104
Thurs.	.0982	.0124	May	.0912	.0131	1987	.0757	.0129
Fri.	.0777	.0140	June	.0719	.0129	1988	.0826	.0122
Sat.	.0672	.0130	July	.0431	.0119	TOTAL	.0765	.0130
TOTAL	.0765	.0130	Aug.	.0641	.0138			
			Sept.	.0309	.0127			
			Oct.	.0769	.0108			
			Nov.	.0756	.0098			
			Dec.	.1092	.0160			
			TOTAL	.0765	.0130			

Taken as a whole, the results of the analysis of validation indicate that the model provides unbiased forecasts for every moment of time and that it is not possible to improve it by considering linear functions of the available set of information. This is the major achievement of our work and the reason that explains why the model has beaten alternative forecasting procedures for this series. However, they also show that our results may be improved by taking into account non-normality and heteroskedasticity: the forecasting performance of the resulting model would remain almost unchanged but it would be possible to adapt the confidence intervals to get a more accurate measure of the uncertainty related to each forecast.

In any case, one would expect to face serious problems of non-linearity before beginning to work with the data, up to the point that a linear model as ours should be a very bad approximation to the underlying generating process. Quite on the contrary, the analysis of the residuals seems to indicate that our model is a not so bad approximation; further research on the subject should provide a measure of how far the linear approximation of this paper is from the «best» nonlinear model.

7. Conclusions

We think that the methodology proposed in this paper improves significantly the usual approach for modelling this kind of series. Its main characteristics are:

1) A satisfactory single equation model is achieved:

The usual procedure consists of splitting the sample into homogeneous subsamples and building separate models for each one of them. In our approach all observations are integrated in the same model, which has the following advantages:

- a) To simplify the management of the database and the forecasting procedure, an essential requirement for the model to be useful.
 - b) To analyze the behavior of the demand in the transitions from one type of day to another, from one season to another, etc., avoiding the bad performances that separate models show in transitions.
 - c) The aggregation of different special days into a smaller number of groups leads to more efficient estimators, and to a better forecast of the effect of an anomaly that happens for the first time.
- 2) The parsimonious way of obtaining: adaptative trend and seasonal, and stable effects for special days and weather variables.

Most models are specified in such a way that they are not flexible enough to capture the effect of changing conditions on the basic load or on the weekly seasonal: see for instance the univariate forecasting models in Engle *et al.* (1992). This is the main argument in the literature for considering a short sample: the less flexible the model is, the shorter the sample must be to avoid contaminating the estimates with old data.

We think that the procedure presented in this paper allows to extract all the advantages related to a large sample (greater efficiency) without its drawbacks (using information no longer relevant): estimated effects of special days and weather variables measure relative changes with respect to the present level of the series, and they are expected to remain more stable in the medium term than estimates of absolute changes; and the trend and seasonal components are modelled in such a way that they adapt quickly to the most recent information.

