APPLIED STOCHASTIC MODELS IN BUSINESS AND INDUSTRY
Appl. Stochastic Models Bus. Ind., 2006; **22**:113–125
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/asmb.622

The seasonal forecast of electricity demand: A hierarchical Bayesian model with climatological weather generator

Sergio Pezzulli[1,‡], Patrizio Frederic[2,*,†], Shanti Majithia[3], Sal Sabbagh[3], Emily Black[4], Rowan Sutton[4] and David Stephenson[4]

[1] *The University of Reading, School of Applied Statistics, Harry Pitt Building, Whiteknights Road, P.O. Box 240, Reading RG6 6FN, U.K.*
[2] *Dipartimento di Scienze Economia Politica, Facoltà di Economia, Università di Modena e Reggio Emilia, viale Berengario 11, 41100 Modena, Italy*
[3] *National Grid Transco (NGT) Group, U.K.*
[4] *Department of Meteorology, The University of Reading, U.K.*

SUMMARY

In this paper we focus on the one year ahead prediction of the electricity peak-demand daily trajectory during the winter season in Central England and Wales. We define a Bayesian hierarchical model for predicting the winter trajectories and present results based on the past observed weather. Thanks to the flexibility of the Bayesian approach, we are able to produce the marginal posterior distributions of all the predictands of interest. This is a fundamental progress with respect to the classical methods. The results are encouraging in both skill and representation of uncertainty. Further extensions are straightforward at least in principle. The main two of those consist in conditioning the weather generator model with respect to additional information like the knowledge of the first part of the winter and/or the seasonal weather forecast. Copyright © 2006 John Wiley & Sons, Ltd.

Accepted 10 October 2005

KEY WORDS: electricity load; mixed model; hierarchical Bayesian model; weather generator; seasonal forecast

1. INTRODUCTION

In Central England and Wales, winter peak-demand of electricity usually occurs between 5.00 pm and 5.30 pm. These data, among others, are routinely collected and investigated

*Correspondence to: P. Frederic, Dipartimento di Scienze Economia Politica, Facoltà di Economia, Università di Modena e Reggio Emilia, viale Berengario 11, 41100 Modena, Italy.
† E-mail: frederic.patrizio@unimore.it
‡ E-mail: S.Pezzulli@reading.ac.uk

Contract/grant sponsor: NERC; contract/grant number: NER/T/S/2002/00154: H1792500

by the National Grid Trasco group (NGT), operating the high-voltage electricity transmission system.

NGT operates the high voltage network and is responsible for residual short term system balancing 24 hours a day. To this aim, demand's prediction is essential. In particular the *seasonal* prediction, that is the forecast from one month up to one year ahead, is crucial for medium-term planning of energy production and trade. In collaboration with NGT, in this paper we focus on the one-year ahead prediction of the peak-demand daily trajectory of the winter season.

Different approach have been taken for electricity demand (or load) forecasting, including time-varying splines (see, e.g. Reference [1]), artificial neural networks [2] and multiple regression models [3, 4]. Most of the literature concerns short term forecasting, in which many different methods have been attempted (see, e.g. Reference [5]). The Bayesian approach [6, 7] has been used much less frequently and at our knowledge for short term forecast only.

The standard NGT seasonal model regresses the demand against a calendar cycle, three weather transforms and the Service sector index, a main component of the National Gross Product which is issued quarterly by the National Statistical Office. Obviously, services and weather series are unknown explanatory variables in extrapolating the demand prediction. One year ahead, the standard NGT seasonal forecast assumes an average (*climatological*) weather and the last year services. Later on, until six months ahead, a constant value is periodically added to the estimated demand's trajectory in order to take account of the expected growth of the services and of other economical factors.

A main drawback of this method is that the resulting forecasts cannot express the uncertainty on weather and services. For this aim, we define a Bayesian hierarchy as a generalization of the NGT regression model. For the weather, we assume a probability model based on past data: the climatological weather generator described in Section 3.1. For the Service series, finally, because it grows approximately linearly in time, we avoid the explicit use of the explanatory variable by including a random effect following a linear pattern (as shown in Section 3). In summary, the posterior distribution for next winter load trajectory is the mixture of the predictive of the mixed model averaged over the climatological weather distribution.

While for day-to-day prediction the engineering constraints produce a quadratic loss function, the NGT long-term forecasting, instead, is much more complex and subjected to change from year to year. The general aim is multipurpose. Several predictands can be relevant to the planner. The most important ones are the general level of the winter trajectory, the highest demand's value and its location in time. Thanks to the flexibility of the Bayesian approach, we are able to produce the marginal posterior distributions of those predictands. This is a fundamental progress achieved by the Bayesian procedure with respect to the classical methods. In other words, the primary aim of this paper is not to improve the performances of the NGT method. This is because the operational NGT model has been proven successful in several years of usage, so that the Bayesian model introduced here is based on essentially the same likelihood and the same climatological data. Rather, the aim is to show how the Bayesian approach can be used to include all the sources of uncertainty into the same framework, and to give a realistic representation of uncertainty in all the predictions needed for decision making.

