



MONASH University

Department of Econometrics and Business Statistics

<http://www.buseco.monash.edu.au/depts/ebs/pubs/wpapers/>

**Forecasting time series with
complex seasonal patterns
using exponential smoothing**

Alysha M De Livera, Rob J Hyndman
and Ralph D Snyder

October 2010

Working Paper 15/09

Forecasting time series with complex seasonal patterns using exponential smoothing

Alysha M De Livera

Department of Econometrics and Business Statistics,
Monash University, VIC 3800
Australia.

Email: alysha.delivera@monash.edu

Division of Mathematics, Informatics and Statistics,
Commonwealth Scientific and Industrial Research Organisation,
Clayton, VIC 3168
Australia.

Email: alysha.delivera@csiro.au

Rob J Hyndman

Department of Econometrics and Business Statistics,
Monash University, VIC 3800
Australia.

Email: rob.hyndman@monash.edu

Ralph D Snyder

Department of Econometrics and Business Statistics,
Monash University, VIC 3800
Australia.

Email: ralph.snyder@monash.edu

28 October 2010

Forecasting time series with complex seasonal patterns using exponential smoothing

Abstract

An innovations state space modeling framework is introduced for forecasting complex seasonal time series such as those with multiple seasonal periods, high frequency seasonality, non-integer seasonality and dual-calendar effects. The new framework incorporates Box-Cox transformations, Fourier representations with time varying coefficients, and ARMA error correction. Likelihood evaluation and analytical expressions for point forecasts and interval predictions under the assumption of Gaussian errors are derived, leading to a simple, comprehensive approach to forecasting complex seasonal time series. A key feature of the framework is that it relies on a new method that greatly reduces the computational burden in the maximum likelihood estimation. The modeling framework is useful for a broad range of applications, its versatility being illustrated in three empirical studies. In addition, the proposed trigonometric formulation is presented as a means of decomposing complex seasonal time series, and it is shown that this decomposition leads to the identification and extraction of seasonal components which are otherwise not apparent in the time series plot itself.

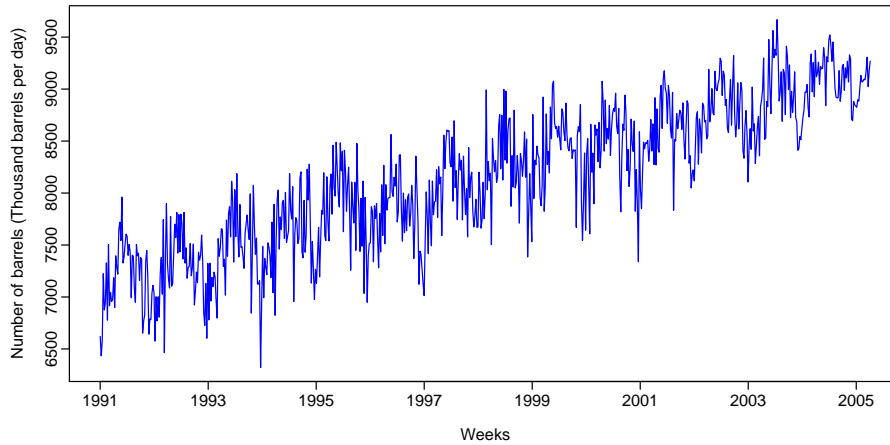
Keywords: exponential smoothing, Fourier series, prediction intervals, seasonality, state space models, time series decomposition.

1 Introduction

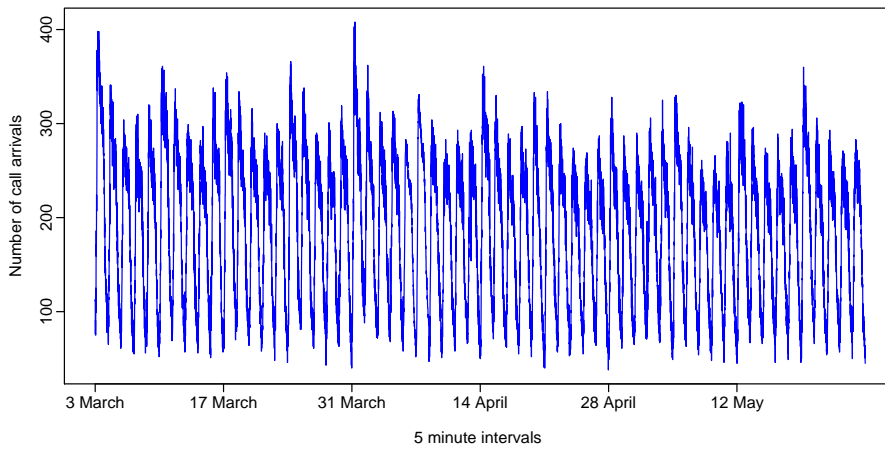
Many time series exhibit complex seasonal patterns. Some, most commonly weekly series, have patterns with a non-integer period. Weekly United States finished motor gasoline products in thousands of barrels per day, as shown in Figure 1(a), has an annual seasonal pattern with period $365.25/7 \approx 52.179$.

Other series have high frequency multiple seasonal patterns. The number of retail banking call arrivals per 5-minute interval between 7:00am and 9:05pm each weekday, as depicted in Figure 1(b), has a daily seasonal pattern with period 169 and a weekly seasonal pattern with period $169 \times 5 = 845$. A longer version of this series might also exhibit an annual seasonal pattern. Further examples where such multiple seasonal patterns can occur include daily hospital admissions, requests for cash at ATMs, electricity and water usage, and access to computer web sites.

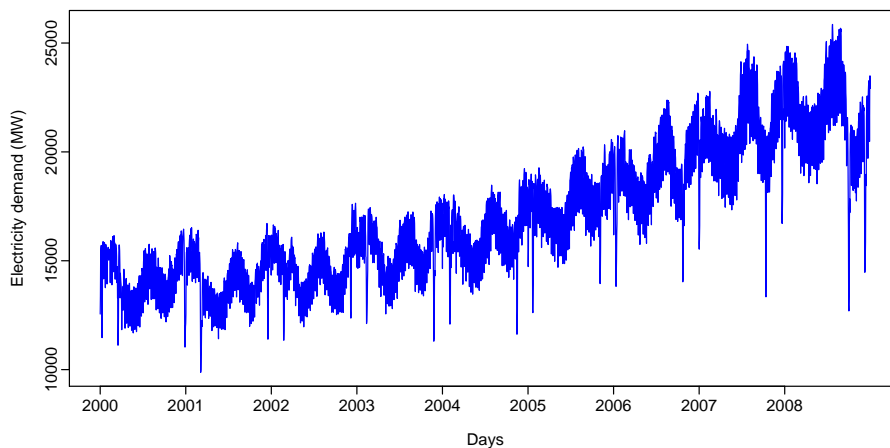
Yet other series may have dual-calendar seasonal effects. Daily electricity demand in Turkey over nine years from 1 January 2000 to 31 December 2008, shown in Figure 1(c), has a weekly seasonal pattern and two annual seasonal patterns: one for the Hijri calendar with a period of 354.37; and the other for the Gregorian calendar with a period of 365.25. The Islamic Hijri calendar is based on lunar cycles and is used for religious activities and related holidays. It is approximately 11 days shorter than the Gregorian calendar. The Jewish, Hindu and Chinese calendars create similar effects that can be observed in time series affected by cultural and social events (e.g., electricity demand, water usage, and other related consumption data), and need to be accounted for in forecasting studies (Lin & Liu 2002, Riazuddin & Khan 2005). Unlike the multiple periods seen with hourly and daily data, these dual calendar effects involve non-nested seasonal periods.



