Dordonnat, V.; Koopman, S.J.; Ooms, M.; Dessertaine, A.; Collet, J.

Working Paper

An Hourly Periodic State Space Model for Modelling French National Electricity Load

Tinbergen Institute Discussion Paper, No. 08-008/4

Provided in Cooperation with:
Tinbergen Institute, Amsterdam and Rotterdam

TI 2008-008/4
Tinbergen Institute Discussion Paper

An Hourly Periodic State Space Model for Modelling French National Electricity Load

V. Dordonnat[a,b]
S.J. Koopman[a]
M. Ooms[a]
A. Dessertaine[b]
J. Collet[b]

[a] *VU University Amsterdam, the Netherlands;*
[b] *Electricité de France, Clamart, France.*

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

An Hourly Periodic State Space Model for Modelling French National Electricity Load

V. Dordonnat [a,b] S.J. Koopman [a] M. Ooms [a]
A. Dessertaine [b] J. Collet [b]

[a] *VU University Amsterdam, Department of Econometrics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands*

[b] *Electricité de France, 1 avenue du general de Gaulle, 92140 Clamart, France*

Abstract

We present a model for hourly electricity load forecasting based on stochastically time-varying processes that are designed to account for changes in customer behaviour and in utility production efficiencies. The model is periodic: it consists of different equations and different parameters for each hour of the day. Dependence between the equations is introduced by covariances between disturbances that drive the time-varying processes. The equations are estimated simultaneously. Our model consists of components that represent trends, seasons at different levels (yearly, weekly, daily, special days and holidays), short-term dynamics and weather regression effects including nonlinear functions for heating effects. The implementation of our forecasting procedure relies on the multivariate linear Gaussian state space framework and is applied to national French hourly electricity load. The analysis focuses on two hours, 9 AM and 12 AM, but forecasting results are presented for all twenty-four hours. Given the time series length of nine years of hourly observations, many features of our model can be readily estimated including yearly patterns and their time-varying nature. The empirical analysis involves an out-of-sample forecasting assessment up to seven days ahead. The one-day ahead forecasts from fourty-eight bivariate models are compared with twenty-four univariate models for all hours of the day. We find that the implied forecasting function strongly depends on the hour of the day.

Key words: Kalman filter; Maximum likelihood estimation; Seemingly Unrelated Regression Equations.

Preprint *16 January 2008*

1 Introduction

The need for accurate forecasting of electricity load goes from short to long-term. Short-term forecasting is important because the national grid requires a balance between produced and consumed electricity at any moment in the day. Long-term forecasting is relevant for the planning of new electricity utilities. Inaccurate forecasts have important financial costs. This paper aims to develop an effective new method for short-term forecasting of hourly electricity loads.

During the past years, many papers have been dedicated to methods and models for hourly load forecasting. Contributions can be distinguished between statistical models and exponential smoothing methods, between univariate models and models with explanatory variables, between linear models and nonlinear models. Earlier papers have developed both single-equation models and multiple-equation models with different equations for the different hours of the day. Both independent multiple-equation models and correlated multiple-equation models have been specified. The time dependence of hourly loads has been captured in observation driven VARIMA type models and in parameter driven models with unobserved components. In this paper we develop a forecasting model based on an interpretable decomposition of electricity loads in a trend effect, time-varying seasonal effects, calendar effects and weather dependent effects.

Our model is inspired by Ramanathan, Engle, Granger, Vahid-Araghi & Brace (1997) who built an extensive multiple regression model with separate forecasting equations for each hour of the day. Their observation driven model included calendar effects and weather effects and outperformed a wide range of alternative models in a forecasting competition. Taking part in the same competition, Harvey & Koopman (1993) developed an unobserved components model with time-varying splines to capture the evolution of intradaily seasonal patterns of hourly electricity loads, thereby integrating the equations for the different hours of the day. In a model for the New South Wales electricity load Cottet & Smith (2003) also used a multi-equation approach to capture the intradaily pattern and developed long and short-term forecast models within a Bayesian framework, but they assumed a diagonal vector autoregressive structure for the error terms. In our paper we follow Smith & Kohn (2002) in allowing cross-correlation between the stochastic terms of the equations for the different hours of the day.

Our model for the French load includes all well-known features in electricity consumption, see e.g. the papers in Bunn & Farmer (1985) : different levels of seasonality (yearly, weekly, daily), calendar events effects and weather dependence, see also Cancelo & Espasa (1996) who built a single equation model but for daily electricity loads thoroughly investigating the effect of spe-

cial days and the relationship between electricity and temperature. Bruhns, Deurveilher & Roy (2005) gave a detailed description of a non-linear forecasting model of French load in use at Electricite de France (EDF), also allowing for different levels of seasonality and weather dependence. In this paper, we present a different multiple-equation linear time-varying regression model for French national hourly electricity load, with one equation for each hour, like the approach of Ramanathan et al. (1997), and more recently, Soares & Souza (2006). We do not include periodic seasonal ARIMA components as they are difficult to interpret from an economic point of view. Instead, like Harvey & Koopman (1993) we capture the dynamics in a time-varying parameter regression in order to understand possible causes of changing trends and seasonal patterns. Pedregal & Young (2006) find periodic parameter changes at different frequencies in their analysis of twelve weeks of four-hourly load data in a dynamic harmonic regression model. As in Young, Pedregal & Tych (1999), they were unable to identify yearly movements. In addition to changes according to the hour of the day we discover a yearly pattern in the effect of temperature, which we partly model as a nonlinear heating effect.

We do not claim that univariate methods deliver bad forecasts. Structural univariate modelling of hourly demand has been suggested by Martin-Rodriguez & Caceres-Hernandez (2005), using unobserved component models and splines to capture the different levels of seasonality in the data. Other authors, like Soares & Medeiros (2005), prefer to build a model without weather variables, arguing that availability and accuracy of weather data forecasts can be problematic. Taylor, De Menezes & McSharry (2006) compared the forecasting performance of a wide range of univariate time series methods for intradaily load forecasting. Taylor & McSharry (2007) studied the effectiveness of these methods for forecasting hourly and halfhourly loads in 10 European countries in the period May-September 2005 and found the performance of recent univariate methods quite promising. Taylor & Buizza (2003) exploit different scenarios in temperature forecasting to estimate its effect on load forecasting uncertainty. Following Taylor & McSharry (2007) we use a simple univariate weekly random walk model as a benchmark in our forecast evaluation.

The first aim of our study is to examine the evolution of the effect of explanatory variables over a long period via the time-varying structure of our model. This evolution may be related to the gradual market penetration and slowly changing efficiency of electricity utilities, for example for the heating and cooling effects. These adjustments are also determined by changes in users' behaviour, for example for the Friday afternoon effect. Such changes are not independent across the different hours of the day, so we study the different hourly loads in one joint model allowing for cross-equation correlations in the innovations. The second modelling aim is to provide accurate short-term forecasts, from one day to one week ahead. The interpretation of the time-varying effects helps us in understanding possible forecast inaccuracies.

We show that our model fits in the multivariate linear Gaussian state space models framework. It implies that we can use Kalman filtering and associated algorithms to estimate the different components of electricity load and to do short-term forecasting. The model parameters are estimated by likelihood maximisation. The interpretation of the results is consistent with expert analysis from EDF, French national electricity producer and provider. Except for some special thresholds in the temperature effect we confine ourselves to linear methods, in contrast to Engle, Granger, Rice & Weiss (1986), Liu, Chen, Liu & Harris (2006), Cottet & Smith (2003) and Hippert, Bunn & Souza (2005), who consider semiparametric methods and artificial neural networks to model meteorological effects and seasonal patterns.

The plan of the remainder of this paper is as follows. Section 2 describes the dataset, Section 3 details the building of the model, describes the construction of the explanatory variables and relates the model to the linear state space framework. Section 4 presents the estimation results and interprets the various time-varying patterns. Section 5 discusses the absolute and relative forecasting performance. Section 6 concludes.

2 Description of French national hourly electricity loads

The dataset used in this study concerns French national hourly electricity consumption from September 1, 1995 until August 31, 2004. This hourly time series is for nine years and consists of $3,288$ daily or $78,912$ hourly observations. The dataset is compiled by Electricité de France (EDF) and is complete, that is, no missing observations are present. However, some days are intentionally considered as missing and excluded from the analysis in this study. On these days, which are known at EDF as EJP (Effacement Jour de Pointe: Peak Day Withdrawal) days, the load supply is subject to special tariffs. These financial incentives are introduced to cut heavy consumption. The EJP tariffs are activated when high levels of consumption and/or when problems with electricity supply (production) are expected to occur. They can only be set in place between November and March and for working days only. Clearly, on these days, the daily load and the hourly load curve are severely affected and standard models will overestimate electricity demand. Special treatments for the forecasting of the EJP days are outside the scope of this paper. The number of EJP days in our dataset is 249, that is 7.5% of the total number of days.

Different averages of French electricity consumption are presented in Figure 1. Display (a) shows the monthly averages for the complete sample and it shows that January consumption is highest and consumption in the holiday month of August is lowest. The other displays present the hourly averages for (b)

the complete sample, (c) the winter months and (d) the summer months. In all cases, the lowest electricity consumption is observed at 5 AM while the consumption is highest at 13 AM except in the winter months when at 7 PM the consumption is highest. Apart from the levels of consumption, the intra-daily electricity load curves in winter and summer are quite similar except in the early evening hours of 6 – 9 PM when consumption increases during the winter period only.

Apart from calendar information for holidays and other special days, the dataset also includes hourly temperature, cloud cover measures and one-day ahead forecasts of hourly temperature. The source of these data is Meteo France and they provide the data for different regions in France. Measures of cloud cover are based on human observations. Forecasts of cloud cover are not provided. EDF weights the temperature and cloud cover data for the different regions to construct a national average for hourly average temperature, their one-day forecasts and cloud cover. The use of weather variables in our forecasting model is regarded as crucial at EDF since much heating is generated by electricity in France. Approximately 28% of the private homes in France have electric heating.

Figure 2 provides some further graphical insights in our dataset. Display (a) presents the daily French national electricity consumption at 9 AM for the sample period of our dataset. The yearly seasonal cycle is clearly visible and a positive trend is also detectable from the data. These features also appear to various extents in the other hours of the day. Display (b) presents the hourly load for the year 2002. In this graph the weekly seasonality of electricity load becomes apparent as well as the effect of the summer holiday in August and the effect of the special tariff EJP days. The time series with the last three weeks of the hourly loads in our dataset, that is from Monday, August 8, 2004 until Sunday, August 29, 2004, is presented in Display (c). The magnitude of the weekly and daily load curves becomes apparent in this way. The load decreases during Saturday, Sunday and the night hours of all days in the week are distinct from each other and from the working day hours generally. Display (d) presents the well-known non-linear relationship between electricity load and average national temperature at 9 AM. The break of the regression curve appears to be at around 15 ^oC.

3 Model specification and parameter estimation

Denote $y_{t,i}$ as the electricity load at day t and hour i as measured in Megawatt. The basic model for the electricity loads that we consider in our study is a seemingly unrelated regression equations (SURE) model where each equation

represents a particular hour i. The hourly regression equation is given by

$$y_{t,i} = f_{t,i} + X'_{t,i}\beta_{t,i} + W'_t\gamma_{t,i} + \varepsilon_{t,i} \qquad \varepsilon_{t,i} \sim NID(0, \sigma^2_{\varepsilon,i}), \tag{1}$$

where $f_{t,i}$ is the trend component, $X_{t,i}$ is a vector of explanatory variables that change with day t and hour i whereas W_t is a vector of explanatory variables that only changes with day t, for $t = 1, \ldots, n$ and $i = 1, \ldots, k$ with k typically equal to 24 in the case of hourly data. Examples of variables in $X_{t,i}$ are temperature and cloud cover since these variables change with the hour. Examples of variables in W_t are dummies for day-type and dummies for holidays, these variables do not change with the hour. We note that (1) can also be considered for a subset of hourly data so that k can take a value of 2 when two specific hours are considered. The disturbance (or irregular) $\varepsilon_{t,i}$ is a random term with mean zero and variance σ_i^2 that can be different for different hours. Irregulars of different hours in the same day can also be correlated with each other, that is $E(\varepsilon_{t,i}\varepsilon_{t,j}) \neq 0$ for $i \neq j = 1, \ldots, k$. Irregulars of different days are not correlated, that is $E(\varepsilon_{t,i}\varepsilon_{s,j}) = 0$ for $t \neq s = 1, \ldots, n$ and $i, j = 1, \ldots, k$.