After introducing the standard NGT regression model in the next section, we define, in Section 3, the Bayesian model. Results and conclusions are in Sections 4 and 5, respectively.

DOI: 10.1002/asmb

2. THE STANDARD NGT PREDICTION METHOD

2.1. The data

The available data include the daily peak demand fluxes (in MW) collected by NGT for $k = 17$ consecutive winters, from 1986–1987 to 2002–2003. Those trajectories start from about the last Sunday of October to the last Saturday of March. The last observed trajectory is represented in Figure 1(a). A weekly cycle can be clearly spotted, with less demand on the weekend due to less industrial activity. For the same reason there is a low during Christmas holidays.

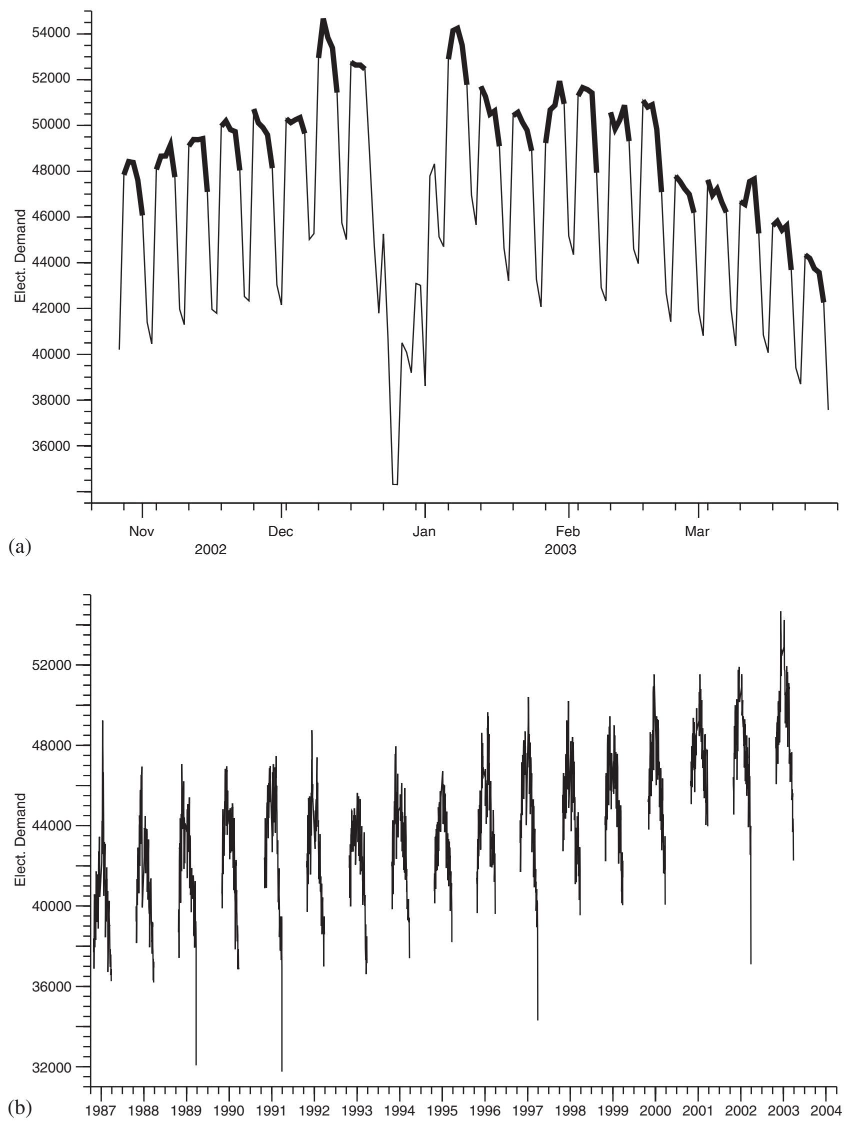

Figure 1. (a) The 2002–2003 winter trajectory of the daily peak electricity demand (MW); and (b) all the winter trajectories (used data only).

DOI: 10.1002/asmb

Because high demand is more important, weekends, bank holidays, Christmas holidays and other special days are omitted from the fitting process. Some anomalous day, like in case of severe storms causing network interruptions, and consumption fall, is also excluded in order to avoid erroneous contaminations of the results. The data actually used correspond to the bold part of the curve in Figure 1(a). These are also represented in Figure 1(b) for all available winters. Here, the general growth of the winter peak demand can be clearly observed.

On the other hand, even the form of the seasonal trajectory has changed in the last decades. Electricity demand is subjected to fast structural changes in time, because the consumer can choose different options within the energy sources. The weekly cycle is also changing because of the increasing activity of the service sector during the weekends. Therefore, the fitting process is usually restricted to four consecutive winters only.

2.2. *The NGT model*

The NGT model is a linear model whose explanatory variables are divided into 3 groups:

(i) the calendar component C: including dummies for Thursday and Friday and a winter cycle described by a cubic polynomial.
(ii) the economic component S: that is the series of the Service Sector Index.
(iii) the weather component W: including three opportune transforms of ground temperature, wind speed and ground solar radiation.