(a) US finished motor gasoline products supplied (thousands of barrels per day), from February 1991 to July 2005.



(b) Number of call arrivals handled on weekdays between 7am and 9:05pm from March 3, 2003, to May 23, 2003 in a large North American commercial bank.



(c) Turkish electricity demand data from January 1, 2000, to December 31, 2008.

Figure 1: Examples of complex seasonality showing (a) non-integer seasonal periods (b) multiple nested seasonal periods, and (c) multiple non-nested and non-integer seasonal periods.

27 Most existing time series models are designed to accommodate simple seasonal patterns
28 with a small integer-valued period (such as 12 for monthly data or 4 for quarterly data).
29 Important exceptions ([Harvey & Koopman 1993](#), [Harvey et al. 1997](#), [Pedregal & Young](#)
30 [2006](#), [Taylor 2003](#), [Gould et al. 2008](#), [Taylor & Snyder 2009](#), [Taylor 2010](#)) handle some but
31 not all of the above complexities. [Harvey et al. \(1997\)](#), for example, uses a trigonometric
32 approach for single seasonal time series within a traditional multiple source of error state
33 space framework. The single source of error approach adopted in this paper is similar in
34 some respects, but admits a larger effective parameter space with the possibility of better
35 forecasts (see [Hyndman et al. 2008](#), Chap 13), allows for multiple nested and non-nested
36 seasonal patterns, and handles potential nonlinearities. [Pedregal & Young \(2006\)](#) and [Harvey](#)
37 [& Koopman \(1993\)](#) have models for double seasonal time series, but they have not been
38 sufficiently developed for time series with more than two seasonal patterns, and are not
39 capable of accommodating the nonlinearity found in many time series in practice. Similarly,
40 in modeling complex seasonality, the existing exponential smoothing models (e.g., [Taylor](#)
41 [2003](#), [Gould et al. 2008](#), [Taylor & Snyder 2009](#), [Taylor 2010](#)) suffer from various weaknesses
42 such as over-parameterization, and the inability to accommodate both non-integer period
43 and dual-calendar effects. In contrast, we introduce a new innovations state space modeling
44 framework based on a trigonometric formulation which is capable of tackling all of these
45 seasonal complexities. Using the time series in [Figure 1](#), we demonstrate the versatility of
46 the proposed approach for forecasting and decomposition.

47 In [Section 2.1](#) we review the existing seasonal innovations state space models including
48 an examination of their weaknesses, particularly in relation to complex seasonal patterns.
49 We then introduce in [Sections 2.2](#) and [2.3](#) two generalizations designed to overcome some
50 or all of these these problems, one relying on trigonometric representations for handling
51 complex as well as the usual single seasonal patterns in a straightforward manner with fewer
52 parameters. [Section 3](#) contains a new method for the calculation of maximum likelihood

estimators, formulae for point and interval predictions, and the description of the model selection methodology. It will be seen that the proposed estimation procedure is sufficiently general to be applied to *any* innovations state space model while possessing some important advantages over an existing approach. The proposed models are then applied in Section 4 to the time series from Figure 1 where it will be seen that the trigonometric formulation leads to better forecasts and may be used for decomposition. Some conclusions are drawn in Section 5.

2 Exponential smoothing models for seasonal data

2.1 Traditional approaches

Single seasonal exponential smoothing methods, which are among the most widely used forecasting procedures in practice (Snyder et al. 2002, Makridakis et al. 1982, Makridakis & Hibon 2000), have been shown to be optimal for a class of innovations state space models (Ord et al. 1997, Hyndman et al. 2002). They are therefore best studied in terms of this framework because it then admits the possibility of likelihood calculation, the derivation of consistent prediction intervals and model selection based on information criteria. The single source of error (innovations) state space model is an alternative to its common multiple source of error analogue (Harvey 1989) but it is simpler, more robust, and has several other advantages (Hyndman et al. 2008) .

The most commonly employed seasonal models in the innovations state space framework include those underlying the well-known Holt-Winters' additive and multiplicative methods. Taylor (2003) extended the linear version of the Holt-Winters method to incorporate a second seasonal component as follows:

$$y_t = \ell_{t-1} + b_{t-1} + s_t^{(1)} + s_t^{(2)} + d_t \quad (1a)$$

$$\ell_t = \ell_{t-1} + b_{t-1} + \alpha d_t \quad (1b)$$

$$b_t = b_{t-1} + \beta d_t \quad (1c)$$

$$s_t^{(1)} = s_{t-m_1}^{(1)} + \gamma_1 d_t \quad (1d)$$

$$s_t^{(2)} = s_{t-m_2}^{(2)} + \gamma_2 d_t, \quad (1e)$$

71 where m_1 and m_2 are the periods of the seasonal cycles and d_t is a white noise random
 72 variable representing the prediction error (or disturbance). The components ℓ_t and b_t
 73 represent the level and trend components of the series at time t , respectively, and $s_t^{(i)}$
 74 represents the i th seasonal component at time t . The coefficients α, β, γ_1 and γ_2 are the
 75 so-called smoothing parameters, and $\ell_0, b_0, \{s_{1-m_1}^{(1)}, \dots, s_0^{(1)}\}$ and $\{s_{1-m_2}^{(2)}, \dots, s_0^{(2)}\}$ are the
 76 initial state variables (or “seeds”).

The parameters and seeds must be estimated, but this can be difficult when the number of seasonal components is large. This problem is partly addressed by noting that there is a redundancy when m_2 is an integer multiple of m_1 , something that seems to have previously gone unnoticed. Consider a time series $\{r_t\}$ consisting of repeated sequences of the constants c_1, \dots, c_{m_1} , one for each season in the smaller cycle. Then the seasonal equations can be written as

$$s_t^{(1)} + r_t = (s_{t-m_1}^{(1)} + r_t) + \gamma_1 d_t \quad (2a)$$

$$s_t^{(2)} - r_t = (s_{t-m_2}^{(2)} - r_t) + \gamma_2 d_t. \quad (2b)$$

77 When these are summed, the effect of r_t disappears. This suggests that the m_1 seed seasonal
 78 effects for the smaller seasonal cycle can be set to zero without constraining the problem
 79 in any way. Alternatively, each sub-season repeats itself m_2/m_1 times within the longer
 80 seasonal pattern. We can impose the constraint that the seed seasonal effects associated with
 81 each sub-season must sum to zero. For example, the period 10:00-11:00 AM repeats itself 7
 82 times in a week. We can insist that the 7 seasonal effects associated with this particular hour

83 sum to zero and that this is repeated for each of the 24 hour periods in a day. Analogues of
84 these restrictions can be developed when there are three or more seasonal patterns.

85 Despite this correction, a large number of initial seasonal values remain to be estimated
86 when some of the seasonal patterns have large periods, and such a model is likely to be
87 over-parameterized. For double seasonal time series [Gould et al. \(2008\)](#) attempted to reduce
88 this problem by dividing the longer seasonal length into sub-seasonal cycles that have
89 similar patterns. However, their adaptation is relatively complex and can only be used for
90 double seasonal patterns where one seasonal length is a multiple of the other. To avoid
91 the potentially large optimization problem, the initial states are usually approximated with
92 various heuristics ([Taylor 2003](#), [Gould et al. 2008](#), [Taylor 2010](#)), a practice that does not
93 lead to optimized seed states. We will propose an alternative estimation method, one that
94 relies on the principle of least squares to obtain optimized seed states — see [Section 3](#).