In case of

(a) the trend component $f_{t,i}$ is a deterministic function of time, for example, $f_{t,i} = a_i + b_i t$ with unknown and fixed regression coefficients a_i and b_i for $i = 1, \ldots, k$;

(b) the regression coefficients in vectors $\beta_{t,i}$ and $\gamma_{t,i}$ are unknown, fixed and the same for each day, that is $\beta_{t,i} = \beta_i$ and $\gamma_{t,i} = \gamma_i$, for $i = 1, \ldots, k$;

the SURE system (1) is standard and the estimation of the unknown regression coefficients can take place using generalised least squares methods. Since the regression coefficients are different for different hours (equations), we can refer to the SURE system as a periodic model, periodic in hours.

One focus of our study is the variation of regression coefficients over the days. The time variation can be explicitly modelled. For example, we can specify the trend component $f_{t,i}$ as a stochastic function of time, the details are given below. The regression coefficients in the vectors $\beta_{t,i}$ and $\gamma_{t,i}$ of (1) become time-varying when we specify these as random walk coefficients. In the remainder of this section we will discuss the details of the model that we adopt in our empirical study.

3.1 Stochastic trend component

The trend component $f_{t,i}$ represents the long-term changes in electricity consumption. A flexible stochastic specification of a time-varying trend compo-

nent is given by the local linear trend model

$$\begin{cases} f_{t+1,i} = f_{t,i} + g_{t,i} + v_{t,i}, & v_{t,i} \sim NID(0, \sigma^2_{v,i}), \quad i = 1, \ldots, k, \\ g_{t+1,i} = g_{t,i} + w_{t,i}, & w_{t,i} \sim NID(0, \sigma^2_{w,i}), \quad t = 1, \ldots, n, \end{cases} \tag{2}$$

where $v_{t,i}$ and $w_{t,i}$ are mutually and serially uncorrelated Gaussian noise terms with mean zero and variances $\sigma^2_{v,i}$ and $\sigma^2_{w,i}$, respectively, for $i = 1, \ldots, k$. The disturbances $v_{t,i}$ and $v_{s,j}$ can only be correlated for $i \neq j = 1, \ldots, k$ and $t = s = 1, \ldots, n$. This correlation is also allowed for $w_{t,i}$ and $w_{s,j}$. Special cases of the local linear trend model (2) include the random walk (with $\sigma^2_{w,i} = 0$ and $g_{1,i} = 0$), the random walk with fixed drift (with $\sigma^2_{w,i} = 0$ and $g_{1,i} \neq 0$), the integrated random walk (with $\sigma^2_{v,i} = 0$) and the linear fixed trend (with $\sigma^2_{v,i} = 0$ and $\sigma^2_{w,i} = 0$). The local linear trend model has $y_{t,i} = f_{t,i} + \varepsilon_{t,i}$ but this specification can be extended with other stochastic components for stationary (cyclical) processes and time-varying seasonal components. Such models are referred to as structural time series models or unobserved components time series models and are discussed at length in Harvey (1989). From this textbook treatment, we can, for example, learn that the forecasting function of the local linear trend model (2) is the well-known non-seasonal Holt-Winters method. The discount coefficients of this forecasting scheme are determined by the variances of the local linear trend model, $\sigma^2_{v,i}$ and $\sigma^2_{w,i}$.

3.2 Fixed and time-varying regressions for hourly explanatory variables

The hourly explanatory variables in $X_{t,i}$ concern weather variables that are based on temperature and cloud cover. In the model of the empirical study below we include four $X_{t,i}$ variables of which three are related to temperature and one is related to cloud cover. The three constructed temperature variables are designed to approximate the non-linear relationship between electricity load and temperature into a linear relationship. Denote $T_{t,i}$ as the national average temperature in ^oC at day t and hour i. The first three variables are based on $T_{t,i}$ and a smoothed temperature variable $T^{smo}_{t,i}$. Smoothed temperature is computed recursively by an exponentially weighted moving average of temperature $T_{t,i}$ of previous hours, that is

$$\begin{aligned} T^{smo}_{t,i+1} &= \kappa T^{smo}_{t,i} + (1-\kappa) T_{t,i+1}, & i = 1, \ldots, k-1, \\ T^{smo}_{t+1,1} &= \kappa T^{smo}_{t,k-1} + (1-\kappa) T_{t+1,1}, & i = k, \end{aligned} \tag{3}$$

with κ being typically close to 1, for example, 0.98. The smoothed temperature $T^{smo}_{t,i}$ is designed to take account of, for example, the physical inertia of

buildings. The first three variables in $X_{t,i}$ are constructed by

$$\begin{cases} X^1_{t,i} = \max(0, 15 - T_{t,i}), \\ X^2_{t,i} = \max(0, 15 - T^{smo}_{t,i}), \qquad i = 1, \ldots, k, \;\; t = 1, \ldots, n, \\ X^3_{t,i} = \max(0, T^{smo}_{t,i} - 18), \end{cases} \tag{4}$$

where we refer to $X^1_{t,i}$ as the heating-degrees variable, $X^2_{t,i}$ as the smoothed-heating-degrees variable and $X^3_{t,i}$ as the smoothed-cooling-degrees variable. The threshold temperatures for heating (15 ^{o}C) and cooling (18 ^{o}C) have been fixed at values determined internally at EDF. The last weather variable used in our model is the cloud cover variable $X^4_{t,i}$ that represents national cloud cover.

The vector of hourly explanatory variables is therefore

$$X_{t,i} = \left(X^1_{t,i} \; X^2_{t,i} \; X^3_{t,i} \; X^4_{t,i} \right)' .$$

The regression coefficients that determine the total hourly weather effect are collected in the vector of unknown coefficients

$$\beta_{t,i} = \left(\beta^1_{t,i} \; \beta^2_{t,i} \; \beta^3_{t,i} \; \beta^4_{t,i} \right)' .$$

A novelty in the modelling of hourly electricity loads is the evolution of these regression coefficients over time. The time-varying coefficients are modelled by

$$\begin{cases} \beta^j_{t,i} \quad = \beta^{*j}_{t,i} + \lambda^j_i X^j_{t-1,i}, \;\; j = 1, 2, \\ \beta^3_{t,i} \quad = \beta^{*3}_{t,i}, \\ \beta^{*j}_{t+1,i} = \beta^{*j}_{t,i} + u^j_{t,i}, \qquad j = 1, 2, 3, \end{cases} \tag{5}$$

where the disturbance $u^j_{t,i}$ is distributed as $NID(0, \sigma^2_{u^j,i})$, is serially uncorrelated and can only be correlated with $u^j_{t,m}$, for $t = 1, \ldots, n$, $i \neq m = 1, \ldots, k$ and $j = 1, 2, 3$. The coefficient $\beta^4_{t,i} = \beta^4_i$ for cloud cover is constant (it does not vary over the days t), for $i = 1, \ldots, k$, because there is no clear reasons why the effect of cloud cover should change over time. The fixed and unknown regression coefficient λ^j_i determines the dependence of the heating regression coefficient on the temperature of the previous day at the same hour, that is $X^j_{t-1,i}$ for $j = 1, 2$ and $i = 1, \ldots, k$. In this way we introduce a mild nonlinear temperature effect into the model since $X^j_{t,i}\beta^j_{t,i} = X^j_{t,i}\beta^{*j}_{t,i} + \lambda^j_i X^j_{t,i} X^j_{t-1,i}$. In case temperature does not change heavily between days, we have $X^j_{t,i} X^j_{t-1,i} \approx (X^j_{t,i})^2$. Furthermore, we introduce a yearly periodic dependence in the model. Since a time series of average temperature has a strong yearly cycle, the coefficient $\beta^j_{t,i}$ does also change with the yearly seasons of winter and summer temperatures,

$j = 1, 2$. This seasonal dependence of coefficients is also referred to as periodic. The model is periodic within the day (different coefficients for different hours) but also within the year (some coefficients depend on a yearly cycle via the temperature variable). Finally, the regression coefficient of the cooling effect is time-varying and modelled by a random walk process.

3.3 Fixed and time-varying regressions for daily calendar variables

The vector of explanatory variables that only changes by day (not by hour) is denoted by W_t and it is mainly concerned with the measurement of yearly, weekly and daily seasonal effects. With respect to the yearly seasonal effect in electricity load that is not captured by the temperature effects in $X_{t,i}$, we consider the following Fourier terms as explanatory variables:

$$a_{s,t} = \cos\left(\tau_t \frac{2\pi s}{365.25}\right), \quad b_{s,t} = \sin\left(\tau_t \frac{2\pi s}{365.25}\right), \quad s = 1, \ldots, 4, \tag{6}$$

where τ_t is the number of days elapsed since the 1st of January in the year in which day t falls for $t = 1, \ldots, n$. Moreover, we make a distinction between weekdays on the one side and weekends/holidays on the other side. For this purpose, we specify

$$a_{s,t}^{WD}, b_{s,t}^{WD} = \begin{cases} a_{s,t}, b_{s,t}, & \text{if day } t \text{ is a weekday;} \\ 0, & \text{if day } t \text{ is a weekend;} \end{cases} \tag{7}$$

and

$$a_{s,t}^{WE}, b_{s,t}^{WE} = \begin{cases} 0, & \text{if day } t \text{ is a weekday;} \\ a_{s,t}, b_{s,t}, & \text{if day } t \text{ is a weekend.} \end{cases} \tag{8}$$

As a result, the yearly cycle for electricity load is modelled by 4 Fourier series that requires 8 coefficients (for the cosine and sine parts) for the weekday yearly cycle and another 8 coefficients for the weekend yearly cycle. The variables $a_{s,t}^{WD}$, $b_{s,t}^{WD}$, $a_{s,t}^{WE}$ and $b_{s,t}^{WE}$ for $s = 1, 2, 3, 4$ are the first 16 explanatory variables in the vector W_t.

The typical weekdays of Tuesday, Wednesday and Thursdays (if not a holiday) are taken as the default day effect in the model and, obviously, no explanatory variable is introduced for this default day to avoid multicollinearity problems. For the weekly seasonal effect and other calendar effects, we introduce a range of dummy variables that correspond to different day types and are based on the operational practices at EDF, see Table 1.

The summer holiday period in France has a pronounced effect on electricity loads. The load levels decrease heavily in this period since many production

Table 1
Daily explanatory variables

$W_t^{1,\ldots,16}$	Fourier series for weekdays and weekends;
	Dummy variables for day types:
W_t^{17}	Mondays (if not a holiday or bridge day);
W_t^{18}	Fridays (if not a holiday or bridge day);
W_t^{19}	Saturdays;
W_t^{20}	Sundays;
W_t^{21}	Holiday (Easter Monday, Ascension Day, Whit Monday, May 1st, May 8th, July 14th, August 15th, November 1st, November 11th, if not a Saturday or Sunday);
W_t^{22}	December 25th;
W_t^{23}	January 1st;
W_t^{24}	December 24th (if not a bridge day);
W_t^{25}	Bridge day: Monday before a holiday or Friday after a holiday;
	Other effects:
W_t^{26}	August weekend trend 1: number of days since end of July;
W_t^{27}	August weekend trend 2: number of days since 2nd half August;
W_t^{28}	Dummy variable to indicate daylight saving period.

facilities are not operating at their full capacities and families live more outdoors. The load level decrease is gradual and can be characterized as follows. The differences in load levels between regular weekdays and weekends decrease progressively during the first half of the summer holiday period while the differences increase during the second half of the summer holidays. We model this effect with the following two variables. The first variable W_t^{26} is always zero except in the weekends of the last days of July and the first two weeks of August, when it equals the number of days since the last Friday in July. The second variable W_t^{27} is always zero except in August in weekend-days after the first two weeks of the month, when it equals the number of days since the last Friday in the second week of August. The last dummy variable of the model, W_t^{28}, is created for the daylight saving period. It distinguishes periods of winter-time and summer-time.