The original weather variables, provided by the UK Met Office, are weighted averages at few selected stations. The transforms are called effective temperature: TE; cooling power of the wind: CP; and (solar) effective illumination: EI. These transforms aim to depict how the external weather can affect the energy needs inside the buildings. The effective temperature, TE, is an exponentially smoothed form of another variable, TO, which is the mean spot temperature during the 4 hours before the peak time. The wind's cooling power, CP, is a non-linear transform of TO and wind speed. Finally, the effective illumination, EI, was originally a complex function of visibility, number and type of cloud and type of precipitation, but recently it has been replaced by a cubic function of the ground solar radiation. Calling:

$d_j(t)$ electricity demand
$\mathrm{Th}_j(t)$ dummy for thursday
$\mathrm{Fr}_j(t)$ dummy for friday
$\mathrm{TE}_j(t)$ effective temperature
$\mathrm{CP}_j(t)$ cooling power of the wind
$\mathrm{EI}_j(t)$ effective illumination

for day $t = 1, 2, \ldots, T$ ($T = 183$), of winter $j = 1, 2, \ldots, k$ ($k = 17$), the NGT model has the form

$$d_j(t) = \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \alpha_3 t^3 + \alpha_4 \mathrm{Th}_j(t) + \alpha_5 \mathrm{Fr}_j(t) + \beta S_j(t) + \gamma_1 \mathrm{TE}_j(t) + \gamma_2 \mathrm{CP}_j(t) + \gamma_3 \mathrm{EI}_j(t) + u_j(t) \quad (1)$$

$$= \alpha_0 + \boldsymbol{\alpha}' \mathbf{C}_j(t) + \beta S_j(t) + \boldsymbol{\gamma}' \mathbf{W}_j(t) + u_j(t) \quad (2)$$

where $\mathbf{C}_j(t) = (t, t^2, t^3, \mathrm{Th}_j(t), \mathrm{Fr}_j(t))'$ is the *cycle* vector, $S_j(t)$ is the Service Index, and $\mathbf{W}_j(t) = (\mathrm{TE}_j(t), \mathrm{CP}_j(t), \mathrm{EI}_j(t))'$ is the weather vector; α_0, $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_5)'$, β, and

DOI: 10.1002/asmb

$\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \gamma_3)'$ are the regression coefficients; and $u_j(t)$ is an ARMA modelled 'error' variable.

2.3. NGT basic and shifted forecasts

The *basic* NGT forecast is produced in late March, when the previous winter is just finishing. At that moment, not only weather and services are unknown for the next year, but the latter are unknown even for the current winter. In fact, the winter services series are issued in May by the Statistical Office. In order to solve this additional problem, the basic (one year ahead) forecast uses a simple trick that manages to use approximately the weather average and the last winter services as follows.

Let $\overline{\mathrm{TE}}(t)$, $\overline{\mathrm{CP}}(t)$ and $\overline{\mathrm{EI}}(t)$ be the average-trajectories of the weather variables for the winter day t. Any statistical measure of location can in principle be chosen, like the daily means or medians. Here we prefer the cubic splines shown in the next section.

Let $\overline{\mathbf{W}}(t) = (\overline{\mathrm{TE}}(t), \overline{\mathrm{CP}}(t), \overline{\mathrm{EI}}(t))'$. If k is the current winter and $\hat{\boldsymbol{\gamma}}$ is the current estimate of the weather coefficient $\boldsymbol{\gamma}$ (fitted on the winters from $k-1$ to $k-4$), let us consider the estimator of $d_{k+1}(t)$

$$D_{k+1}(t) = d_k(t) + \hat{\boldsymbol{\gamma}}'(\overline{\mathbf{W}}(t) - \mathbf{W}_k(t)) \tag{3}$$

Assuming model 1 and $\hat{\boldsymbol{\gamma}} \approx \boldsymbol{\gamma}$, should be

$$D_{k+1}(t) \approx \alpha_0 + \boldsymbol{\alpha}'\mathbf{C}_k(t) + \beta S_k(t) + \boldsymbol{\gamma}'\overline{\mathbf{W}}(t) + u_k(t)$$

Thus D_{k+1} is based on $S_k(t)$ although that series is unknown. Moreover, the weather component is fixed to the climatology and the last year error term replaces the next year error. Also, both the cyclic and economic coefficients enter (3) with their *true* coefficient values, that is a quite appealing property for an estimator.

Note, finally, that the winter cycle is aligned with the last winter, instead of the $(k+1)$th. In other words, there is $C_k(t)$ instead of $C_{k+1}(t)$. However, this is only a minor drawback because $D_{k+1}(t)$ can be shifted for few days in order to align the weekdays in C_k to the next year calendar. This final re-alignment gives the basic forecast: NGT0, say. It approximately assumes no economical growth, that is persistence (or stagnation) of economy, and average weather conditions. Then the inter-annual growth of the demand is periodically estimated by separate econometric modelling. More precisely, a constant inter-annual shift is predicted and summed up to the basic forecast.