95 A further problem is that none of the approaches based on (1) can be used to handle complex
96 seasonal patterns such as non-integer seasonality and calendar effects, or time series with
97 non-nested seasonal patterns. One of our proposed models will allow for all these features.

98 The non-linear versions of the state space models underpinning exponential smoothing,
99 although widely used, suffer from some important weaknesses. [Akram et al. \(2009\)](#) showed
100 that most non-linear seasonal versions can be unstable, having infinite forecast variances
101 beyond a certain forecasting horizon. For some of the multiplicative error models which
102 do not have this flaw, [Akram et al. \(2009\)](#) proved that sample paths will converge almost
103 surely to zero even when the error distribution is non-Gaussian. Furthermore, for non-linear
104 models, analytical results for the prediction distributions are not available.

105 The models used for exponential smoothing assume that the error process $\{d_t\}$ is serially
106 uncorrelated. However, this may not always be the case. In an empirical study, using the
107 Holt-Winters' method for multiplicative seasonality, [Chatfield \(1978\)](#) showed that the error

108 process there is correlated and can be described by an AR(1) process. Taylor (2003) in a
 109 study of electricity demand forecasting using a double-seasonal Holt-Winters' multiplicative
 110 method found a similar problem. Others such as Gardner (1985), Reid (1975), and Gilchrist
 111 (1976) have also mentioned this issue of correlated errors, and the possibility of improving
 112 forecast accuracy by explicitly modeling it. The source of this autocorrelation may be due
 113 to features of the series not explicitly allowed for in the specification of the states. Annual
 114 seasonal effects may impact on the call center data, for example, but the limited sample size
 115 means that it cannot be explicitly modeled.

116 2.2 Modified models

117 We now consider various modifications of the state space models for exponential smoothing
 118 to handle a wider variety of seasonal patterns, and to also deal with the problems raised
 119 above.

120 To avoid the problems with non-linear models that are noted above, we restrict attention
 121 to linear homoscedastic models but allow some types of non-linearity using Box-Cox
 122 transformations (Box & Cox 1964). This limits our approach to only positive time series,
 123 but most series of interest in practice are positive. The notation $y_t^{(\omega)}$ is used to represent
 124 Box-Cox transformed observations with the parameter ω , where y_t is the observation at time
 125 t .

We can extend model (1) to include a Box-Cox transformation, ARMA errors and T seasonal patterns as follows.

$$y_t^{(\omega)} = \begin{cases} \frac{y_t^\omega - 1}{\omega}; & \omega \neq 0 \\ \log y_t & \omega = 0 \end{cases} \quad (3a)$$

$$y_t^{(\omega)} = \ell_{t-1} + \phi b_{t-1} + \sum_{i=1}^T s_{t-m_i}^{(i)} + d_t \quad (3b)$$

$$\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha d_t \quad (3c)$$

$$b_t = (1 - \phi)b + \phi b_{t-1} + \beta d_t \quad (3d)$$

$$s_t^{(i)} = s_{t-m_i}^{(i)} + \gamma_i d_t \quad (3e)$$

$$d_t = \sum_{i=1}^p \varphi_i d_{t-i} + \sum_{i=1}^q \theta_i \varepsilon_{t-i} + \varepsilon_t, \quad (3f)$$

126 where m_1, \dots, m_T denote the seasonal periods, ℓ_t is the local level in period t , b is the
 127 long-run trend, b_t is the short-run trend in period t , $s_t^{(i)}$ represents the i th seasonal component
 128 at time t , d_t denotes an ARMA(p, q) process and ε_t is a Gaussian white noise process with
 129 zero mean and constant variance σ^2 . The smoothing parameters are given by α , β and γ_i
 130 for $i = 1, \dots, T$. We adopt the [Gardner & McKenzie \(1985\)](#) damped trend with damping
 131 parameter ϕ , but follow the suggestion in [Snyder \(2006\)](#) to supplement it with a long-run
 132 trend b . This change ensures that predictions of future values of the short-run trend b_t
 133 converge to the long-run trend b instead of zero. The damping factor is included in the level
 134 and measurement equations as well as the trend equation for consistency with [Gardner &](#)
 135 [McKenzie \(1985\)](#), but identical predictions are obtained (see [Snyder 2006](#)) if it is excluded
 136 from the level and measurement equations.

137 The identifier BATS is an acronym for key features of the model: **B**ox-Cox trans-
 138 form, **A**RMA errors, **T**rend, and **S**easonal components. It is supplemented with argu-
 139 ments $(\omega, \phi, p, q, m_1, m_2, \dots, m_T)$ to indicate the Box-Cox parameter, damping parame-
 140 ter, ARMA parameters (p and q), and the seasonal periods (m_1, \dots, m_T) . For example,
 141 BATS(1, 1, 0, 0, m_1) represents the underlying model for the well-known Holt-Winters' addi-
 142 tive single seasonal method. The double seasonal Holt-Winters' additive seasonal model
 143 described by [Taylor \(2003\)](#) is given by BATS(1, 1, 0, 0, m_1, m_2), and that with the residual
 144 AR(1) adjustment in the model of [Taylor \(2003, 2008\)](#) is given by BATS(1, 1, 1, 0, m_1, m_2).
 145 The Holt-Winters' additive triple seasonal model with AR(1) adjustment in [Taylor \(2010\)](#) is
 146 given by BATS(1, 1, 1, 0, m_1, m_2, m_3).

147 The BATS model is the most obvious generalization of the traditional seasonal innovations
 148 models to allow for multiple seasonal periods. However, it cannot accommodate non-integer
 149 seasonality, and it can have a very large number of states; the initial seasonal component
 150 alone contains m_T non-zero states. This becomes a huge number of values for seasonal
 151 patterns with high periods.

152 2.3 Trigonometric seasonal models

In the quest for a more flexible parsimonious approach, we introduce the following trigono-
 metric representation of seasonal components based on Fourier series (West & Harrison
 1997, Harvey 1989):

$$s_t^{(i)} = \sum_{j=1}^{k_i} s_{j,t}^{(i)} \quad (4a)$$

$$s_{j,t}^{(i)} = s_{j,t-1}^{(i)} \cos \lambda_j^{(i)} + s_{j,t-1}^{*(i)} \sin \lambda_j^{(i)} + \gamma_1^{(i)} d_t \quad (4b)$$

$$s_{j,t}^{*(i)} = -s_{j,t-1}^{(i)} \sin \lambda_j^{(i)} + s_{j,t-1}^{*(i)} \cos \lambda_j^{(i)} + \gamma_2^{(i)} d_t, \quad (4c)$$

153 where $\gamma_1^{(i)}$ and $\gamma_2^{(i)}$ are smoothing parameters and $\lambda_j^{(i)} = 2\pi j/m_i$. We describe the stochastic
 154 *level* of the i th seasonal component by $s_{j,t}^{(i)}$, and the stochastic *growth* in the *level* of the i th
 155 seasonal component that is needed to describe the change in the seasonal component over
 156 time by $s_{j,t}^{*(i)}$. The number of harmonics required for the i th seasonal component is denoted
 157 by k_i . The approach is equivalent to index seasonal approaches when $k_i = m_i/2$ for even
 158 values of m_i , and when $k_i = (m_i - 1)/2$ for odd values of m_i . It is anticipated that most
 159 seasonal components will require fewer harmonics, thus reducing the number of parameters
 160 to be estimated. A deterministic representation of the seasonal components can be obtained
 161 by setting the smoothing parameters equal to zero.