The values of these 28 calendar variables are collected in the vector

$$W_t = \left(W_t^1 \; W_t^2 \; \ldots \; W_t^{28} \right)' .$$

The unknown regression coefficients for the total calendar effect are in the

vector

$$\gamma_{t,i} = \left(\gamma_{t,i}^1 \; \gamma_{t,i}^2 \; \ldots \; \gamma_{t,i}^{28} \right)'.$$

We allow all these coefficients to change over time. However, it has been clear from the beginning of this study that the variables concerned with Christmas, New Year and daylight saving (W_t^j for $j = 22, 23, 24, 28$) can not be made time-varying since these variables refer to yearly events and our data set only spans a limited number of years. We therefore have

$$\begin{cases} \gamma_{t+1,i}^j = \gamma_{t,i}^j + e_{t,i}^j, & j = 1, \ldots, 21, 25, 26, 27, \\ \gamma_{t,i}^j \quad = \gamma_i^j, & j = 22, 23, 24, 28, \end{cases} \tag{9}$$

where the disturbance $e_{t,i}^j \sim N(0, \sigma_{e^j,i}^2)$ is serially uncorrelated and can only be correlated with $e_{t,m}^j$ for $i \neq m = 1, \ldots, k$, $j = 1, \ldots, 21, 25, 26, 27$ and $t = 1, \ldots, n$.

3.4 Vector representation of model with correlated errors

The time-varying regression model (1) can be formulated in vector form. Define $y_t = (y_{t,1}, \ldots, y_{t,k})'$ as the vector with hourly electricity loads for day t. The model for y_t is given by

$$y_t = f_t + X_t^* \beta_t + W_t^* \gamma_t + \varepsilon_t, \qquad \varepsilon_t \sim NID(0, \Sigma_\varepsilon), \tag{10}$$

where $f_t = (f_{t,1}, \ldots, f_{t,k})'$ is the vector of trends as modelled in (2) for $t = 1, \ldots, n$. The regression effects are represented by the parameter vectors β_t and γ_t with (fixed and time-varying) regression coefficients and the matrices X_t^* and W_t^* consisting of k rows with the explanatory variables in $X_{t,i}$ ($i = 1, \ldots, k$) and W_t, respectively. The specification of the coefficient vector β_t is implied by (5) and γ_t is implied by (9). The disturbance vector $\varepsilon_t = (\varepsilon_{t,1}, \ldots, \varepsilon_{t,k})'$ is serially uncorrelated. The variance matrix Σ_ε is possibly a full matrix such that the k equations in (10) can be correlated with each other. This also applies to the trends in f_t and to each time-varying parameter in β_t and γ_t. The disturbance vectors driving these multivariate dynamic processes have mean zero and full variance matrices. It is natural to assume that the regression effects for different hours of the day have similar or related impacts on the electricity loads for these hours. Hence we expect that these disturbances for different hours are correlated. The same arguments apply to the trend components. In case $k = 24$, variance matrices become relatively large and this obviously leads to many unknown parameters. Therefore, various restrictions on the variance matrices may need to be imposed when k is high. In our empirical study below, we will assume that Σ_ε is diagonal. The hourly periodic model in which we allow disturbances associated with

the stochastic processes for trend components and/or time-varying regression coefficients for different hours to be correlated is a novelty in the forecasting of electricity loads. We discuss estimation and signal extraction next.

3.5 Estimation, signal extraction and forecasting

The multivariate model (10) with the dynamic processes for f_t and elements in β_t and γ_t can be framed in a linear Gaussian state space model as given by

$$y_t = Z_t \alpha_t + \varepsilon_t, \qquad \alpha_{t+1} = T_t \alpha_t + R_t \eta_t, \qquad t = 1, \ldots, n, \tag{11}$$

where the vector of hourly electricity loads y_t and the disturbance vector ε_t are the same as in (10). The state vector α_t contains the trend components ($f_{t,i}$ and $g_{t,i}$) in (2) and the (partly) time-varying regression coefficients $\beta_{t,i}$ in (5) and $\gamma_{t,i}$ in (9). The dynamic processes of the trend components and the time-varying regression coefficients can be generally represented by the vector Markov process (or Vector Autoregressive process) for α_t in (11) as given by

$$\alpha_t = \Big(f_t' \; g_t' \; \beta_t' \; \gamma_t' \; \lambda' \Big)',$$

where $g_t = (g_{t,1}, \ldots, g_{t,k})'$ is the vector of slope (or growth) terms associated with $f_{t,i}$ in (2) and λ is the vector with elements λ_i^j in (5) for $j = 1, 2$ and $i = 1, \ldots, k$. The (partly) time-varying system matrices Z_t, T_t and R_t are fixed and known. For our model (10), we have

$$Z_t = \Big[I \; 0 \; X_t^* \; W_t^* \; 0 \Big], \qquad T_t = \begin{bmatrix} I & I & 0 & 0 & 0 \\ 0 & I & 0 & 0 & 0 \\ 0 & 0 & I & 0 & X_t^+ \\ 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & I \end{bmatrix},$$

where the matrix X_t^+ consists of zeroes and elements in $X_{t,i}^j$ to capture the terms $\lambda_i^j X_{t,i}^j$, for $j = 1, 2$ and $i = 1, \ldots, k$, in β_t of (5). The matrices I are identity matrices with appropriate but possibly different dimensions. The selection matrix R_t consists of zeros and ones; it links the appropriate elements of η_t with α_{t+1} for $t = 1, \ldots, n$. The vector η_t contains the disturbances associated with the trends, $v_{t,i}$ and $w_{t,i}$ in (2), with $u_{t,i}$, the innovations for the time-varying regression coefficients in (5) and with $e_{t,i}$ in (9) for $i = 1, \ldots, k$ and $t = 1, \ldots, n$. Since the corresponding disturbances of different hours are correlated, the variance matrix $\mathrm{Var}(\eta_t)$ is block-diagonal with $k \times k$ blocks of variance matrices.

The fixed regression coefficients are also placed in the state vector α_t. Fixed elements such as λ in α_t have the corresponding rows of R_t in (11) equal to zero. A more detailed discussion of how linear Gaussian time series models with regression effects can be formulated in state space is provided by Durbin & Koopman (2001, Chapter 3). Once the model is placed in state space form, the Kalman filter is used for predicting the state vector α_t and the associated smoothing algorithms produce estimates of α_t based on the whole sample. We have implemented filtering and smoothing algorithms that account for unknown (diffuse) initial conditions concerning α_1. The Kalman filter also computes the loglikelihood function that depends on the unknown fixed parameters in $\mathrm{Var}(\varepsilon_t)$ and $\mathrm{Var}(\eta_t)$. The loglikelihood maximization with respect to the unknown parameters in these variance matrices is based on the quasi-Newton optimization method and is typically a high-dimensional problem. We obtain good starting values via the Expectation Maximization (EM) algorithm, see Shumway & Stoffer (1982) and Koopman (1993) for further details. Furthermore, quasi-Newton methods rely on the score vector which can be evaluated numerically or analytically. In our framework, the analytical score is obtained from a single smoothing algorithm, see Koopman & Shephard (1992) for further details. We should note that the variance matrices $\mathrm{Var}(\varepsilon_t)$ and $\mathrm{Var}(\eta_t)$ are transformed such that they are always positive semi-definite and maximization takes place without further constraints.

Once the variance matrices $\mathrm{Var}(\varepsilon_t)$ and $\mathrm{Var}(\eta_t)$ are estimated, the Kalman filter and smoothing algorithms are used for signal extraction and allow us to draw time series plots of the estimated trend components as well as the estimated time-varying regression coefficients (both with associated standard errors). The fixed regression coefficients can also be calculated in this way. Since the Kalman filter equations can deal with missing observations in a natural way, forecasting is also straightforward in our framework. By extending the data sample $y_1, \ldots, y_n$ with missing values for y_t with $t = n+1, n+2, \ldots$ and by applying the Kalman filter to this extended sample, the forecasts are produced as by-products. Formally we define the forecast of y_{n+l} by $F_{n+\ell}$. The confidence interval of the forecast $F_{n+\ell}$ can be computed using $\mathrm{Var}(F_{n+\ell})$. The econometric computations have been implemented for the object-oriented matrix programming environment of `Ox`, see Doornik (2006), with state space routines from `SsfPack` as described in Koopman, Shephard & Doornik (1999).

4 Empirical results

In this section we report the results of the implementation of model (10), with $k = 2$, for a bivariate daily time series of electricity loads for the morning hour of 9 AM ($i = 1$) and the noon hour 12 AM ($i = 2$). The variance matrices for the bivariate trend components and each time-varying regression

coefficient are full. Only the variance matrix $\mathrm{Var}(\varepsilon_t)$ of the irregular ε_t is taken as diagonal. The estimation is based on the first eight years of the data from September 1, 1995 until August 31, 2003. The last year from September 1, 2003 until August 31, 2004, is used for post-sample forecast evaluation. The sample size of the two daily time series is $n = 2,922$. The post-sample length is 366 days.

4.1 Estimation results

The estimated standard deviations of the disturbances in (2), (5), (9) and (10) are presented in Table 2 together with their ratios with respect to the estimated standard deviation of the irregular (the so-called q-ratios). In case the q-ratio is large, say, $q > 1$, the component or regression coefficient varies heavily over time. Generally, the q-ratios are smaller for 12 AM than for 9 AM. This confirms the common belief that the load at noon is more predictable than the load in the morning hours. Large q-ratios are obtained for the time-varying cosine coefficients of the yearly cycle, all for weekdays. These values suggest that the yearly load cycles vary more for the weekdays than for the weekends. The time-varying holiday dummy coefficients have particular high q-ratios which will lead to some inaccurate forecasts for holidays in our model. It is interesting that all q-ratios are sufficiently large so that most coefficients do vary over time in our model.

The last column of Table 2 reports the disturbance correlations between 9 AM and 12 AM that are implied by the full variance matrix estimates. Whereas many standard deviations of the disturbances are clearly different for the two hours, the implied correlations are mostly estimated at values close to unity. The exceptions are the time-varying dummy effects for Saturdays and Sundays, with respective correlations of -0.70 and -0.30. This indicates substitution effects in the mornings of weekends. A relatively low electricity load at 9 AM is compensated with a higher electricity load at 12 AM and vice versa.

Table 3 reports the estimated λ coefficients and the fixed regression coefficient estimates of model (10). Almost all estimates are significant. The coefficients λ_i^j in (5) determine the importance of the nonlinearity and the yearly periodic effect in the time-varying regression coefficients of the two heating effects $X_{t,i}^j$ ($j = 1, 2$). All the estimated coefficients are significant, giving evidence that temperature effects are subject to yearly (periodic) nonlinear behaviour. The fixed regression estimates show that Cloud cover has a significant effect on the load, whereas daylight saving effect is not significant. The latter result could have been expected for 12 AM. The special effects for Christmas and New Year events are, unsuprisingly, highly significant.

4.2 In-sample signal extraction: trends and time-varying coefficients

Based on the parameter estimates as in Table 2, we apply the Kalman filter and smoothing algorithms to produce smoothed estimates of the state vector α_t that contains the trend components and the time-varying coefficients. Our time series plots of these estimates start only in January 1, 1997 because the initialisation period for the filtering and smoothing is somewhat unstable due to dummy variables that have zeroes for a long period of time (holiday effects). Figure 3 shows the estimated local linear trends for 9 AM (a) and 12 AM (b): they appear very smooth. They indicate that all other systematic effects in electricity loads have been captured by model (10). Even over a long period of various years, the underlying trends show no structural changes.

Figure 4 presents the time-varying regression coefficients for heating, smoothed heating and smoothed cooling at 9AM - displays (a) (c) and (e) - and 12 AM - displays (b), (d) and (f). The heating coefficients exhibit a seasonal pattern, most notably for the heating effect in displays (a) and (b). Here the coefficient is largest for the winter period while it increases from September onwards and it decreases from February onwards. Because temperature values also have a yearly cycle, this confirms the nonlinear and periodic nature of this time-varying coefficient. During the summer period the heating coefficients are implicitly interpolated because the corresponding explanatory variables are zero, giving no information about these coefficients. The associated standard errors are higher in the summer as a result. The seasonal patterns of the heating coefficients are an interesting and novel feature of our model.

The cooling regression coefficients in display (e) and (f) of Figure 4 increase slowly over time with noticeable changes between the summers of 1999, 2000 and 2001. These increases of the cooling effect may be attributed to the growing number of installations of cooling utilities for business in France, especially from 1999 onwards. The cooling coefficients are implicitly interpolated during colder periods when the explanatory cooling variable is zero. Since the number of days with a smoothed temperature larger than 18^oC is relatively small and the significance levels of the cooling coefficients are lower than those of the heating coefficients, the standard errors are relatively constant through summer and winter periods. It is a satisfactory empirical finding that the time-varying cooling coefficients show a clear upward trend in the period when air conditioning became more intensively used.