References

- Bogard, C.; George, G.; Jenkins, G. M. and McLeod, G. (1982): «Analysing a large number of energy time series for a utility company», in *Case Studies in Time Series Analysis*, Jenkins, G. M. and G. McLeod (eds.), Gwilym Jenkins & Partners Ltd., Lancaster, chapter 5.
- Bolzern, P. and Fronza, G. (1986): «Role of weather inputs in short-term forecasting of electric load», in *Electrical Power and Energy Systems*, Butterworth & Co Publishers, pp. 42-46.
- Box, G. E. P. and Tiao, G. C. (1975): «Intervention analysis with applications to economic and environmental problems», *Journal of the American Statistical Association* 70, pp. 70-79.
- Bunn, D. W. and Farmer, E. D. (1985): «Economic and operational context of electric load prediction», in *Comparative Models for Electrical Load Forecasting*, Bunn, D. W. and Farmer, E. D. (eds.), John Wiley & Sons, New York, chapter 1.
- Cancelo, J. R. and Espasa, A. (1991): «Forecasting daily demand for electricity with multiple input, nonlinear transfer function models», Working Paper 91-21, Departamento de Economía, Universidad Carlos III de Madrid.
- Cancelo, J. R. and Espasa, A. (1955): «Modelización del efecto temperatura en el consumo de electricidad: un ejercicio de búsqueda de especificación en relaciones dinámicas no lineales», *Estadística Española* 139, pp. 183-200.
- EPRI (1983): *Weather Normalization of Electricity Sales*, EA-3143, Palo Alto, CA.
- Engle, R. F.; Granger, C. W. J.; Rice, J. and Weiss, A. (1986): «Semiparametric estimates of the relation between weather and electricity sales», *Journal of the American Statistical Association* 81, pp. 310-320.
- Engle, R. F.; Mustafá, C. and Rice, J. (1992): «Modelling peak electricity demand», *Journal of Forecasting* 11, pp. 241-251.
- Ernoul, M. and Mattatia, R. (1984): «Prévision à court terme de la courbe de charge de la consommation d'énergie électrique: le modèle PRECOCE», *E.D.F. Bulletin de la Direction des Etudes et Recherches*, serie B, 3, pp. 37-44.
- Ernoul, M. and Meslier, F. (1982): «Analysis and forecast of electrical energy demand», *Revue Générale de l'Electricité* 4, pp. 257-278.
- Espasa, A. (1993): «Modelling daily series of economic activity», *Proceedings of the Business and Economic Statistics Section*, American Statistical Association, pp. 313-318.
- Espasa, A. and Cancelo, J. R. (1995): «Algunas consideraciones sobre la modelización de series diarias de actividad económica», *Actas de las X Jornadas de Economía Industrial*, pp. 195-201.
- Espasa, A.; Revuelta, J. M. and Cancelo, J. R. (1996): «Automatic modelling of daily series of economic activity», in *Proceedings in Computational Statistics COMPSTAT 1996*, Prat, A. (ed.), Physica-Verlag, Heidelberg, pp. 51-63.
- Fischler, E. B. and Nelson, R. F. (1986): «Integrating time-series and end-use methods to forecast electricity sales», *Journal of Forecasting* 5, pp. 15-30.
- Goh, T. N.; Ong, H. L. and Lee, Y. O. (1986): «A new approach to statistical forecasting of daily peak power demand», *Electrical Power Systems Research* 10, pp. 145-148.
- Gross, G. and Galiana, F. D. (1987): «Short-term load forecasting», *Proceedings of the IEEE* 75, pp. 1558-1573.
- Hesterberg, T. C. and Papalexopoulos, A. (1988): «Short term electric load forecasting using linear regression», *Proceedings of the Business and Economic Statistics Section*, American Statistical Association, pp. 608-612.
- Mood, A. M.; Graybill, F. A. and Boes, D. C. (1974): *Introduction to the Theory of Statistics*, McGraw-Hill, New York.

- O'Connor, M. J. and Kapoor, S. G. (1984): «Time series analysis of building electrical load», in *Time Series Analysis: Theory and Practice* 5, Anderson, O. D. (ed.), North-Holland, New York, pp. 135-144.
- Piggott, J. L. (1985): «Short-term forecasting at British Gas», in *Comparative Models for Electrical Load Forecasting*, Bunn, D. W. and Farmer, E. D. (eds.), John Wiley & Sons, New York, chapter 11.
- Scientific Computing Associates (1986): *The SCA Statistical System: Forecasting and Time Series Analysis*, Scientific Computing Associates Corp., Chicago.
- Schneider, A. M.; Takenawa, T. and Schiffman, D. A. (1985): «24-Hour electric utility load forecasting», in *Comparative Models for Electrical Load Forecasting*, Bunn, D. W. and Farmer, E. D. (eds.), John Wiley & Sons, New York, chapter 7.
- Spanos, A. (1986): *Statistical Foundations of Econometric Modelling*, Cambridge University Press, Cambridge.
- Thompson, R. P. (1976): «Weather sensitive electric demand and energy analysis on a large geographically diverse power system», *IEEE Transactions on Power Apparatus and Systems*, PAS-95, 1, pp. 385-393.

Resumen

En este trabajo se presenta un modelo para predecir el consumo diario de energía eléctrica en España. No se sigue la práctica habitual de dividir la muestra en submuestras homogéneas y construir modelos específicos para cada una de ellas, sino que se propone un modelo único para toda la muestra. Así pues, el trabajo se centra en cómo obtener un modelo uniecuacional aplicable a datos muy heterogéneos que se mantenga estable durante un período de tiempo relativamente largo. Esta aplicación también es de interés como ejemplo de una metodología general para la modelización de series diarias de actividad económica.

Recepción de original, octubre de 1995
Versión final, septiembre de 1996