162 A new class of innovations state space models is obtained by replacing the seasonal
163 component $s_t^{(i)}$ in equation (3c) by the trigonometric seasonal formulation, and the mea-
164 surement equation by $y_t^{(\omega)} = \ell_{t-1} + \phi b_{t-1} + \sum_{i=1}^T s_{t-1}^{(i)} + d_t$. This class is designated by
165 TBATS, the initial **T** connoting “trigonometric”. To provide more details about their
166 structure, this identifier is supplemented with relevant arguments to give the designation
167 $\text{TBATS}(\omega, \phi, p, q, \{m_1, k_1\}, \{m_2, k_2\}, \dots, \{m_T, k_T\})$.

168 A TBATS model requires the estimation of $2(k_1 + k_2 + \dots + k_T)$ initial seasonal values, a
169 number which is likely to be much smaller than the number of seasonal seed parameters
170 in a BATS models. Because it relies on trigonometric functions, it can be used to model
171 non-integer seasonal frequencies. A TBATS model should be distinguished from two other
172 related (Proietti 2000) multiple source of error seasonal formulations presented by Hannan
173 et al. (1970) and Harvey (1989). Some of the key advantages of the TBATS modeling
174 framework are: (i) it admits a larger effective parameter space with the possibility of better
175 forecasts (Hyndman et al. 2008, Chap 13); (ii) it allows for the accommodation of nested
176 and non-nested multiple seasonal components; (iii) it handles typical nonlinear features that
177 are often seen in real time series; (iv) it allows for any autocorrelation in the residuals to be
178 taken into account; and (v) it involves a much simpler, yet efficient estimation procedure
179 (see Section 3).

180 2.4 Innovations state space formulations

The above models are special cases of the linear innovations state space model (Anderson & Moore 1979) adapted here to incorporate the Box-Cox transformation to handle nonlinearities. It then has the form:

$$y_t^{(\omega)} = \mathbf{w}' \mathbf{x}_{t-1} + \varepsilon_t \quad (5a)$$

$$\mathbf{x}_t = \mathbf{F} \mathbf{x}_{t-1} + \mathbf{g} \varepsilon_t, \quad (5b)$$

181 where w' is a row vector, g is a column vector, F is a matrix and x_t is the unobserved state
 182 vector at time t .

183 **TBATS model**

184 The state vector for the TBATS model with a non-stationary growth term can be defined
 185 as $x_t = (\ell_t, b_t, s_t^{(1)}, \dots, s_t^{(T)}, d_t, d_{t-1}, \dots, d_{t-p+1}, \varepsilon_t, \varepsilon_{t-1}, \dots, \varepsilon_{t-q+1})'$ where $s_t^{(i)}$ is the row
 186 vector $(s_{1,t}^{(i)}, s_{2,t}^{(i)}, \dots, s_{k_i,t}^{(i)}, s_{1,t}^{*(i)}, s_{2,t}^{*(i)}, \dots, s_{k_i,t}^{*(i)})$. Let $\mathbf{1}_r = (1, 1, \dots, 1)$ and $\mathbf{0}_r = (0, 0, \dots, 0)$
 187 be row vectors of length r ; let $\gamma_1^{(i)} = \gamma_1^{(i)} \mathbf{1}_{k_i}$, $\gamma_2^{(i)} = \gamma_2^{(i)} \mathbf{1}_{k_i}$, $\gamma^{(i)} = (\gamma_1^{(i)}, \gamma_2^{(i)})$, $\gamma =$
 188 $(\gamma^{(1)}, \dots, \gamma^{(T)})$, $\varphi = (\varphi_1, \varphi_2, \dots, \varphi_p)$ and $\theta = (\theta_1, \theta_2, \dots, \theta_p)$; let $O_{u,v}$ be a $u \times v$ matrix of
 189 zeros, let $I_{u,v}$ be a $u \times v$ rectangular diagonal matrix with element 1 on the diagonal, and let
 190 $a^{(i)} = (\mathbf{1}_{k_i}, \mathbf{0}_{k_i})$ and $a = (a^{(1)}, \dots, a^{(T)})$. We shall also need the matrices $B = \gamma' \varphi$, $C = \gamma' \theta$,

$$A_i = \begin{bmatrix} C^{(i)} & S^{(i)} \\ -S^{(i)} & C^{(i)} \end{bmatrix}, \quad \tilde{A}_i = \begin{bmatrix} \mathbf{0}_{m_i-1} & 1 \\ I_{m_i-1} & \mathbf{0}'_{m_i-1} \end{bmatrix},$$

191 and $A = \bigoplus_{i=1}^T A_i$, where $C^{(i)}$ and $S^{(i)}$ are $k_i \times k_i$ diagonal matrices with elements $\cos(\lambda_j^{(i)})$
 192 and $\sin(\lambda_j^{(i)})$, respectively, for $j = 1, 2, \dots, k_i$ and $i = 1, \dots, T$, and where \bigoplus denotes the
 193 direct sum of the matrices. Let $\tau = 2 \sum_{i=1}^T k_i$.

194 Then the matrices for the TBATS model can be written as $w = (1, \phi, \mathbf{a}, \varphi, \theta)'$, $g =$
 195 $(\alpha, \beta, \gamma, 1, \mathbf{0}_{p-1}, 1, \mathbf{0}_{q-1})'$, and

$$F = \begin{bmatrix} 1 & \phi & \mathbf{0}_\tau & \alpha\varphi & \alpha\theta \\ 0 & \phi & \mathbf{0}_\tau & \beta\varphi & \beta\theta \\ \mathbf{0}'_\tau & \mathbf{0}'_\tau & A & B & C \\ 0 & 0 & \mathbf{0}_\tau & \varphi & \theta \\ \mathbf{0}'_{p-1} & \mathbf{0}'_{p-1} & O_{p-1,\tau} & I_{p-1,p} & O_{p-1,q} \\ 0 & 0 & \mathbf{0}_\tau & \mathbf{0}_p & \mathbf{0}_q \\ \mathbf{0}'_{q-1} & \mathbf{0}'_{q-1} & O_{q-1,\tau} & O_{q-1,p} & I_{q-1,q} \end{bmatrix}.$$

196 These matrices apply when all of the components are present in the model. When a
 197 component is omitted, the corresponding terms in the matrices must be omitted.

198 BATS model

199 The state space form of the BATS model can be obtained by letting $s_t^{(i)} =$
 200 $(s_t^{(i)}, s_{t-1}^{(i)}, \dots, s_{t-(m_i-1)}^{(i)})$, $\mathbf{a}^{(i)} = (\mathbf{0}_{m_i-1}, 1)$, $\gamma^{(i)} = (\gamma_i, \mathbf{0}_{m_i-1})$, $A = \bigoplus_{i=1}^T \tilde{A}_i$, and by re-
 201 placing $2k_i$ with m_i in the matrices presented above for the TBATS models.