Here we will consider the basic forecast NGT0 and the last shifted forecast, NGT1 say, which is computed just before the start of the winter, that is 6 month ahead in our terminology. Thus we retain the NGT standard methods under the worst and the best possible information states.

Pezzulli *et al.* [8] review the standard NGT procedure, where the ARMA modelling of the error term was not found much relevant. This motivates the following model.

3. THE BAYESIAN HIERARCHICAL MODEL

As shown in Figure 2, the Service index growth is approximately linearly in time. Therefore the linear part of the calendar cycle can adjust for the S contribution to the demand and makes S itself become redundant. This will be achieved by means of a random winter effect α_{0j} in our

DOI: 10.1002/asmb

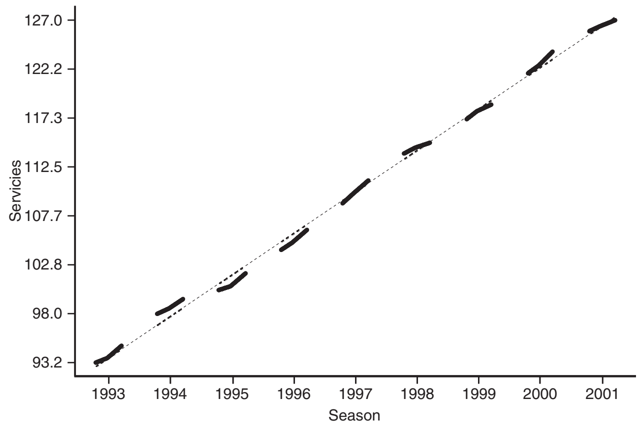

Figure 2. The Service Index trajectories for the last 9 observed winters (broken lines) versus an interpolating line (continuous line).

model. Therefore we assume:

$$(d_j(t)|\theta_j(t), \sigma) \sim \mathrm{N}(\theta_j(t); \sigma^2) \tag{4}$$

with

$$\theta_j(t) = \alpha_{0j} + \boldsymbol{\alpha}'\mathbf{C}_j(t) + \boldsymbol{\gamma}'\mathbf{W}_j(t)$$

and where α_{0j} represents the random effect for winter j

$$(\alpha_{0j}|\lambda_0, \lambda_1, \sigma_0) \sim \mathrm{N}(\lambda_0 + \lambda_1 \cdot j, \sigma_0^2)$$

For the computations we used WinBUGS (v. 1.4, Imperial College and Medical Research Council (U.K.)) software. As well known, WinBUGS requires proper priors [9], which invariably produces proper posteriors.

For stable numerical computations we worked on standardized variables. For the prior distribution, the regression coefficients $\boldsymbol{\alpha}$, $\boldsymbol{\gamma}$ and $\boldsymbol{\lambda}$ were assumed to follow independent Normal vague distribution (i.e. with zero mean and with variance 1000).

Moreover, the vector (σ^2, σ_0^2) were modelled independently from the regression coefficients by using the following specification. Let $\tau = 1/\sigma^2$ and $\tau_0 = 1/\sigma_0^2$. We fixed $\tau \sim \mathrm{G}(1, 0.1)$ which is a rather flat distribution with mean 10 and variance 100, and $(\tau_0|\tau) \sim \mathrm{G}(1, 1/\tau)$, having conditional mean and variance τ and τ^2, so that marginal mean and variance are 10 and 300. Thus, the resulting joint distribution reflects the vague prior belief that the two variance components have the same order of magnitude. The robustness of the results has been checked when the prior for τ is $\mathrm{G}(0.1, 0.1)$ or even the most conventional $\mathrm{G}(0.001, 0.001)$.

By the previous settings we can infer on all the unknown parameters and hyperparameters of the model. However, in order to predict the electricity demand for the next year (e.g. for year $k+1$) we need to define a (multivariate) probability distribution for the three weather vectors $\mathbf{TE}_{k+1}$, $\mathbf{CP}_{k+1}$, $\mathbf{EI}_{k+1}$ which are currently unknown. This weather distribution is an informative proper distribution depending on all the past available winters and is defined in the next subsection. The predictive distribution for the next winter trajectory is then evaluated by averaging

DOI: 10.1002/asmb

the demand formula (4) for year $k+1$ over all the unknown quantities, that is the parameters, the hyperparameters and the weather vectors.

3.1. The climatological weather generator

First, define the $3T$-dimensional composite-weather vector $\mathbf{W}_j$ such that

$$\mathbf{W}_j' = (\mathbf{TE}_j', \mathbf{CP}_j', \mathbf{EI}_j')$$

We assume the $\mathbf{W}_j$ exchangeable for all $j = 1, 2, \ldots, k+1$. This means, for example, that the effect of climate change is ignorable in such a short period of time, which is quite reasonable.