Figure 5 shows the estimated time-varying heating and cooling effects on the electricity load, that is $\beta^j_{t,i}X^j_{t,i}$ for $j = 1, 2, 3$ in displays (a),(c),(e) for 9 AM, i.e. $i = 1$ (respectively (b),(d),(f) for 12 AM, i.e. $i = 2$). Naturally, the heating effects on the load are most pronounced in the winter periods. The smoothed heating effects in display (c) and (d) have a clearer impact on the load than

the actual heating effects in display (a) and (b). The cooling effects in display (e) and (f) have a lower impact. However, the alternating heating and cooling effects in winter and summer periods are clear from Figure 5.

Figures 6 and 7 show the estimated effects for the different day types of (a) Mondays, (b) Fridays, (c) Saturdays, (d) Sundays, (e) Holidays and (f) Bridge days, at 9 AM and 12 AM, respectively. The effects are negative for all these dayt types and therefore the levels for those days are lower compared to the regular (default) days of Tuesday, Wednesday and Thursday, when there are more business activities. For example, the Monday effect is -1100 MWh at 9 AM but only -525 MWh at 12 AM. For the other day type effects, the differences between the two hours are smaller. In general, the day type effects for 9 AM are stronger than for noon. However, the Friday effect for 12 AM is stronger than for 9 AM and it becomes stronger over time: from -125 MWh to around -250 MWh at the end of the summer 2003, see Figure 7, display (b). This change may be explained by the decrease of the official number of working hours in France. The day type effects are, not surprisingly, most pronounced for weekends and holidays. The Saturday effect decreases slowly from a around -8000 to -9500 MWh at 9 AM. A similar long-term decrease from -7000 to -7250 MWh is observed at 12 AM but this change is much smoother over time, compare also the q-ratios for W_t^{19} in Table 2. The Sunday effect at 9AM varies between -13000 MWh and -15500 MWh. The strong variation in the Sunday effects may be due to the various holidays that occur around a Sunday. Special modelling of such effects may be considered. The significant Bridge day effect turns out to be rather constant over time. The holiday effects are as important as the Sunday effect and also relatively constant.

Figure 8 presents the global yearly effects for (a) 9 AM and (b) 12 AM. The yearly effect consists of the impact via the time-varying Fourier coefficients (separated for weekdays and weekends) together with the daylight saving (fixed effect) and August trend effects, i.e. $\gamma_i^{28} W_t^{28} + \sum_{j \in \{1,..,16,26,27\}} \gamma_{t,i}^j W_t^j$, $i = 1, 2$. The time-varying structure of the model enables the coefficients to adapt to periods with a fast change in the yearly pattern, especially at the end of the year and in August. The result is a more parsimonious model since it avoids the inclusion of more special dummy variables to capture these specific effects. This adaptiveness is captured by large q-ratios for some of the Fourier coefficients in Table 2.

4.3 In-sample diagnostics

The standardised prediction errors are obtained from the Kalman filter. When the model (10) is well-specified, the standardised prediction errors are serially independent and normally distributed. Various diagnostics can be used to

check whether the prediction errors can be regarded as independent random deviates from a standard normal distribution. Since our main focus is on forecasting, we concentrate on the dynamic features of the prediction errors. A particular interesting diagnostic is the sample autocorrelation function of the (in-sample) standardised prediction errors. Figure 9 displays the correlogram for lags 1 up to 365 for (a) 9 AM and (b) 12 AM. The correlations are mostly within the 95% confidence bands which is satisfactory. For low lags, the correlations are outside the interval but we view these values as acceptable. The one exception is the correlation at lag 365 (the one-year lag) which shows that the model does not capture all dynamics well with respect to the yearly cycle.

5 Forecasting performance

The forecasting performance of model (10) is investigated for the post-sample observations from September 1, 2003 until August 31, 2004, using the parameter estimates of Section 4. Since model (10) includes weather variables as explanatory variables, the forecasting accuracy can be based either on realised hourly values of the weather variables or on on their one-day ahead forecasts. The former may be preferred to avoid discussing external inaccuracies due to weather forecast errors, while the latter is preferable to compare models as they would be used in real situations. Since we have both realised temperature and one-day ahead temperature forecasts available in our data base, we evaluate the forecasts of our model for both situations. Parameter estimation and Kalman filter updating are not affected by using temperature forecasts since at EDF the realised temperature is always available on the next day. The one-day ahead prediction error based on realised temperature can then be computed and the likelihood function can be adjusted. We do not have the forecasts of the cloud cover in our data base, so here we use realised data only.

We use the mean absolute percentage forecast error (MAPE), the root mean squared forecast error (RMSE), and the mean percentage forecast error (MPE) to assess forecasting performance. For hour i, they are given by:

$$\begin{aligned} \mathrm{MAPE}(i) &= N^{-1} \textstyle\sum_{t=1}^{N} 100 |E_{t,i}^{h} \;/\; y_{t,i}|, \\ \mathrm{RMSE}(i) &= \sqrt{N^{-1} \textstyle\sum_{t=1}^{N} (E_{t,i}^{h})^2}, \\ \mathrm{MPE}(i) &= N^{-1} \textstyle\sum_{t=1}^{N} 100 E_{t,i}^{h} \;/\; y_{t,i}, \end{aligned} \tag{12}$$

where $F_{t,i}^{h}$ and $E_{t,i}^{h} = y_{t,i} - F_{t,i}^{h}$ are the actual forecast and forecast error (not standardised), respectively, at day t and hour i for $t = n+1, \ldots, n+N$ and $i = 1, 2$ with N as the number of available forecasts, h being the forecasting horizon, in our case $h = 1, \ldots, 7$.

We first analyse the results of one-day ahead forecasts for the hours of 9 AM and 12 AM. This is particularly interesting since maximum likelihood estimation optimizes one-step ahead prediction errors. Then we look at forecasting the hourly loads up to 7 days ahead for 9 AM and 12 AM. Finally we present and analyze one-day ahead forecasts for all 24 hours based on one univariate model and two bivariate models.

5.1 One-day ahead forecasts for 9 AM and 12 AM

Figure 10 presents the one-day ahead forecast errors for the electricity load at (a) 9 AM and (b) 12 AM as well as (c) their standard errors for non-EJP days. On the whole, the forecasts seem unbiased. The largest forecast errors correspond to holidays (in November, December, May, July and August). Standard errors are large during winter periods and weekends but particular large values are obtained for specific holidays and August weekends. These large standard errors are typically associated with effects for which only a few observations are available for estimation.

Figure 11 presents the correlogram of the daily prediction errors for the post-sample observations at (a) 9 AM and (b) 12 AM. The correlations for the lags 1 up to 7 are somewhat larger compared to the in-sample correlogram but they are not significant. However, we do not find a pattern that leads us to believe that we have missed a structural dynamic feature in the time series. The correlogram values for higher order lags vanish to zero since the number of values that can be used to compute the higher lag correlations becomes smaller (the post-sample period consists of 366 days). Therefore we cannot use Figure 11 to comment on the long-term forecasting ability of the model.

Table 5 presents the overall MAPEs for one-day ahead forecasts of non-EJP days for 9 AM and 12 AM. The MAPE for our full model is 1.34% and 1.31% when using realised temperatures, and 1.44% and 1.50% when using temperature forecasts. To place these measures in perspective, we also present the results for separate subperiods and we also consider four benchmark models. The first benchmark model is the weekly random walk (RW). This forecast method was one of the best benchmarks in the study by Taylor & McSharry (2007). Therefore we also report the MAPE for the RW. The basic forecast function of the RW in our study is simply $F^h_{t,i} = y_{t-7,i}$, $h = 1, .., 7$. The observed load of a week ago at the same hour is the forecast for today. We take the value of two weeks ago if there was a holiday one week ago. Various problems arise with special days including holidays. We have deleted these forecasts for the RW since we only need the RW to serve as a benchmark. We consider three restricted variants of our general model as additional benchmarks. In summary, we consider five different forecasting models:

RW: Weekly random walk;

Reg: Model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = \sigma_{u^j,i} = \sigma_{e^j,i} = 0$ in (2),(5) and (9) ;

Univ: Univariate version of model (10), i.e. $k = 1$ (trend components and regression coefficients time-varying);

TVR: Model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = 0$ in (2), (regression coefficients time-varying);

TTR: Model (10) as it is, with $k = 2$ (trend components and regression coefficients time-varying).

The overall forecasting performances for non-EJP days of all five models are reported in Table 5, with separate evaluations by day type.

We first discuss the forecasting results of Table 4 where we report the MAPE as defined in (12) for the five models separately for each month of the forecasting period and where we compare forecasts based on realised temperature (left part of the table) with those based on one-day ahead temperature forecasts (right part of the table). The RW is obviously not affected by this choice of temperature values. The number of forecast errors for each month N is indicated in the third column of Table 4. The number of forecasts produced by RW may be less than N. The ninth column gives the RMSE of the one-day ahead temperature forecasts. The most interesting aspects of Table 4 are:

- The RW is performing rather poorly for all months and both hours except for June, July and, in a less pronounced way, September. However, the RW only outperforms the fixed regression (Reg) model for these months.
- Te univariate model, Univ, outperforms both RW and Reg, with the exception of February where Reg is the best for both 9 AM and 12 AM.
- The TVR and TTR models outperform the three other models overall. Time-varying trends do not necessarily lead to better forecasts. Most of the differences between the MAPEs of TTR and TVR models are small.
- The forecasting results for TTR and TVR are disappointing for January, May, August and December. The most accurate forecasts are for June and September. Forecasting results for both hours of the day are comparable.

Comparing the forecasting accuracy using realised temperatures and forecasted temperatures, we find the following. With respect to the effect of temperature forecast errors on load forecast errors, we note that the one day ahead temperature forecast RMSE of 1.16^oC in March at 9 AM seems to generate a large MAPE for TVR and TTR and, to a lesser extent, for Univ. A similar effect is observed in November and December. In April at 9 AM and in February at 12 AM, Univ is more affected by the temperature forecast errors than TVR and TTR. These findings illustrate the importance of temperature

forecast accuracy. However, temperature forecast errors have a smaller impact on forecasting accuracy in the summer months. It confirms that cooling effects have a smaller impact on electricity load than heating effects. Forecasts from time-varying regression models still outperform those of Reg and RW when the forecasts are based on one-day ahead temperature forecasts.

Table 5 presents the MAPEs for each day type for 9 AM and 12 AM. It reports the forecast results for all special days (holidays, bridge days) together in one category. For obvious reasons, forecasts for the RW model are missing for these days. Interesting findings of Table 5 are:

- The holiday loads are the most difficult to forecast since they come in small numbers and the observed loads vary more than on other days.
- Overall the forecast accuracy measures for Univ, TVR and TTR are smaller than for Reg and RW. More specifically, the forecasts of TVR and TTR outperform those of Univ at 9 AM. At 12 AM, the forecasts of the univariate model Univ generally outperform those of TVR and TTR.
- Default days and Fridays (Saturdays) are forecasted most accurately at 9 AM (12 AM). Loads on Mondays and Sundays are more difficult to predict.
- The forecast accuracies obtained for 9 AM and 12 AM are comparable.

Qualitatively, these findings do not alter when forecasts are based on realised temperatures, rather than one-day ahead temperature forecasts. For completeness, Table 5 also reports the mean percentage forecast error (MPE) for EJP days. The EJP days are treated as missing for the estimation of parameters in Section 4 but the Kalman filter can still produce forecasts for these days. The bias in forecasting the EJP days is clear and we conclude that the model systematically over-estimates the realised electricity loads for these days.

Overall we are satisfied with the post-sample forecasting performance of our model. We have shown that time-varying and periodic regression effects are important in accurately forecasting hourly loads.

5.2 Multi-day ahead forecasts

Table 6 shows the forecast precisions for multiple day ahead forecasts (one to seven days) and for 9 AM and 12 AM. The RW has the same MAPE values for all horizons because the one-step ahead forecast is based on the load of a week ago. The forecasts in this table are computed using realised values of weather variables as forecasts for the 7-day horizon were not available in our data set.

For all forecast horizons, the RW does badly in terms of MAPE, both at 9 AM and 12 AM. All time-varying regression models perform better than model

Reg upto five days ahead. TTR is best at 9 AM for one upto five days ahead forecasting. At 12 AM, the Univ model is best for one upto six days ahead forecasting. Model Reg is best for six and seven days ahead forecasting at 9 AM and for seven days ahead forecasting at 12 AM. These results confirm that our time-varying model is primarily designed for short-term forecasting.