202 Reduced forms

203 It is well known that linear forecasting systems have equivalent ARIMA (Box & Jenkins
 204 1970) reduced forms, and it has been shown that the forecasts from some exponential
 205 smoothing models are identical to the forecasts from particular ARIMA models (Chatfield
 206 & Yar 1991, McKenzie 1984). The reduced forms of BATS and TBATS models can be
 207 obtained by,

$$\varphi_p(L)\eta(L)y_t^{(\omega)} = \theta_q(L)\delta(L)\varepsilon_t, \quad (6)$$

208 where L is the lag operator, $\eta(L) = \det(\mathbf{I} - \mathbf{F}^*L)$, $\delta(L) = \mathbf{w}^* \text{adj}(\mathbf{I} - \mathbf{F}^*L) \mathbf{g}^*L + \det(\mathbf{I} -$
 209 $\mathbf{F}^*L)$, $\phi_p(L)$ and $\theta_q(L)$ are polynomials of length p and q , $\mathbf{w}^* = (1, \phi, \mathbf{a})$, $\mathbf{g}^* = (\alpha, \beta, \boldsymbol{\gamma})'$,
 210 and

$$\mathbf{F}^* = \begin{bmatrix} 1 & \phi & \mathbf{0} \\ 0 & \phi & \mathbf{0} \\ \mathbf{0}' & \mathbf{0}' & \mathbf{A} \end{bmatrix},$$

211 with the corresponding parameters defined as above. (Refer to [De Livera 2010b](#), Chap 4 for
 212 the proofs.)

For BATS models with a non-stationary growth, the reduced form is then given by (6), with

$$\begin{aligned} \eta(L) &= (1 - \phi L)(1 - L) \prod_{j=1}^T (L^{m_j-1} + L^{m_j-2} + \dots + L + 1) \\ \delta(L) &= \prod_{j=1}^T (L^{m_j-1} + L^{m_j-2} + \dots + L + 1) [L^2(\phi - \phi\alpha) + L(\alpha + \phi\beta - \phi - 1) + 1] \\ &\quad + (1 - \phi L) \sum_{i=1}^T \prod_{j=1, j \neq i}^T (L^{m_j-1} + L^{m_j-2} + \dots + L + 1) \gamma_i L^{m_i}. \end{aligned}$$

For TBATS models, with a non-stationary growth the reduced form is then given by (6),

with

$$\begin{aligned} \eta(L) &= (1 - L)(1 - \phi L) \prod_{i=1}^T \prod_{j=1}^{k_i} (1 - 2 \cos \lambda_j^{(i)} L + L^2) \\ \delta(L) &= [L^2 \phi(1 - \alpha) + L(\alpha + \phi\beta - \phi - 1) + 1] \prod_{i=1}^T \prod_{j=1}^{k_i} (1 - 2 \cos \lambda_j^{(i)} L + L^2) + \\ &\quad (1 - L)(1 - \phi L) \sum_{i=1}^T \sum_{\tilde{i}=1, \tilde{i} \neq i}^{k_i} \prod_{\tilde{j}=1, \tilde{j} \neq j}^{k_{\tilde{i}}} (1 - 2 \cos \lambda_{\tilde{j}}^{(\tilde{i})} L + L^2) [(\cos \lambda_j^{(i)} \gamma_{1i} + \sin \lambda_j^{(i)} \gamma_{2i}) L^2 - \gamma_{1i} L^3] \\ &\quad + (1 - L)(1 - \phi L) L \prod_{i=1}^T \prod_{j=1}^{k_i} (1 - 2 \cos \lambda_j^{(i)} L + L^2) \sum_{i=1}^T k_i \gamma_{1i}. \end{aligned}$$

213 The reduced form of the model implies that the TBATS model has a relatively complex
 214 ARIMA structure which is dependent on the number of terms k_i chosen for the i th seasonal

215 component, and encompasses trigonometric coefficients that rely on the frequency of
216 each seasonal pattern. This highlights the fact that the state space form of the model
217 has advantages over its reduced form counterpart, as it leads to a more logical derivation,
218 providing a coherent interpretation of the components with a simpler estimation procedure.

219 **3 Estimation, prediction, and model selection**

220 **3.1 Estimation**

221 The typical approach with linear state space models is to estimate unknown parameters like
222 the smoothing parameters and the damping parameter using the sum of squared errors or the
223 Gaussian likelihood (see [Hyndman et al. 2008](#), Chap 3). In our context it is necessary to also
224 estimate the unknown Box-Cox transformation parameter ω , and the ARMA coefficients.

225 The seed states of state space models are usually treated as random vectors. Given trial
226 values of the unknown parameters, the joint steady state distributions of stationary states
227 are derived, and then assigned to associated seed states. Thus, for given values of ϕ and σ^2 ,
228 the seed short-run growth rate would be assigned an $N(0, \sigma^2/(1 - \phi^2))$ distribution. Most
229 states, however, are non-stationary, and they are presumed to have Gaussian distributions
230 with arbitrarily large variances ([Ansley & Kohn 1985](#)). A Kalman filter is typically used
231 to obtain one-step ahead prediction errors and associated variances needed for evaluating
232 fitting criteria for given trial values of the parameters. The Kalman filter in [Snyder \(1985b\)](#)
233 would be appropriate for innovations state space models in particular. However, like all
234 Kalman filters it would need to be augmented with additional equations ([De Jong 1991](#)) to
235 handle the non-stationary states.

236 A simpler alternative is available in the context of innovations state space models. By condi-
237 tioning on all the seed states and treating them as unknown fixed parameters, exponential

238 smoothing can be used instead of an augmented Kalman filter to generate the one-step ahead
 239 prediction errors needed for likelihood evaluation. In this case both the parameters and seed
 240 states are selected to maximize the resulting *conditional* likelihood function. If not for the
 241 different treatment of stationary states, the exponential smoothing and augmented Kalman
 242 filter approaches yield the same conditional distribution of the final state vector and so yield
 243 identical prediction distributions of future series values (Hyndman et al. 2008).

244 The conditional likelihood of the observed data $\mathbf{y} = (y_1, \dots, y_n)$ is derived on the as-
 245 sumption that $\varepsilon_t \sim N(0, \sigma^2)$. This implies that the density of the transformed series is
 246 $y_t^{(\omega)} \sim N(\mathbf{w}'\mathbf{x}_{t-1}, \sigma^2)$ so that the density of the transformed data is

$$p(\mathbf{y}^{(\omega)} | \mathbf{x}_0, \boldsymbol{\vartheta}, \sigma^2) = \prod_{t=1}^n p(y_t^{(\omega)} | \mathbf{x}_{t-1}, \boldsymbol{\vartheta}, \sigma^2) = \prod_{t=1}^n p(\varepsilon_t) = \frac{1}{(2\pi\sigma^2)^{\frac{n}{2}}} \exp\left(\frac{-1}{2\sigma^2} \sum_{t=1}^n \varepsilon_t^2\right)$$

where $\boldsymbol{\vartheta}$ is a vector containing the Box-Cox parameter, smoothing parameters and ARMA
 coefficients. Therefore, the density of the original series, using the Jacobian of the Box-Cox
 transformation, is

$$\begin{aligned} p(y_t | \mathbf{x}_0, \boldsymbol{\vartheta}, \sigma^2) &= p(y_t^{(\omega)} | \mathbf{x}_0, \boldsymbol{\vartheta}, \sigma^2) \left| \det\left(\frac{\partial y_t^{(\omega)}}{\partial y}\right) \right| = p(y_t^{(\omega)} | \mathbf{x}_0, \boldsymbol{\vartheta}, \sigma^2) \prod_{t=1}^n y_t^{\omega-1} \\ &= \frac{1}{(2\pi\sigma^2)^{\frac{n}{2}}} \exp\left(\frac{-1}{2\sigma^2} \sum_{t=1}^n \varepsilon_t^2\right) \prod_{t=1}^n y_t^{\omega-1}. \end{aligned}$$