Secondly, we assume the $\mathbf{W}_j$ to be multivariate normally distributed

$$\mathbf{W}_j \sim \mathrm{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

The normality assumption is forced by the computational burdens, but it is also in agreement with our previous investigations [3]. This allows for a very complex structure in terms of means, variances and correlations. In fact, the number of parameters is so huge $p = 3T + 3T(3T+1)/2 = 151\,524$ that the probabilistic structure is even too much complex. Since we have observed just 17 winters, the maximum likelihood estimate of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ is strongly ill-conditioned. Moreover, this great flexibility of the multivariate normal family is against our prior knowledge. For example, it allows very different mean and variances for the same weather transform (e.g. TE) on two consecutive days. It also supports very different correlations when moving from the pair (t, t') of days to the pair $(t+1, t'+1)$. We know, instead, that some sort of quasi stationarity in the mean, in the variances and correlations should hold. Thus, in order to take account of the temporal structure of the weather trajectories, we used the following smoothed approach.

For the mean vector $\boldsymbol{\mu}$, we compose the 3 smoothing splines of the daily mean trajectories (with subjective choice of the smoothing parameters). The result is shown in Figure 3, where the raw daily means (e.g. the maximum likelihood solutions) can be compared to the smoothed estimates. Of course we have prior reasons to believe much more in the latter. The same method has been used for the variances. For the correlations, working separately on each of the 6 variance-blocks **TE** vs **TE**, **TE** vs **CP**, …, **EI** vs **EI**, the same smoothing spline estimation is used to 'correct' the empirical correlations of lag 1, then the lag 2 correlations and so on. At the end of this rather extensive process, positive definiteness is achieved by the Higham algorithm on the estimated correlation matrix [10].

In conclusion, we found that smoothing improves greatly the definiteness of the variance matrix. While the raw estimate has $k = 17$ positive eigenvalues out of $3T = 459$ dimensions of the matrix, the smoothed estimate identify more than 300 dimensions. The eigenvectors interpretation (not shown here) is also more clear and credible in the latter case. A comparison between raw and smoothed approaches is shown in Figures 4(a) and (b), respectively, for the **TE** vs **TE** correlation block. As shown in Reference [11], the assumption of continuity of the covariance kernel is essential for the predictability of the stochastic process. This is equivalent to impose regularity conditions on the weather trajectories (see, e.g. Reference [12]).

For simplicity, the weather generator has been fitted by means of all the available winters and then used to evaluate either the standard NGT and the Bayesian models in forecasting some of the last observed trajectories. This is not completely correct because we should use the weather observed in the previous winters only, excluding of course the targeted one. However, there are two main reasons for accepting this procedure. First, the (smoothed) weather generator does not

DOI: 10.1002/asmb

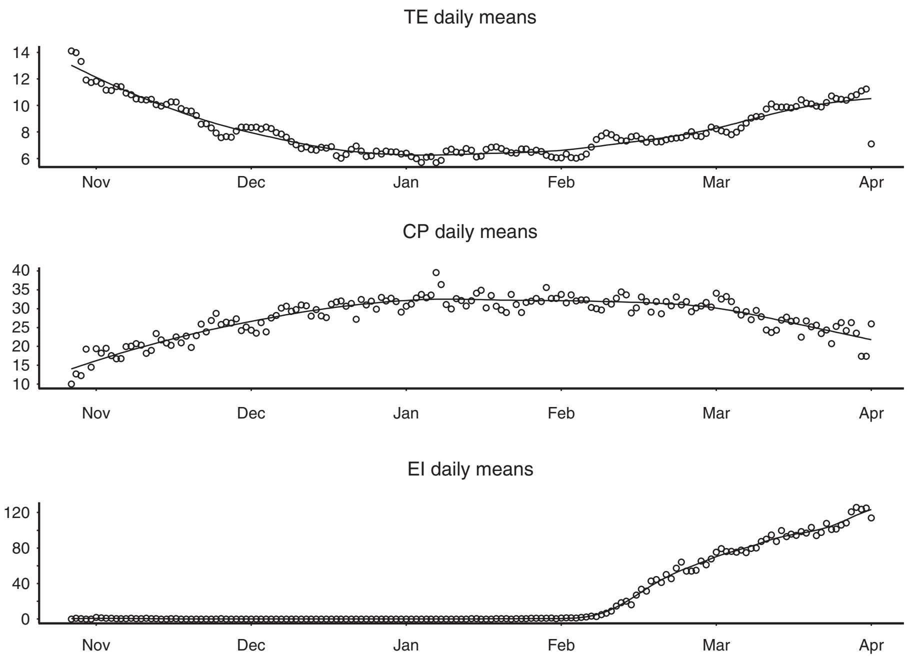

Figure 3. Observed daily means of effective temperature (TE), cooling power of the wind (CP) and effective illumination (EI), denoted by points, in comparison with the corresponding smoothing splines (lines).

sensibly change if we exclude the last few winters. Secondly, we are comparing the models under the same conditions, so that the relative performances are likely unchanged. The model parameters, on the other hand, are estimated over the last four winters that precede the targeted trajectory, either for the standard NGT methods and the Bayesian forecast.