5.3 One-day ahead forecasts comparison for all hours

For a more general assessment of our methodology, we consider forecasting all twenty-four hours of the day based on the models

Univ	:	model (10) with $k = 1$;
TTR-1	:	model (10) with $k = 2$, for consecutive hours $(i-1, i)$;
TTR+1	:	model (10) with $k = 2$, for consecutive hours $(i, i+1)$,

where $i = 0, \ldots, 23$ and $i - 1 = -1$ refers to the last hour of the previous day.

In Table 7 we compare the one-day ahead forecasting accuracy for each hour i, with $i = 0, 1, \ldots, 23$, in terms of RMSE and MAPE. The load forecasts are based on one day ahead temperature forecasts. Only regular days (non-holidays, non-bridge-days, non-EJP days) are considered. In terms of RMSE, Univ outperforms the bivariate models for seven hours: 13 upto 17, 21 and 22. However, in terms of MAPE Univ outperforms those models only for two hours: 1 AM and 2 PM. We prefer the MAPE over the RMSE as it less sensitive to outliers and easier to compare across different applications.

In view of these forecasting results, we comment on the estimated correlations in the bivariate models, which we do not present here to save space. In the morning hours, the forecasts produced by the bivariate models are better than those produced by the univariate models. The estimated disturbance correlations for each morning hour are all very close to one which may indicate that these high correlations lead to more accurate forecasts. Correlations for the disturbances for heating, smoothed heating, smoothed cooling and level components are all close to unity in all bivariate models for all 24 hours. These estimation results will be useful in the specification search for a more parsimonious forecasting model for all 24 hours.

We conclude from Table 7 that a bivariate modelling approach where we allow strong correlations between time-varying regression effects for different hours can improve general forecasting performance.

6 Conclusion

We present a linear multivariate periodic state space model for the forecasting of hourly electricity loads. The model includes a stochastic trend component together with fixed and time-varying regression effects. Each equation in the model is associated with a specific hour and has different coefficients and different time-varying processes which are possibly correlated through the disturbances that drive them. Kalman filter methods are used for estimation, signal extraction and forecasting.

Our linear Gaussian time series model has a relatively simple structure: it can be multivariate, it has trend components and regression effects (fixed and time-varying). The EDF data set of French national loads consists of a long time series with hourly observations for loads, temperature, cloud cover and one-day ahead temperature forecasts. We capture interesting trends and time-varying regression coefficients from our empirical study. Some of our empirical findings have been known by experts at EDF but have not been properly measured earlier. For example, the slow increase of the cooling effect on loads, the yearly patterns in the heating regression coefficients and the strong correlations between the effects for different hours (9 AM and 12 AM).

The main purpose of our model is however short-term load forecasting. The forecasting results are satisfactory for 1, 2 and 3 days ahead. Some improvements can be made for longer forecast horizons. We may need to focus on finding more appropriate dynamic specifications for the intra-yearly variations in loads. The next challenge of our periodic model-based time-varying parameter approach to forecasting loads is to extend the model to more than two hours estimating all components simultaneously. We should be able to find parsimonious formulations for multivariate models, because we have found that load components of different hours are highly correlated.

References

Bruhns, A., Deurveilher, G. & Roy, J. S. (2005), A non linear regression model for mid-term load forecasting and improvements in seasonality, *in* 'Proceedings of the 15th Power Systems Computation Conference 2005, Liege Belgium'.

Bunn, D. W. & Farmer, E. D., eds (1985), *Comparative Models for Electrical Load Forecasting*, John Wiley, New York.

Cancelo, J. R. & Espasa, A. (1996), 'Modelling and forecasting daily series of electricity demand', *Investigaciones economicas* **20**, 359–376.

Cottet, R. & Smith, M. (2003), 'Bayesian modeling and forecasting of intraday electricity load', *Journal of the American Statistical Association* **98**, 839–849.

Doornik, J. A. (2006), *Ox: An Object-Oriented Matrix Language*, Timberlake Consultants Press, London, U. K.

Durbin, J. & Koopman, S. J. (2001), *Time Series Analysis by State Space Methods*, Oxford University Press, Oxford, UK.

Engle, R. F., Granger, C. W. J., Rice, J. & Weiss, A. (1986), 'Semiparametric estimates of the relation between weather and electricity', *Journal of the American Statistical Association* **81**, 310–320.

Harvey, A. C. (1989), *Forecasting, structural time series models and the Kalman Filter*, Cambridge University Press, Cambridge, UK.

Harvey, A. C. & Koopman, S. J. (1993), 'Forecasting hourly electricity demand using time-varying splines', *Journal of the American Statistical Association* **88**, 1228–1237.

Hippert, H. S., Bunn, D. W. & Souza, R. W. (2005), 'Large neural networks for electricity load forecasting: Are they overfitted?', *International Journal of Forecasting* **21**, 425–434.

Koopman, S. J. (1993), 'Disturbance smoother for state space models', *Biometrika* **80**, 117–126.

Koopman, S. J. & Shephard, N. (1992), 'Exact score for time series models in state space form', *Biometrika* **79**, 823–826.

Koopman, S. J., Shephard, N. & Doornik, J. A. (1999), 'Statistical algorithms for models in state space using SsfPack 2.2', *The Econometrics Journal* **2**, 107–160, www.ssfpack.com.

Liu, J. M., Chen, R., Liu, L.-M. & Harris, J. L. (2006), 'A semi-parametric time series approach in modeling hourly electricity loads', *Journal of Forecasting* **25**, 537–559.

Martin-Rodriguez, G. & Caceres-Hernandez, J. J. (2005), 'Modelling the hourly Spanish electricity demand', *Economic Modelling* **22**, 551–569.

Pedregal, D. J. & Young, P. C. (2006), 'Modulated cycles, an approach to modelling periodic components from rapidly sampled data', *International Journal of Forecasting* **22**, 181–194.

Ramanathan, R., Engle, R., Granger, C. W. J., Vahid-Araghi, F. & Brace, C. (1997), 'Short-run forecasts of electricity loads and peaks', *International Journal of Forecasting* **13**, 161–174.

Shumway, R. H. & Stoffer, D. S. (1982), 'An approach to time series smoothing and forecasting using the EM algorithm', *Journal of Time Series Analysis* **3**, 253–263.

Smith, M. & Kohn, R. (2002), 'Parsimonious covariance matrix estimation for longitudinal data', *Journal of the American Statistical Association* **97**, 1141–1153.

Soares, L. J. & Medeiros, M. C. (2005), Modeling and forecasting short-term electricity load demand: A two-step approach methodology, Technical Report Textos para discussão 495, Department of Economics, PUC, Rio de Janeiro, Brazil.

Soares, L. J. & Souza, L. R. (2006), 'Forecasting electricity demand using generalized long memory', *International Journal of Forecasting* **22**, 17–28.

Taylor, J. W. & Buizza, R. (2003), 'Using weather ensemble predictions in electricity demand forecasting', *International Journal of Forecasting* **19**, 57–70.

Taylor, J. W., De Menezes, L. M. & McSharry, P. E. (2006), 'A comparison of univariate methods for forecasting electricity demand up to a day ahead', *International Journal of Forecasting* **22**, 1–16.

Taylor, J. W. & McSharry, P. E. (2007), 'Short-term load forecasting methods: An evaluation based on european data', *IEEE Transactions on Power Systems* **22**, 2213 – 2219.

Young, P. C., Pedregal, D. J. & Tych, W. (1999), 'Dynamic harmonic regression', *Journal of Forecasting* **18**, 369–394.

Table 2
Estimation results for model (10) for the in-sample period September 1, 1995 until August 31, 2003, with estimated standard deviations (st.dev.) of irregular, of disturbances driving the stochastic processes of trends (level and slope) in (2) and of time-varying regression coefficients in (5) and (9). The q-ratio is the standard deviation divided by the one of the irregular. The estimates are presented for both 9 AM and 12 AM. The estimated correlations (corr.) are reported in the last column.

		9 AM ($i = 1$)		12 AM ($i = 2$)		
Cmp / Expl	par	st.dev.	q-ratio	st.dev.	q-ratio	corr.
irregular	$\varepsilon_{t,i}$	185.4	1.000	283.2	1.000	–
level	$f_{t,i}$	52.7	0.284	52.4	0.185	1
slope	$g_{t,i}$	13.0	0.070	13.0	0.046	1
$X^1_{t,i}$ Heating	$\beta^1_{t,i}$	25.7	0.138	19.6	0.069	0.99
$X^2_{t,i}$ SmoHeating	$\beta^2_{t,i}$	95.3	0.513	68.2	0.241	1
$X^3_{t,i}$ SmoCooling	$\beta^3_{t,i}$	4.3	0.023	4.4	0.016	1
$W^1_t = a^{WD}_{1,t}$	$\gamma^1_{t,i}$	471.9	2.545	468.8	1.655	1
$W^2_t = b^{WD}_{1,t}$	$\gamma^2_{t,i}$	210.4	1.135	213.2	0.753	1
$W^3_t = a^{WE}_{1,t}$	$\gamma^3_{t,i}$	62.4	0.337	57.1	0.202	1
$W^4_t = b^{WE}_{1,t}$	$\gamma^4_{t,i}$	114.7	0.618	118.5	0.418	1
$W^5_t = a^{WD}_{2,t}$	$\gamma^5_{t,i}$	105.1	0.567	117.4	0.414	1
$W^6_t = b^{WD}_{2,t}$	$\gamma^6_{t,i}$	105.9	0.571	103.6	0.366	1
$W^7_t = a^{WE}_{2,t}$	$\gamma^7_{t,i}$	33.9	0.183	26.8	0.095	0.97
$W^8_t = b^{WE}_{2,t}$	$\gamma^8_{t,i}$	50.0	0.270	43.2	0.153	1
$W^9_t = a^{WD}_{3,t}$	$\gamma^9_{t,i}$	392.9	2.119	381.5	1.347	1
$W^{10}_t = b^{WD}_{3,t}$	$\gamma^{10}_{t,i}$	6.8	0.037	9.4	0.033	1
$W^{11}_t = a^{WE}_{3,t}$	$\gamma^{11}_{t,i}$	32.2	0.177	33.9	0.120	1
$W^{12}_t = b^{WE}_{3,t}$	$\gamma^{12}_{t,i}$	36.1	0.195	40.1	0.142	1
$W^{13}_t = a^{WD}_{4,t}$	$\gamma^{13}_{t,i}$	375.6	2.026	294.4	1.040	1
$W^{14}_t = b^{WD}_{4,t}$	$\gamma^{14}_{t,i}$	85.8	0.463	85.1	0.301	1
$W^{15}_t = a^{WE}_{4,t}$	$\gamma^{15}_{t,i}$	23.1	0.125	22.2	0.078	1
$W^{16}_t = b^{WE}_{4,t}$	$\gamma^{16}_{t,i}$	26.7	0.144	24.8	0.088	1
W^{17}_t Monday	$\gamma^{17}_{t,i}$	14.4	0.078	0.6	0.002	1
W^{18}_t Friday	$\gamma^{18}_{t,i}$	0.7	0.004	4.3	0.015	1
W^{19}_t Saturday	$\gamma^{19}_{t,i}$	59.2	0.319	7.1	0.025	-0.70
W^{20}_t Sunday	$\gamma^{20}_{t,i}$	133.6	0.721	107.1	0.378	-0.30
W^{21}_t Holiday	$\gamma^{21}_{t,i}$	559.5	3.018	465.8	1.645	1
W^{22}_t Bridge day	$\gamma^{22}_{t,i}$	126.0	0.679	90.7	0.320	1
W^{26}_t August Tr1	$\gamma^{26}_{t,i}$	6.0	0.032	5.0	0.018	1
W^{27}_t August Tr2	$\gamma^{27}_{t,i}$	110.7	0.597	89.6	0.317	1

Table 3
Estimation results for the lagged temperature coefficients of model (10) in the time-varying regression equations (5) for the heating and smoothed heating effects (leading to nonlinear and yearly periodic dependence of temperature on electricity loads) and estimation results for the fixed regression coefficients of model (10) in equations (5) and (9) related to the cloud cover effect and to four calendar effects. The estimates are presented for both 9 AM and 12 AM.