On concentrating out the variance σ^2 with its maximum likelihood estimate

$$\hat{\sigma}^2 = n^{-1} \sum_{t=1}^n \varepsilon_t^2, \quad (7)$$

we obtain the log-likelihood given by

$$\mathcal{L}(\mathbf{x}_0, \boldsymbol{\vartheta}, \sigma^2) = \frac{-n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{t=1}^n \varepsilon_t^2 + (\omega - 1) \sum_{t=1}^n \log y_t. \quad (8)$$

Substituting (7) into (8), multiplying by -2 , and omitting constant terms, we get

$$\mathcal{L}^*(\mathbf{x}_0, \boldsymbol{\vartheta}) = n \log \left(\sum_{t=1}^n \varepsilon_t^2 \right) - 2(\omega - 1) \sum_{t=1}^n \log y_t. \quad (9)$$

247 The quest is to minimize the quantity (9) to obtain maximum likelihood estimates, but
 248 the dimension of the seed states vector \mathbf{x}_0 makes this computationally challenging. Our
 249 approach to this problem is based on the observation that ε_t is a linear function of the
 250 seed vector \mathbf{x}_0 . Thus, we show that it is possible to concentrate the seed states out of the
 251 likelihood, and so substantially reduce the dimension of the numerical optimization problem.
 252 This concentration process is the exponential smoothing analogue of de Jong's method for
 253 augmenting Kalman filters to handle seed states with infinite variances.

The innovation ε_t can be eliminated from the transition equation in (5) to give $\mathbf{x}_t = \mathbf{D}\mathbf{x}_{t-1} + \mathbf{g}y_t$ where $\mathbf{D} = \mathbf{F} - \mathbf{g}\mathbf{w}'$. The equation for the state, obtained by back-solving this recurrence equation to period 0, can be used in conjunction with the measurement equation to obtain

$$\begin{aligned} \varepsilon_t &= y_t^{(\omega)} - \mathbf{w}' \sum_{j=1}^{t-1} \mathbf{D}^{j-1} \mathbf{g} y_{t-j}^{(\omega)} - \mathbf{w}' \mathbf{D}^{t-1} \mathbf{x}_0, \\ &= y_t^{(\omega)} - \mathbf{w}' \tilde{\mathbf{x}}_{t-1} - \mathbf{w}'_{t-1} \mathbf{x}_0, \\ &= \tilde{y}_t - \mathbf{w}'_{t-1} \mathbf{x}_0, \end{aligned} \quad (10)$$

where $\tilde{y}_t = y_t^{(\omega)} - \mathbf{w}' \tilde{\mathbf{x}}_{t-1}$, $\tilde{\mathbf{x}}_t = \mathbf{D}\tilde{\mathbf{x}}_{t-1} + \mathbf{g}y_t$, $\mathbf{w}'_t = \mathbf{D}\mathbf{w}'_{t-1}$, $\tilde{\mathbf{x}}_0 = 0$ and $\mathbf{w}'_0 = \mathbf{w}'$ (see Snyder 1985a, for the derivation). Thus, the relationship between each error and the initial state vector \mathbf{x}_0 is linear. It can also be seen from (10) that the seed vector \mathbf{x}_0 corresponds to a regression coefficients vector, and so it may be estimated using conventional linear least squares methods. Thus, the problem reduces to minimizing the following with respect to $\boldsymbol{\vartheta}$:

$$\mathcal{L}^*(\boldsymbol{\vartheta}) = n \log(\text{SSE}^*) - 2(\omega - 1) \sum_{t=1}^n \log y_t, \quad (11)$$

254 where SSE^* is the optimized value of the sum of squared errors for given parameter values.

255 In contrast to the existing estimation procedure (Hyndman et al. 2008) where a numerical
 256 optimizer is used to find the values of both the initial states and the parameters, our approach
 257 concentrates out the initial state values from the likelihood function, leaving only the much
 258 smaller parameter vector for optimization, a tactic in our experience that leads to better
 259 forecasts. It is also effective in reducing computational times instead of invoking the
 260 numerical optimizer to directly estimate the seed state vector.

261 We can constrain the estimation to the *forecastability region* (Hyndman et al. 2007) so that
 262 the characteristic roots of D lie within the unit circle, a concept that is equivalent to the
 263 invertibility condition for equivalent ARIMA models. The coefficients $w'D^{j-1}g$ are the
 264 matrix analogues of the weights in an exponentially weighted average, and this constraint
 265 ensures that their effect is to reduce the importance placed on older data. When some roots
 266 lie on the unit circle, the discounting effect is lost (although this possibility admits some
 267 important special cases). For integer period seasonality, the seasonal values can also be
 268 constrained when optimizing, so that each seasonal component sums to zero.

269 3.2 Prediction

The prediction distribution in the transformed space for future period $n+h$, given the final state vector x_n and given the parameters ϑ, σ^2 , is Gaussian. The associated random variable is designated by $y_{n+h|n}^{(\omega)}$. Its mean $E(y_{n+h|n}^{(\omega)})$ and variance $V(y_{n+h|n}^{(\omega)})$ are given, after allowing for the Box-Cox transformation, by the equations (Hyndman et al. 2005):

$$E(y_{n+h|n}^{(\omega)}) = w'F^{h-1}x_n \quad (12a)$$

$$V(y_{n+h|n}^{(\omega)}) = \begin{cases} \sigma^2 & \text{if } h = 1; \\ \sigma^2 \left[1 + \sum_{j=1}^{h-1} c_j^2 \right] & \text{if } h \geq 2; \end{cases} \quad (12b)$$

270 where $c_j = \mathbf{w}'\mathbf{F}^{j-1}\mathbf{g}$. The prediction distribution of $y_{n+h|n}$ is not normal. Point forecasts
271 and forecast intervals, however, may be obtained using the inverse Box-Cox transformation
272 of appropriate quantiles of the distribution of $y_{n+h|n}^{(\omega)}$. The point forecast obtained this
273 way is the median, a minimum mean absolute error predictor (Pankratz & Dudley 1987,
274 Proietti & Riani 2009). The prediction intervals retain the required probability coverage
275 under back-transformation because the Box-Cox transformation is monotonically increasing.
276 To simplify matters we use the common *plug-in* approach to forecasting. The pertinent
277 parameters and final state are replaced by their estimates in the above formulae. This
278 ignores the impact of estimation error, but the latter is a second-order effect in most practical
279 contexts.

280 3.3 Model selection

281 The use of an information criterion

282 In this paper, the $\text{AIC} = \mathcal{L}^*(\hat{\boldsymbol{\vartheta}}, \hat{\mathbf{x}}_0) + 2K$ is used for choosing between the models, where
283 K is the total number of parameters in $\boldsymbol{\vartheta}$ plus the number of free states in \mathbf{x}_0 , and $\hat{\boldsymbol{\vartheta}}$ and $\hat{\mathbf{x}}_0$
284 denote the estimates of $\boldsymbol{\vartheta}$ and \mathbf{x}_0 . When one of the smoothing parameters takes the boundary
285 value 0, the value of K is reduced by one as the model simplifies to a special case. For
286 example, if $\beta = 0$, then $b_t = b_0$ for all t . Similarly, when either $\phi = 1$ or $\omega = 1$, the value of
287 K is reduced by one in each instance to account for the resulting simplified model. In an
288 empirical study, Billah et al. (2005) indicated that information criterion approaches, such
289 as the AIC, provide the best basis for automated model selection, relative to other methods
290 such as prediction-validation. Alternative information criteria such as the AICc (Burnham &
291 Anderson 2002) may also be used.