Note that in order to forecast the next winter trajectory, the climatological weather generator enters the Bayesian model as a known distribution. However, we could also model our uncertainty about the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ thus depicting a more realistic position. This is computationally very expensive, however, and probably will have only a minor effect on the results and for these reasons it is not implemented in this paper. Finally, it is worth to mention that an alternative approach for constructing the weather generator could be based on dynamic linear modelling as in Reference [13].

4. RESULTS

We compared standard and Bayesian models on the last 6 winters. For each one, we used Splus (Insightful Corp.) software for computing the standard NGT predictions and WinBUGS for the Bayesian one. For the latter, we checked the convergence of the MCMC algorithm by the ANOVA based-method running in WinBUGS [14]. After this, a posterior sample of 1000 predicted trajectories has been collected and compared to the actual observations. The expected

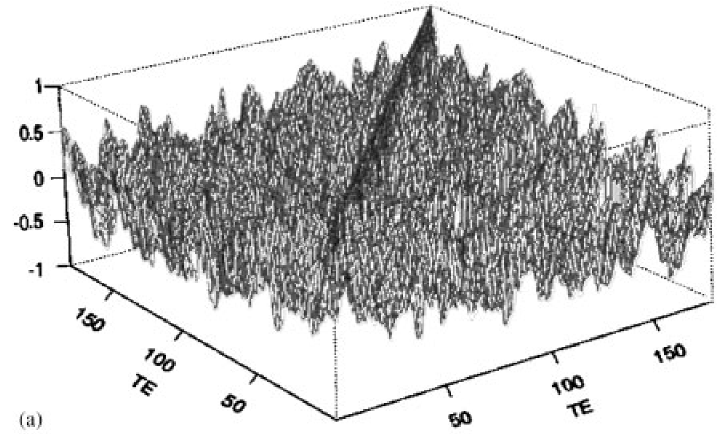

(a)

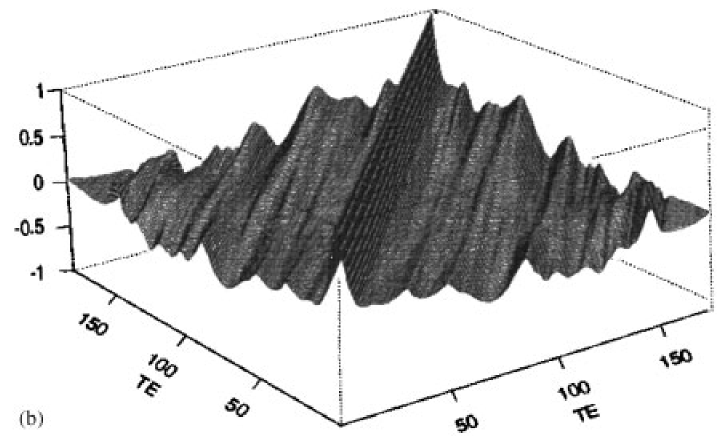

(b)

Figure 4. Raw correlation surface: (a) smooth correlation surface; and (b) for winter effective temperature (TE). Winter days on X and Y axes.

trajectory (composed by the daily averages of the posterior sample) has been used as the Bayesian forecast.

In Table I, we compare the Bayesian forecast with the basic NGT forecast (NGT0) and the shifted one (NGT1) over the last six observed winters. The bias error corresponds to the error in guessing the average winter demand. To this aim, we conclude from the table that both the Bayesian and the shifted-NGT methods are slightly improving the basic NGT forecast. Furthermore, all the methods provide an estimate whose error is less than 2% of the realized average winter load. This error percentage for the mean winter load is similar to short term errors for single day loads, while of course long term predictions can be worse on single days.

From Table I, the performances of NGT1 and Bayes seems rather close in terms of mean square error. Although relative efficiencies find some differences, those are very small in term of MW, which is likely due to the similarity between the underlying models. An important difference, however, is that the Bayesian forecast is available at the end of the previous winter while NGT1 is given just before the starting of the targeted season.

DOI: 10.1002/asmb

Table I. Forecasting performances of NGT0, NGT1 and Bayesian models. Bias and root mean square error (RMSE) units in GW. Efficiencies Eff0 and Eff1 of the Bayesian model versus NGT0 and NGT1 models are ratios of the form NGT-RMSE/Bayesian-RMSE in percent values.

	NGT0		NGT1		Bayes			
Winter	Bias	RMSE	Bias	RMSE	Bias	RMSE	Eff0	Eff1
1997–1998	0.3	1.1	0.0	1.1	−0.4	1.1	96	93
1998–1999	−0.2	1.0	−0.6	1.2	−0.7	1.2	82	93
1999–2000	0.5	0.9	0.6	1.0	−0.2	0.9	97	103
2000–2001	1.3	1.8	0.9	1.5	0.8	1.5	124	106
2001–2002	0.7	1.2	0.3	1.0	0.3	1.0	115	97
2002–2003	0.9	1.4	0.4	1.1	0.3	1.2	117	96