Explanatory variable	hour	coefficient	estimate	stand.err.	t-value
$X^1_{t,i}$ Heating	9 AM	λ^1_1	6.84	1.02	6.69
	12 AM	λ^1_2	8.55	1.22	7.02
$X^2_{t,i}$ Smoothed-heating	9 AM	λ^2_1	16.96	4.42	3.84
	12 AM	λ^2_2	16.58	3.44	4.82
$X^4_{t,i}$ Cloud cover	9 AM	β^4_1	147	7.9	18.6
	12 AM	β^4_2	171	7.8	21.9
W^{23}_t December 25th	9 AM	γ^{23}_1	-14028	326.1	43.0
	12 AM	γ^{23}_2	-10602	278.3	38.1
W^{24}_t January 1st	9 AM	γ^{24}_1	-15629	321.8	48.6
	12 AM	γ^{24}_2	-10847	268.7	40.4
W^{25}_t December 24th	9 AM	γ^{25}_1	-4486	336.8	13.3
	12 AM	γ^{25}_2	-4499	289.1	15.5
W^{28}_t Daylight saving	9 AM	γ^{28}_1	165	108.1	1.5
	12 AM	γ^{28}_2	108	104.6	1.0

Table 4
One-day ahead forecasting results by month and by hour for the post-sample period September 1, 2003 until August 31, 2004, using real (left) and forecast temperature (right). The MAPE in (12) is reported for five different models: weekly random walk (RW); regression model (Reg) , i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = \sigma_{u^j,i} = \sigma_{e^j,i} = 0$ in (2),(5) and (9); univariate model (Univ), i.e. model (10) with $k = 1$; time-varying regression model (TVR), i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = 0$ in (2); local linear trend plus time-varying regression model (TTR) given in (10) with $k = 2$. The RMSE in (12) is reported for temperature forecast errors (^oC). The RW model does not necessarily produce forecasts for all N observations. Overall results for all 343 non-EJP days are in Table 5.

				MAPE (real temperature)				RMSE	MAPE (forecast temperature)			
Hour	Month	N	RW	Reg	Univ	TVR	TTR	^oC	Reg	Univ	TVR	TTR
9	January	25	5.84	3.03	1.97	1.67	**1.65**	0.58	3.07	1.93	1.58	**1.53**
	February	20	5.76	**0.98**	1.30	1.50	1.49	0.77	**0.92**	1.31	1.46	1.46
	March	27	10.79	2.35	1.16	1.17	**1.12**	1.16	2.21	**1.47**	1.71	1.69
	April	30	7.59	1.15	1.24	**1.13**	1.27	1.00	**1.37**	1.52	1.40	1.47
	May	31	5.38	3.24	2.33	**1.86**	2.02	0.75	3.35	2.20	**1.84**	2.03
	June	30	1.14	1.79	**0.60**	0.64	0.63	0.37	1.81	**0.61**	0.64	0.62
	July	31	1.16	2.22	**1.02**	1.12	1.06	0.61	2.24	**1.03**	1.14	1.08
	August	31	5.36	7.07	2.22	1.99	**1.97**	0.56	7.07	2.21	1.99	**1.97**
	September	30	1.78	2.26	**0.66**	0.72	0.72	1.05	2.31	**0.69**	0.74	0.75
	October	31	6.40	2.00	1.26	1.02	**1.00**	1.06	1.92	1.24	1.02	**1.01**
	November	30	5.24	1.65	1.30	1.45	**1.27**	1.00	1.68	**1.48**	1.72	1.52
	December	27	5.70	3.64	**1.78**	1.98	2.04	0.76	3.67	**1.97**	2.23	2.29
12	January	25	5.57	2.75	1.83	1.77	**1.73**	1.19	2.36	1.85	**1.68**	1.74
	February	20	6.82	**1.06**	1.16	1.26	1.31	1.82	**2.02**	2.23	2.08	2.10
	March	27	10.05	1.67	**1.02**	1.17	1.16	1.55	**1.51**	1.60	1.76	1.71
	April	30	6.80	**1.12**	1.41	1.35	1.48	0.67	1.38	**1.38**	1.52	1.55
	May	31	4.50	2.75	**1.72**	1.76	1.94	0.70	2.70	**1.75**	1.82	1.98
	June	30	1.26	2.04	0.66	**0.61**	0.64	0.61	2.04	0.66	**0.61**	0.64
	July	31	1.15	1.89	**1.05**	1.19	1.12	0.80	1.89	**1.05**	1.19	1.12
	August	31	4.71	5.85	**1.53**	1.62	1.60	0.82	5.85	**1.53**	1.62	1.60
	September	30	0.97	2.24	**0.69**	0.75	0.72	0.94	2.24	**0.69**	0.75	0.73
	October	31	5.97	2.06	**1.03**	1.09	1.08	0.97	2.03	**1.13**	1.17	1.16
	November	30	5.18	1.20	1.27	1.34	**1.19**	1.45	**1.16**	1.87	2.01	1.90
	December	27	4.80	2.69	**1.71**	1.77	1.82	1.33	2.65	**1.96**	2.05	2.12

Table 5
One-day ahead forecasting results by day type and by hour for the post-sample period September 1, 2003 until August 31, 2004, using real (left) and forecast temperature (right). The MAPE in (12) is reported for five different models: weekly random walk (RW); regression model (Reg) , i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = \sigma_{u^j,i} = \sigma_{e^j,i} = 0$ in (2),(5) and (9); univariate model (Univ), i.e. model (10) with $k = 1$; time-varying regression model (TVR), i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = 0$ in (2); local linear trend plus time-varying regression model (TTR) given in (10) with $k = 2$. Total MAPE for all 343 non-EJP days is also reported as well as the MPE in (12) for EJP days. The RW model does not necessarily produce forecasts for all N observations.

				MAPE (real temperature)				MAPE (forecast temperature)			
Hour	Day type	N	RW	Reg	Univ	TVR	TTR	Reg	Univ	TVR	TTR
9	Default	132	4.63	2.23	0.93	**0.92**	0.92	2.26	**1.02**	1.06	1.05
	Monday	47	4.88	2.66	1.39	1.39	**1.36**	2.73	1.45	1.47	**1.44**
	Friday	48	4.50	1.93	**1.01**	1.02	1.01	1.86	**1.06**	1.09	1.10
	Saturday	49	5.41	2.38	1.42	**1.22**	1.23	2.38	1.35	**1.22**	1.27
	Sunday	51	6.39	2.73	1.69	1.59	**1.55**	2.77	1.87	1.71	**1.68**
	Special	16	–	9.03	5.49	**5.33**	5.40	9.16	5.47	5.48	**5.44**
	Total	343	5.04	2.66	1.40	1.34	**1.34**	2.68	1.47	1.45	**1.44**
12	Default	132	4.51	1.90	**0.94**	0.96	0.97	1.91	**1.13**	1.18	1.19
	Monday	47	4.27	2.15	**1.35**	1.40	1.42	2.16	**1.53**	1.53	1.53
	Friday	48	4.38	1.92	1.17	1.16	**1.12**	1.89	1.43	1.33	**1.30**
	Saturday	49	4.80	2.08	1.09	1.01	**1.01**	2.12	1.18	1.14	**1.13**
	Sunday	51	5.50	2.81	**1.26**	1.44	1.45	2.81	**1.54**	1.69	1.70
	Special	16	–	6.61	**4.23**	4.68	4.82	6.95	**4.33**	4.97	5.03
	Total	343	4.66	2.34	**1.25**	1.30	1.31	2.34	**1.44**	1.50	1.50
				MPE(real temperature)				MPE (forecast temperature)			
Hour	Day type	N	RW	Reg	Univ	TVR	TTR	Reg	Univ	TVR	TTR
9	EJP	23	–	-1.52	-4.25	-4.87	-4.77	-1.71	-4.52	-5.21	-5.13
12	EJP	23	–	-2.36	-4.55	-4.64	-4.55	-3.66	-6.36	-6.15	-6.09

Table 6
Forecasting results for different forecast horizons, for the post-sample period September 1, 2003 until August 31, 2004 (total non-EJP days is 343, N decreases with forecasting horizon). The MAPE in (12) is reported for five different models: weekly random walk (RW); regression model (Reg) , i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = \sigma_{u^j,i} = \sigma_{e^j,i} = 0$ in (2),(5) and (9); univariate model (Univ), i.e. model (10) with $k = 1$; time-varying regression model (TVR), i.e. model (10) with $k = 2$ and $\sigma_{v,i} = \sigma_{w,i} = 0$ in (2); local linear trend plus time-varying regression model (TTR) given in (10) with $k = 2$. The RW model does not necessarily produce forecasts for all N.

			MAPE (Real temperature)				
Hour	Horizon	N	RW	Reg	Univ	TVR	TTR
9	1 day	343	5.04	2.66	1.40	1.34	**1.34**
	2 days	342	5.04	2.68	1.92	**1.84**	1.85
	3 days	341	5.04	2.70	2.26	**2.19**	2.21
	4 days	340	5.04	2.72	2.47	**2.36**	2.39
	5 days	339	5.04	2.74	2.71	**2.57**	2.57
	6 days	338	5.04	**2.76**	2.87	2.83	2.85
	7 days	337	5.04	**2.77**	3.07	3.06	3.09
12	1 day	343	4.66	2.34	**1.25**	1.30	1.31
	2 days	342	4.66	2.32	**1.58**	1.70	1.72
	3 days	341	4.66	2.36	**1.87**	2.00	2.02
	4 days	340	4.66	2.38	**2.01**	2.11	2.14
	5 days	339	4.66	2.40	**2.10**	2.24	2.25
	6 days	338	4.66	2.41	**2.30**	2.45	2.49
	7 days	337	4.66	**2.42**	2.49	2.67	2.66

Table 7
One-day ahead forecasting results by day type and by hour for the post-sample period September 1, 2003 until August 31, 2004, using forecast temperature. Total non-special and non-EJP days is 327 (326 at 3AM). The RMSE and the MAPE in (12) are reported for three different models: Univariate Model (Univ), i.e. model (10) with $k = 1$; TTR model given in (10) with $k = 2$, for consecutive hours (i-1,i) (TTR-1), TTR model given in (10) with $k = 2$, for consecutive hours (i,i+1) (TTR+1). Total RMSE and MAPE for the 24 hours is also reported (7847 observations).

		RMSE			MAPE		
Hour	N	Univ	TTR-1	TTR+1	Univ	TTR-1	TTR+1
0	327	966	**755**	993	1.43	**1.11**	1.42
1	327	986	984	**971**	**1.40**	1.41	1.40
2	327	992	961	**957**	1.54	1.50	**1.47**
3	326	946	929	**910**	1.52	1.52	**1.50**
4	327	922	894	**846**	1.50	1.49	**1.44**
5	327	859	858	**823**	1.44	1.44	**1.38**
6	327	1028	1067	**1027**	1.56	**1.54**	1.55
7	327	1178	1170	**1164**	1.64	1.61	**1.55**
8	327	983	1016	**979**	1.32	1.32	**1.29**
9	327	958	**954**	964	1.27	1.27	**1.25**
10	327	1097	1043	**1041**	1.44	1.33	**1.31**
11	327	1128	1091	**1070**	1.43	1.37	**1.34**
12	327	1041	1041	**1032**	1.30	1.30	**1.25**
13	327	**1050**	1080	1074	1.38	**1.38**	1.39
14	327	**1025**	1088	1062	**1.39**	1.45	1.43
15	327	**1046**	1108	1057	1.47	1.53	**1.47**
16	327	**1089**	1099	1132	1.61	**1.59**	1.63
17	327	**1139**	1158	1164	1.61	1.60	**1.59**
18	327	1137	1217	**1111**	1.55	1.64	**1.54**
19	327	1057	1065	**1050**	1.46	1.48	**1.42**
20	327	942	**922**	1003	1.34	**1.32**	1.45
21	327	**795**	867	814	1.20	1.26	**1.20**
22	327	**759**	785	759	1.07	1.12	**1.06**
23	327	722	728	**680**	1.06	1.07	**0.92**
Total	7847	993	995	**987**	1.41	1.40	**1.39**

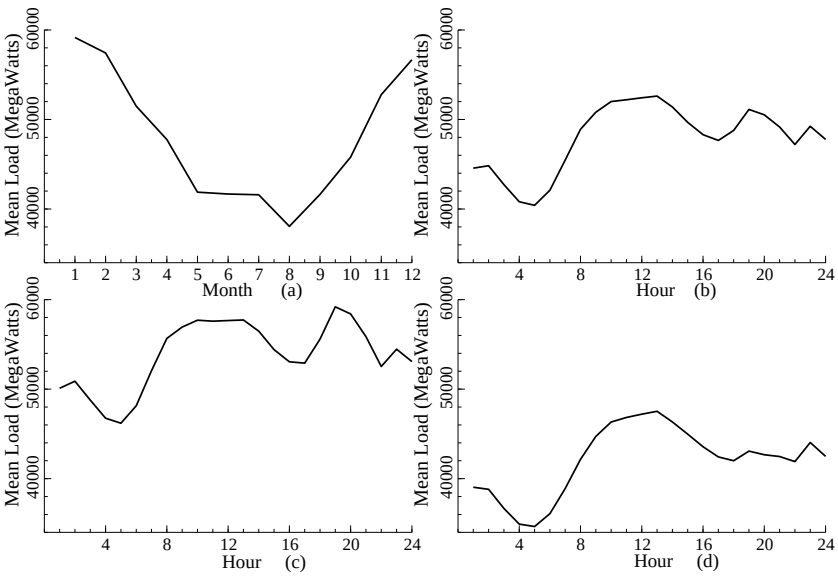

Fig. 1. Data description of French national electricity loads from September 1, 1995 until August 31, 2004: (a) monthly averages; (b) hourly averages; (c) hourly averages for winter months (October – March); (d) hourly averages for summer months (April – September).