292 **Selecting the number of harmonics k_i in the trigonometric models**

293 The forecasts from the TBATS model depend on the number of harmonics k_i used for the
294 seasonal component i . It is impractical to consider all possible combinations in the quest for
295 the best combination. After much experimentation we found that the following approach
296 leads to good models and that further improvement can rarely be achieved (see [De Livera](#)
297 [2010b](#), Chap 3).

298 De-trend the first few seasons of the transformed data using an appropriate de-trending
299 method. In this paper, we employed the method described by [Hyndman et al. \(2008,](#)
300 [Chap 2\)](#). Approximate the resulting de-trended data using the linear regression
301 $\sum_{i=1}^T \sum_{j=1}^{k_i} a_j^{(i)} \cos(\lambda_j^{(i)} t) + b_j^{(i)} \sin(\lambda_j^{(i)} t)$. Starting with a single harmonic, gradually add
302 harmonics, testing the significance of each one using F -tests. Let k_i^* be the number of
303 significant harmonics (with $p < 0.001$) for the i th seasonal component. Then fit the required
304 model to the data with $k_i = k_i^*$ and compute the AIC. Considering one seasonal component at
305 a time, repeatedly fit the model to the estimation sample, gradually increasing k_i but holding
306 all other harmonics constant for each i , until the minimum AIC is achieved. This approach
307 based on multiple linear regression, is preferred over letting $k_i^* = 1$ for each component, as
308 the latter was found to be unnecessarily time consuming.

309 **Selecting the ARMA orders p and q for the models**

310 In selecting a model, suitable values for the ARMA orders p and q must also be found.
311 We do this using a two-step procedure. First, a suitable model with no ARMA component
312 is selected. Then the automatic ARIMA algorithm of [Hyndman & Khandakar \(2008\)](#) is
313 applied to the residuals from this model in order to determine the appropriate orders p and q
314 (we assume the residuals are stationary). The selected model is then fitted again but with
315 an ARMA(p, q) error component, where the ARMA coefficients are estimated jointly with

316 the rest of the parameters. The ARMA component is only retained if the resulting model
317 has lower AIC than the model with no ARMA component. Our subsequent work on the
318 proposed models on a large number of real time series (De Livera 2010a) have indicated that
319 fitting ARMA in a two step approach yielded the best out-of-sample predictions, compared
320 to several alternative approaches.

321 4 Applications of the proposed models

322 The results obtained from the application of BATS and TBATS to the three complex time
323 series in Figure 1 are reported in this section.

324 In addition, it is shown that the TBATS models can be used as a means of decomposing
325 complex seasonal time series into trend, seasonal and irregular components. In decomposing
326 time series, the trigonometric approach has several important advantages over the traditional
327 seasonal formulation. First, seasonal components obtained from the BATS model are not
328 normalized (cp. the seasonal model of Harrison & Stevens 1976). Although normalized
329 components may not be necessary if one is only interested in the forecasts and the prediction
330 intervals, when the seasonal component is to be analyzed separately or used for seasonal
331 adjustment, normalized seasonal components are required (Archibald & Koehler 2003, Hyndman et al. 2008). Thus, BATS models have to be modified, so that the seasonal components
332 are normalized for each time period, before using them for time series decomposition (see
333 De Livera 2010b, Chap 5 for a normalized version of the BATS model). In contrast, the
334 trigonometric terms in TBATS models do not require normalization, and so are more appropriate for decomposition. Second, in estimating the seasonal components using BATS, a
335 large number of parameters are required, which often leads to noisy seasonal components. In
336 contrast, a smoother seasonal decomposition is expected from TBATS where the smoothness
337 of the seasonal component is controlled by the number of harmonics used. Furthermore,
338
339

340 a BATS model cannot be used to decompose time series with non-integer seasonality and
341 dual calendar effects. Using TBATS models for complex seasonal time series, the overall
342 seasonal component can be decomposed into several individual seasonal components with
343 different frequencies. These individual seasonal components are given by $s_t^{(i)}$ ($i = 1, \dots, T$)
344 and the trend component is obtained by ℓ_t . Extracting the trend and seasonal components
345 then leaves behind a covariance stationary irregular component, denoted by d_t . In particular,
346 this approach leads to the identification and extraction of one or more seasonal components,
347 which may not be apparent in the time series plots themselves.

348 4.1 Application to weekly US gasoline data

349 Figure 1(b) shows the number of barrels of motor gasoline product supplied in the United
350 States, in thousands of barrels per day, from February 1991 to July 2005 (see [www.
351 forecastingprinciples.com/files/T_competition_new.pdf](http://www.forecastingprinciples.com/files/T_competition_new.pdf) for details). The
352 data are observed weekly and show a strong annual seasonal pattern. The length of sea-
353 sonality of the time series is $m_1 = 365.25/7 \approx 52.179$. The time series exhibits an upward
354 additive trend and an additive seasonal pattern; that is, a pattern for which the variation does
355 not change with the level of the time series.

356 The series, which consists of 745 observations, was split into two segments: an estimation
357 sample period (484 observations) and test sample (261 observations). The estimation sample
358 was used to obtain the maximum likelihood estimates of the initial states and the smoothing
359 parameters, and to select the appropriate number of harmonics and ARMA orders. Following
360 the procedure for finding the number of harmonics to start with, it was found that only one
361 harmonic was highly significant. The model was then fitted to the whole estimation sample
362 of 484 values by minimizing the criterion equation (11). The values of the AIC decreased
363 until $k_1 = 7$, and then started to increase.

Out-of-sample performance was measured by the Root Mean Square Error (RMSE), defined as

$$\text{RMSE}_h = \sqrt{\frac{1}{p-h+1} \sum_{t=n}^{n+p-h} (y_{t+h} - \hat{y}_{t+h|t})^2}, \quad (13)$$

where $p = 261$ is the length of the test sample, $n = 484$ is the length of the estimation sample and h is the length of the forecast horizon. Further analysis showed that changing the value of k_1 from 7 generated worse out of sample results, indicating that the use of the AIC as the criterion for this model selection procedure is a reasonable choice. In this way, the TBATS(0.9922, 1, 0, 0, {365.25/7, 7}) model was obtained. As a second step, ARMA models were fitted to the residuals with (p, q) combinations up to $p = q = 5$, and it was discovered that the TBATS(0.9922, 1, 0, 1, {365.25/7, 7}) model minimizes the AIC.

The BATS model was considered next with $m_1 = 52$, and following the above procedure, it was discovered that the BATS(0.9875, 1, 0, 1, 52) model minimized the AIC. Figure 2 shows the out-of-sample RMSEs obtained for the two models, and it can be seen that the trigonometric model performs better for all lead times.