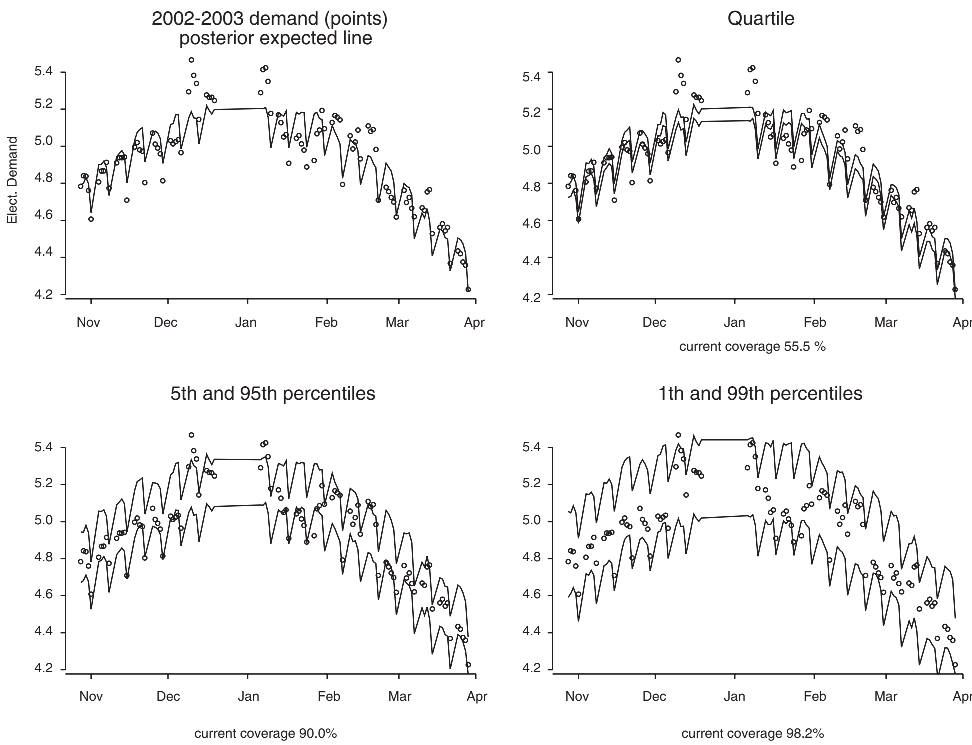

Figure 5. Bayes forecast and 50, 90 and 98% credibility bands (lines) against actual trajectory (points). Units in 10 GW.

A motivating improvement over the standard procedure is that we can assess the uncertainty around the expected trajectory by taking account of both uncertainty in the parameters and in the weather. In order to check the calibration of such a probabilistic forecast, we computed some reference quantile trajectories and then checked the observed coverage. For example, for any day we compute a 50% credibility interval and trace the upper and lower line for all the

DOI: 10.1002/asmb

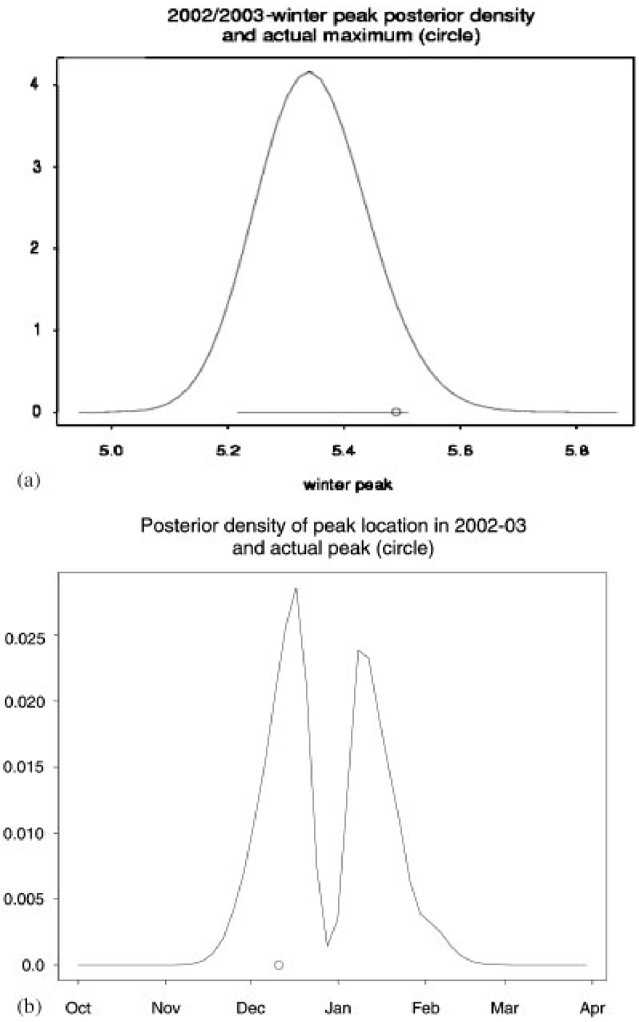

Figure 6. Kernel density representations of the marginal posterior distributions of the: (a) maximum winter demand intensity; and (b) time, compared to the occurred values. In (a), the horizontal line indicate the 95% credibility interval (10 GW units).

days of the winter, thus obtaining a 50% credibility band. Then the model is validated if the observations included in the band are about 50%. In Figure 5, for the year 2002–2003 we show the expected trajectory and the 50, 90 and 98 per cent bands (lines). The actual per cent of observations (points) inside those bands are 55.5, 90.0 and 98.2, respectively. Thus the predicted fractions are quite similar to the observed ones, and we can conclude that the spread of the posterior sample is reasonable. This is not so if we assume a fixed climatological weather, that essentially corresponds to use the confidence intervals on the standard NGT models. The 98% credibility band (not shown here) is about halved in size and contains less than 60% of the data.