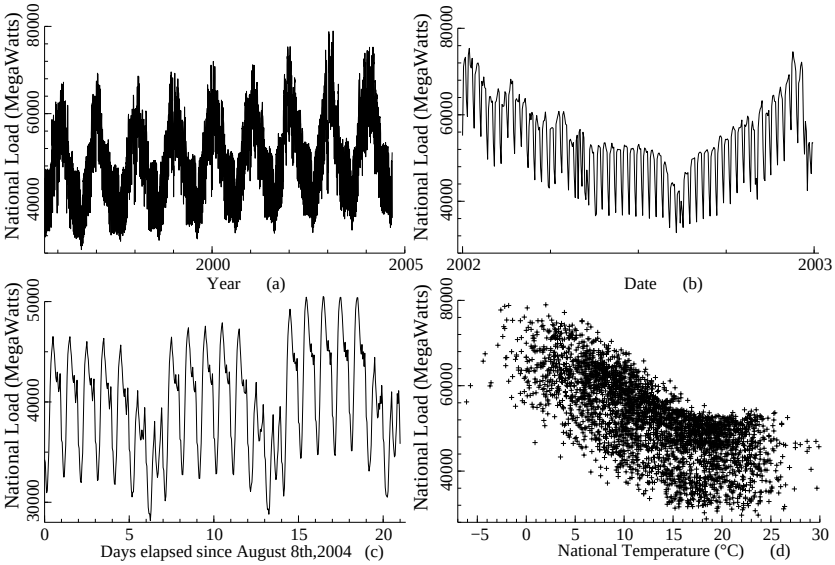

Fig. 2. Data description: (a) French national daily electricity loads from September 1, 1995 until August 31, 2004, at 9 AM; (b) Hourly electricity loads in 2002 (including special tariff EJP days); (c) Hourly electricity loads in three weeks after August 8, 2004; (d) Daily electricity load versus national average temperature from September 1, 1995 until August 31, 2004.

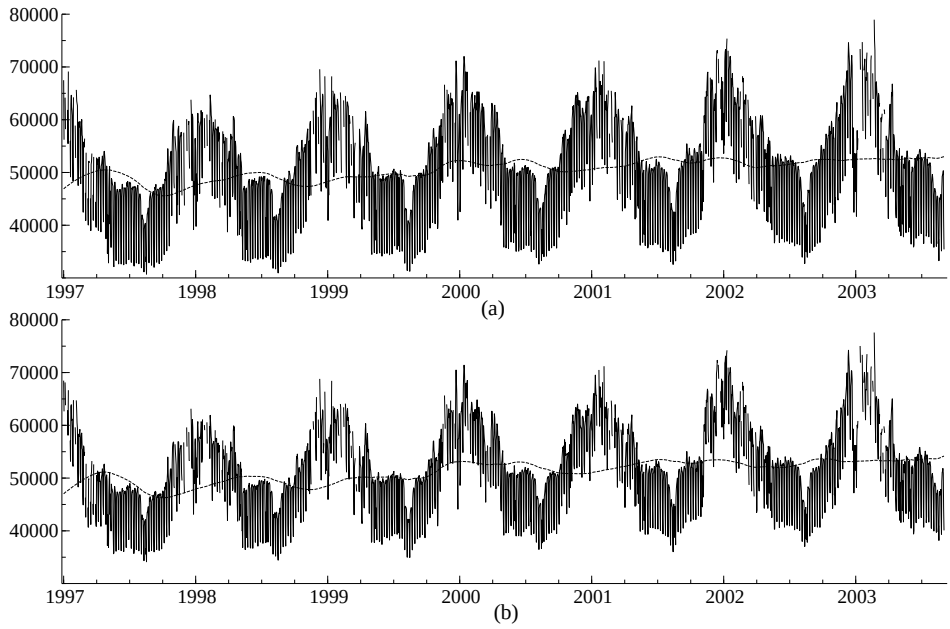

Fig. 3. French national hourly electricity load in MWh from January 1, 1997 until August 31, 2003 and smoothed estimates of the stochastic trend in model (10): (a) at 9 AM, stochastic trend $f_{t,1}$; (b) at 12 AM, stochastic trend $f_{t,2}$.

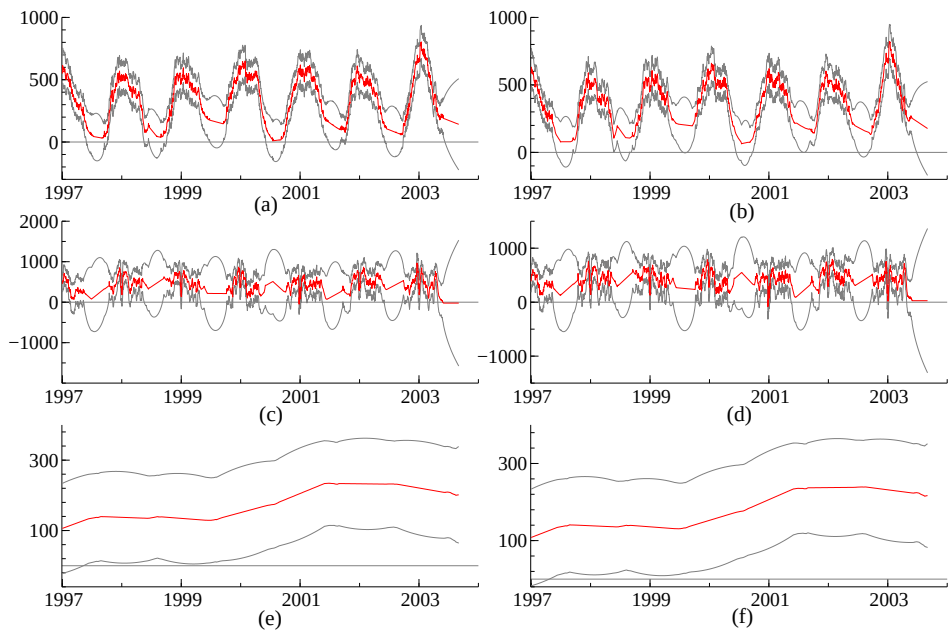

Fig. 4. Stochastic temperature coefficients for 9 AM and 12 AM: smoothed estimates and associated 95% confidence intervals. Coefficient for heating degrees (a) at 9 AM ($\beta^1_{t,1}$) and (b) at 12 AM ($\beta^1_{t,2}$) ; Coefficient for smoothed-heating degrees (c) at 9 AM ($\beta^2_{t,1}$) and (d) at 12 AM ($\beta^2_{t,2}$) ; Coefficient for smoothed-cooling degrees (e) at 9 AM ($\beta^3_{t,1}$) and (f) at 12 AM ($\beta^3_{t,2}$) .

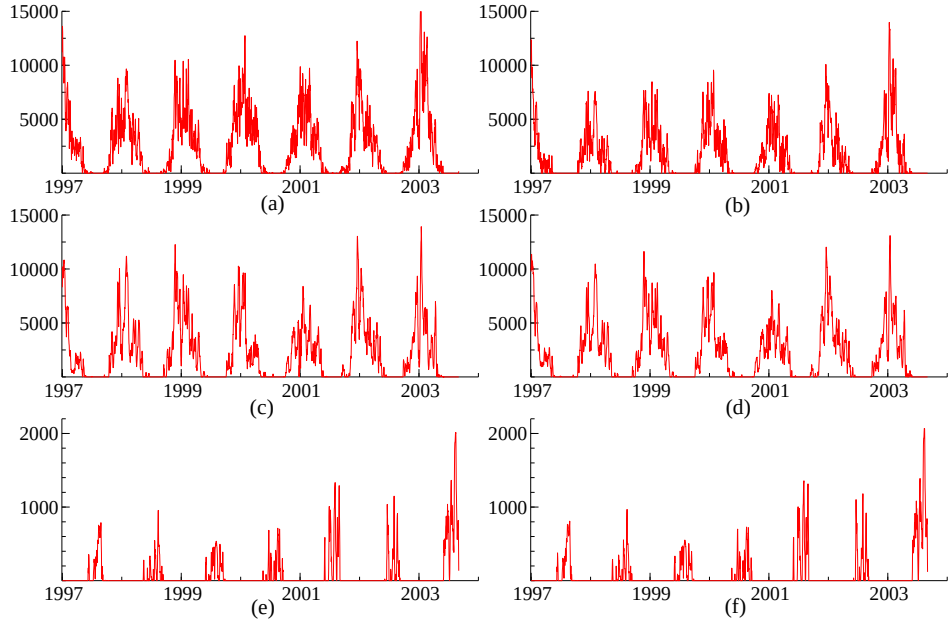

Fig. 5. Stochastic regression effect for heating degrees (a) at 9 AM ($\beta^1_{t,1}X^1_{t,1}$) and (b) at 12 AM ($\beta^1_{t,2}X^1_{t,2}$) ; Stochastic regression effect for smoothed-heating degrees (c) at 9 AM ($\beta^2_{t,1}X^2_{t,1}$) and (d) at 12 AM ($\beta^2_{t,2}X^2_{t,2}$) ; Stochastic regression effect for smoothed-cooling degrees (e) at 9 AM ($\beta^3_{t,1}X^3_{t,1}$) and (f) at 12 AM ($\beta^3_{t,2}X^3_{t,2}$).

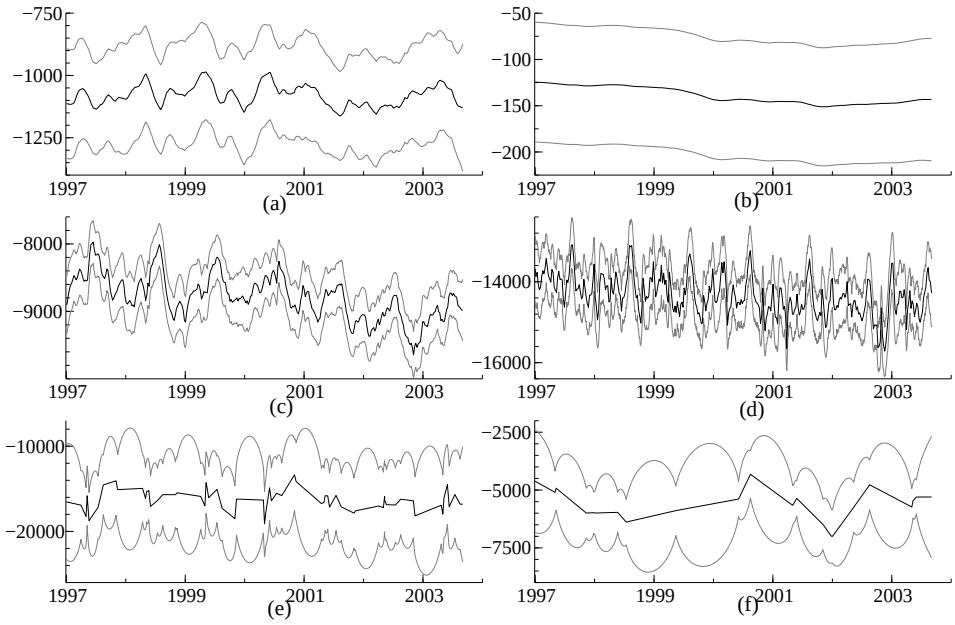

Fig. 6. Time-varying day-type coefficients for 9 AM: smoothed estimates and associated 95% confidence intervals. (a) Coefficient for Mondays, $\gamma^{17}_{t,1}$; (b) Coefficient for Fridays, $\gamma^{18}_{t,1}$; (c) Coefficient for Saturdays, $\gamma^{19}_{t,1}$; (d) Coefficient for Sundays, $\gamma^{20}_{t,1}$; (e) Coefficient for Holidays, $\gamma^{21}_{t,1}$; (f) Coefficient for bridge days, $\gamma^{25}_{t,1}$.