DOI: 10.1002/asmb

A further gain of the Bayesian approach is that we can obtain posterior summaries for any function of the targeted load trajectory. Two important predictands for NGT regard the maximum of the trajectory. These are the intensity and the time of the winter maximum, which are useful for planning the peak production load. The marginal posterior densities for winter 2002–2003 are shown in Figures 6(a) and (b) for intensity and time, respectively. In order to improve the exposition we used the Rosenblat's kernel density method [15]. As shown by the circles, the predictions are rather accurate for that winter. Although on the edge, the load maximum is included in the 95% credibility interval. This happens for all the years considered. Those pictures are rather straightforward to present to the decision maker, who can probably make sensible use of the conveyed information.

5. CONCLUSIONS

We propose a Bayesian hierarchy with normal likelihood, vague priors on the parameters and hyperparameters and a weather generator probability model based on climatology.

The results are encouraging in both skill and representation of posterior uncertainty. The credible intervals for the demand trajectory show to cover a percentage of observations that is close enough to the expected percentage. Compared with the standard operative methods we do not need either the Service Index series and the separate econometric modelling that is used to update the basic NGT-operational forecast.

Also, the Bayesian forecast is available as soon as the previous winter is finished, that is one-year ahead, and can easily provide sensible predictions regarding any functional of the targeted trajectories, like the maximum value and its location in time.

The Bayesian approach is also extremely flexible for using new information. For example, after the starting of the winter, the weather generator can easily be modified in order to condition on the observed part of the weather trajectories.

A further development of our project is to find the opportune modification of the weather generator for conditioning on the seasonal forecast of the weather variables.

ACKNOWLEDGEMENTS

This research was funded by NERC COAPEC project, grant number NER/T/S/2002/00154: H1792500. We thank Krystyna Zietak (Wroclaw University of Technology) for providing suggestions and references on how to achieve the positive definiteness of the correlation matrix.

REFERENCES

1. Harvey AC, Koopman SJ. Forecasting hourly electricity demand using time-varying splines. *Journal of the American Statistical Association* 1993; **88**:1228–1236.
2. Connor JT. A robust neural network filter for electricity demand prediction. *Journal of Forecasting* 1996; **15**: 436–458.
3. Ramanathan R, Engle R, Granger CWJ, Vahid-Araghi F, Brace C. Short-run forecasts of electricity loads and peaks. *International Journal of Forecasting* 1997; **13**:161–174.
4. Taylor JW, Buizza R. Using weather ensemble predictions in electricity demand forecasting. *International Journal of Forecasting* 2003; **19**:57–70.

DOI: 10.1002/asmb

5. Moghram I, Rahman S. Analysis and evaluation of five short-term load forecasting techniques. *IEEE Transactions on Power Systems* 1989; **4**(4):1484–1491.
6. Douglas AP, Breipohl AM, Lee FN, Adapa R. The impacts of temperature forecast uncertainty on Bayesian load forecasting. *IEEE Transactions on Power Systems* 1998; **7**(3):1507–1513.
7. Cottet R, Smith M. Bayesian modeling and forecasting of intraday electricity load. *Journal of the American Statistical Association* 2003; **98**(464):839–849.
8. Pezzulli S, Sabbagh S, Majitia S. Electricity demand seasonal forecast at the National Grid Company: basic procedure and model validation. *NGT Internal Report*, 2003.
9. Spiegelhalter DJ, Thomas A, Best NG, Lunn D. *WinBUGS Version 1.4 User Manual January 2003*. MRC Biostatistics Unit: Cambridge, U.K., 2003. www.mrc-bsu.cam.ac.uk/bugs/winbugs
10. Higham NJ. Computing the nearest correlation matrix: a problem from finance. *Journal of Numerical Analysis* 2002; **22**:329–343.
11. Parzen E. An approach to time series analysis. *Annals of Mathematical Statistics* 1961; **32**:951–989.
12. Wahba G. *Spline Models for Observational Data*. SIAM: Philadelphia, 1990.
13. West M, Harrison PJ. *Bayesian Forecasting and Dynamic Models* (2nd edn). Springer: New York, 1997.
14. Brooks SP, Gelman A. Alternative methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics* 1998; **7**:434–455.
15. Rosenblatt M. Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics* 1956; **27**:832–837.

DOI: 10.1002/asmb