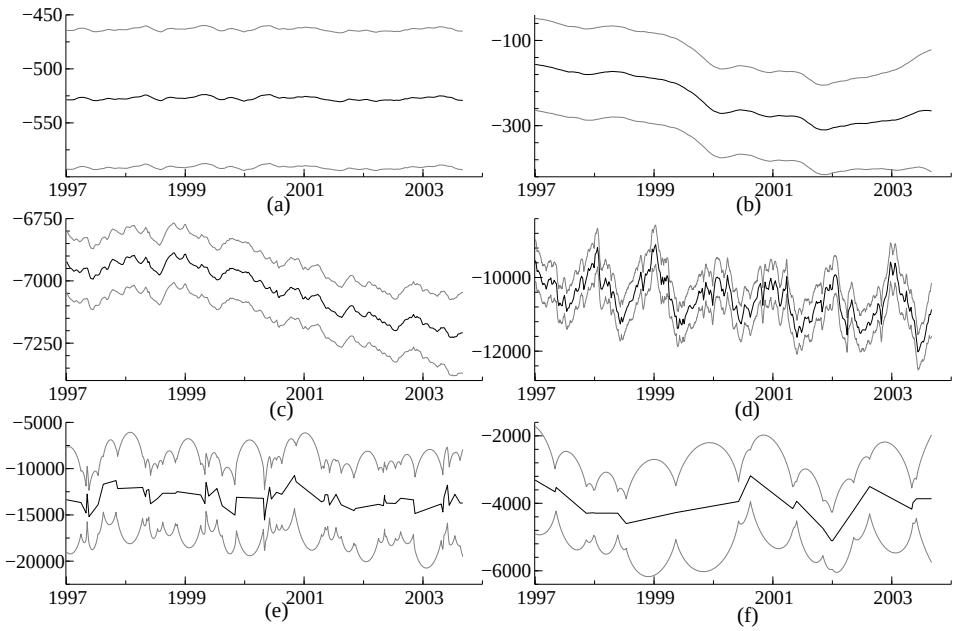

Fig. 7. Time-varying day-type coefficients for 12 AM: smoothed estimates and associated 95% confidence intervals. (a) Coefficient for Mondays, $\gamma_{t,2}^{17}$; (b) Coefficient for Fridays, $\gamma_{t,2}^{18}$; (c) Coefficient for Saturdays, $\gamma_{t,2}^{19}$; (d) Coefficient for Sundays, $\gamma_{t,2}^{20}$; (e) Coefficient for Holidays, $\gamma_{t,2}^{21}$; (f) Coefficient for bridge days, $\gamma_{t,2}^{25}$.

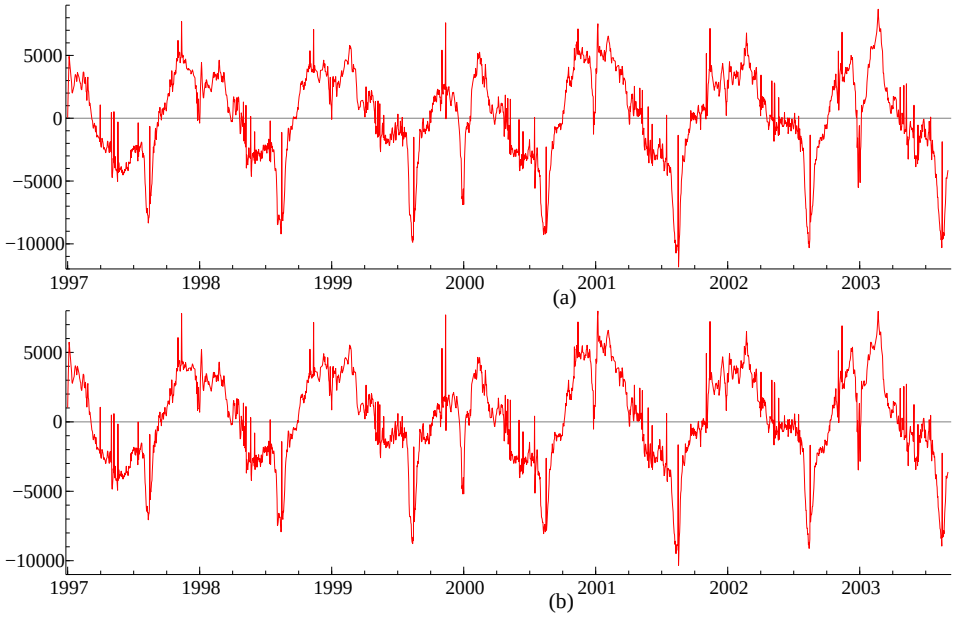

Fig. 8. Stochastic yearly patterns in the hourly loads. Smoothed estimates of the total effect of yearly Fourier series for weekdays and weekends plus the effect of daylight saving and August trends for weekends. (a) Yearly pattern at 9 AM, $\gamma_1^{28} W_t^{28} + \sum_{j \in \{1,..,16,26,27\}} \gamma_{t,1}^j W_t^j$; (b) Yearly pattern at 12 AM, $\gamma_2^{28} W_t^{28} + \sum_{j \in \{1,..,16,26,27\}} \gamma_{t,2}^j W_t^j$;

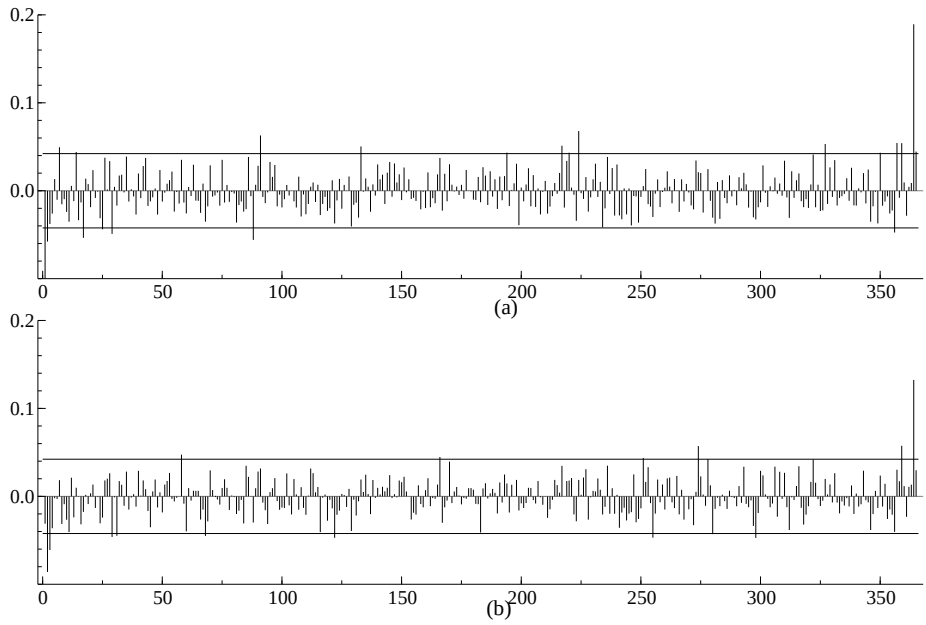

Fig. 9. Empirical ACFs of in-sample scaled one-day ahead prediction errors (standardised) for lags $1, \ldots, 366$, at (a) 9 AM ($i = 1$) and at (b) 12 AM ($i = 2$).

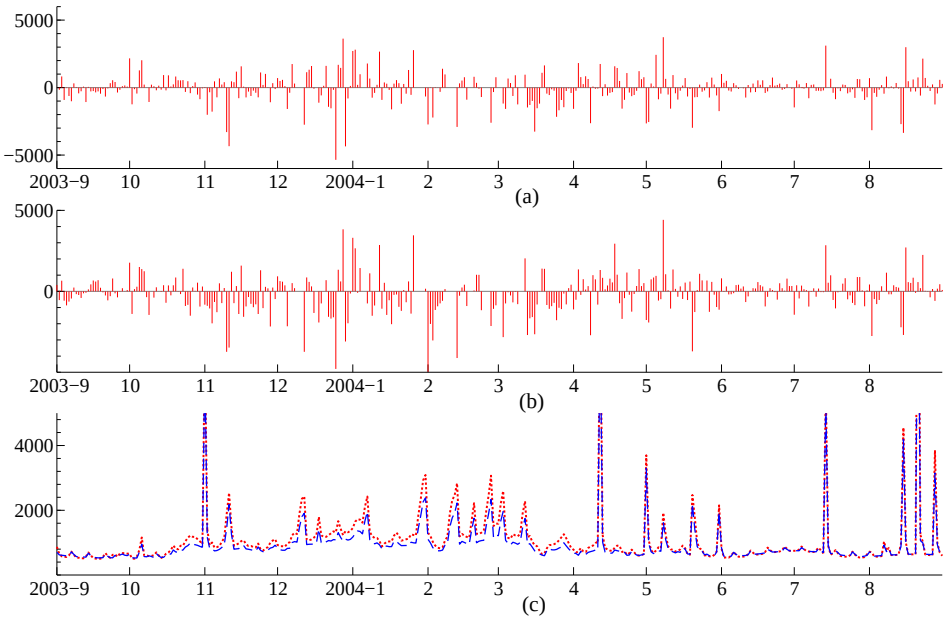

Fig. 10. Post-sample hourly one-day-ahead forecast errors $F^1_{t,i}$ based on weather forecasts at (a) 9 AM ($i = 1$) and at (b) 12 AM ($i = 2$), with (c) associated standard errors, for $t = n + 1, .., n + 366$, .

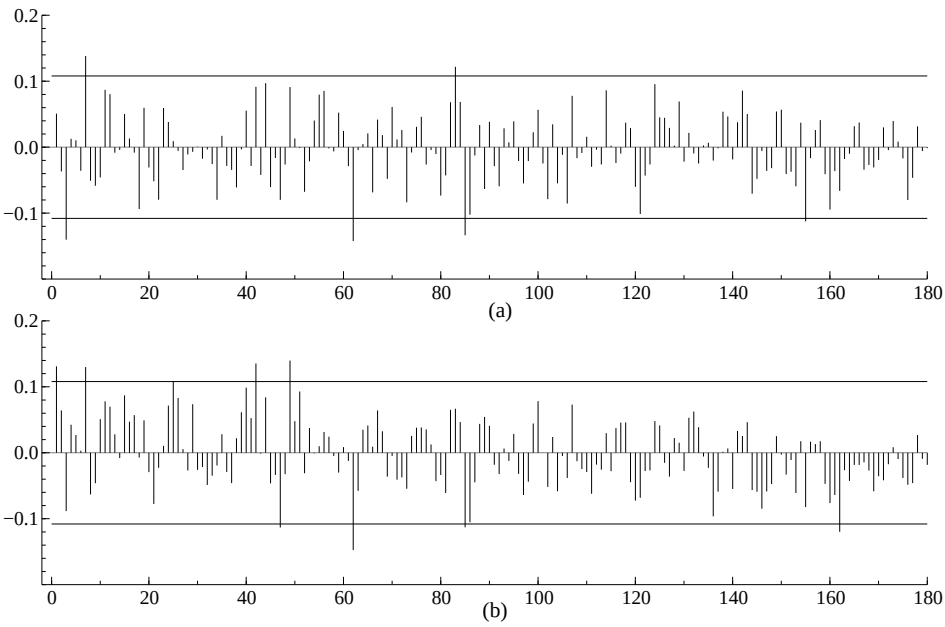

Fig. 11. Empirical autocorrelation function for post-sample one-day ahead forecast errors $y_{t,i} - F^1_{t,i}$ based on weather forecasts, standardised, for $t = n + 1, .., n + 366$, at (a) 9 AM ($i = 1$) and at (b) 12 AM ($i = 2$).