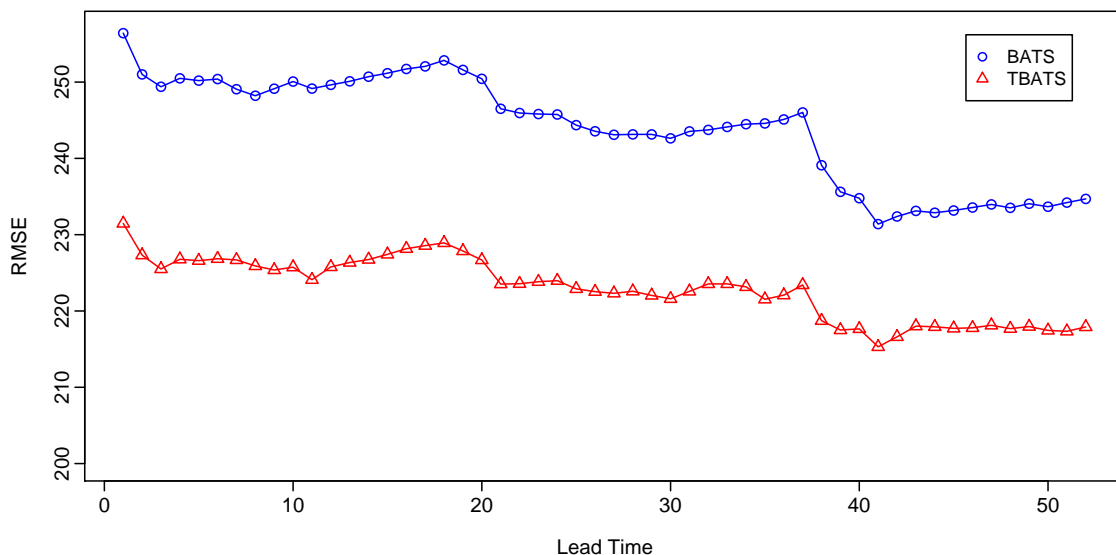


Figure 2: Out-of-sample results for the US gasoline data using BATS(0.9875, 1, 0, 1, 52) and TBATS(0.9922, 1, 0, 1, {365.25/7, 7}).

Table 1: Parameters chosen for each application of the TBATS model.

Data	Parameters															
	ω	ϕ	α	β	γ_1	γ_2	γ_3	γ_4	γ_5	γ_6	θ_1	θ_2	φ_1	φ_2	φ_3	
Gasoline	0.9922	1	0.0478	0	0.0036	-0.0005						-0.2124				
Call center	1		0.0921	0	0.0006	-0.0002	0.0022	0.0020				-0.1353		0.1776	-0.0144	-0.0200
Electricity	0.1393	1	0.8019	0	0.0034	-0.0037	0.0001	0.0001	0.1171	0.251	-1.2610	0.3128	1.0570	-0.2991	-0.1208	

Table 2: Parameters chosen for each application of the BATS model.

Data	Parameters												
	ω	ϕ	α	β	γ_1	γ_2	γ_3	θ_1	θ_2	φ_1	φ_2	φ_3	
Gasoline	0.9875	1	0.0457	0	0.2246					-0.2186			
Call center	0.4306		0.0368	0	0.0001	0.0001				0.0552	0.1573	0.1411	
Electricity	0.0013	1	0.2216	0	0	0	0						

375 The BATS model cannot handle the non-integer periods, and so has to be rounded off to the
 376 nearest integer. It may also be over-parameterized, as 52 initial seasonal values have to be
 377 estimated. Both these problems are overcome in the trigonometric formulation.

378 Tables 1 and 2 show the estimated parameters obtained for the TBATS and BATS models
 379 respectively. The estimated values of 0 for β and 1 for ϕ imply a purely deterministic
 380 growth rate with no damping effect. The models also imply that the irregular component of
 381 the series is correlated and can be described by an ARMA(0,1) process, and that a strong
 382 transformation is not necessary to handle nonlinearities in the series.

383 The decomposition of the Gasoline time series, obtained from the fitted TBATS model, is
 384 shown in Figure 3. The vertical bars at the right side of each plot represent equal heights
 385 plotted on different scales, thus providing a comparison of the size of each component. The
 386 trigonometric formulation in TBATS allows for the removal of more randomness from the
 387 seasonal component without destroying the influential bumps.

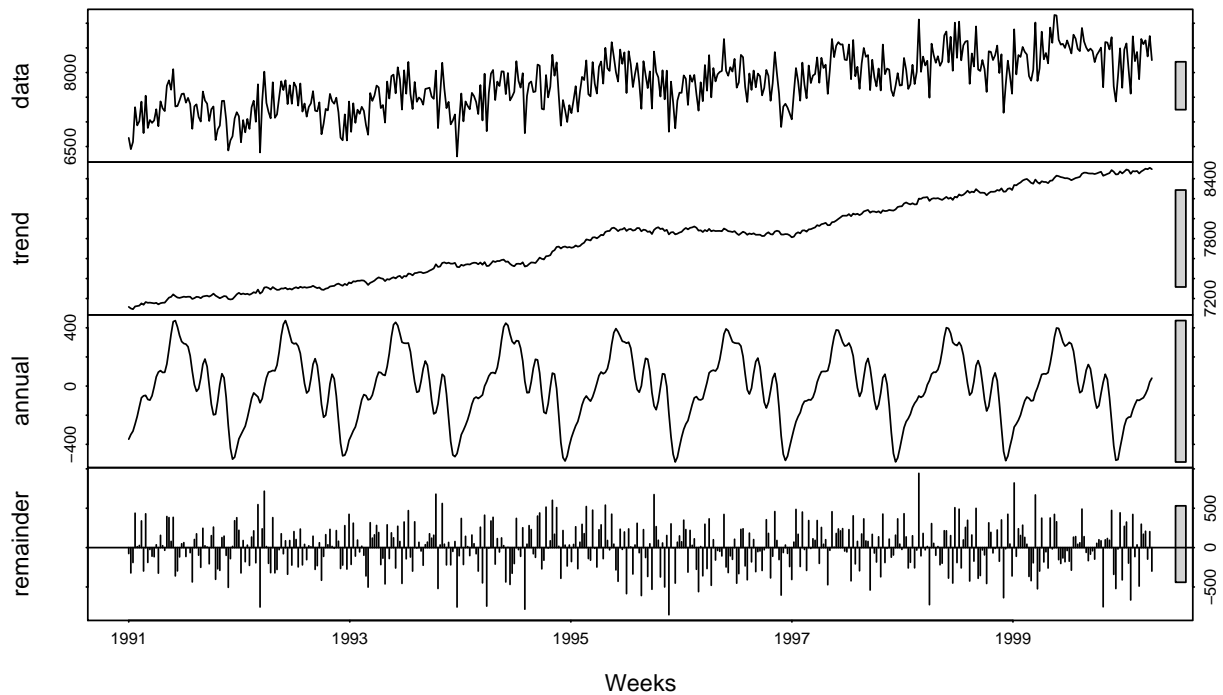


Figure 3: Trigonometric decomposition of the US gasoline data. The within sample RMSE was 279.9.

388 4.2 Application to call center data

389 The call center series in Figure 1(a) consists of 10,140 observations, that is 12 weeks of
 390 data starting from 3 March 2003 (Weinberg et al. 2007). It contains a daily seasonal pattern
 391 with period 169 and a weekly seasonal pattern with period $169 * 5 = 845$. The fitting sample
 392 consists of 7,605 observations (9 weeks). As the trend appears to be close to zero, the growth
 393 rate b_t was omitted from the models.

394 The selection procedure led to the models TBATS(1, NA, 3, 1, {169, 29}, {845, 15}) and
 395 BATS(0.4306, NA, 3, 0, 169, 845). Other BATS models with $\omega = 1$ were also tried, but their
 396 forecasting performance was worse.

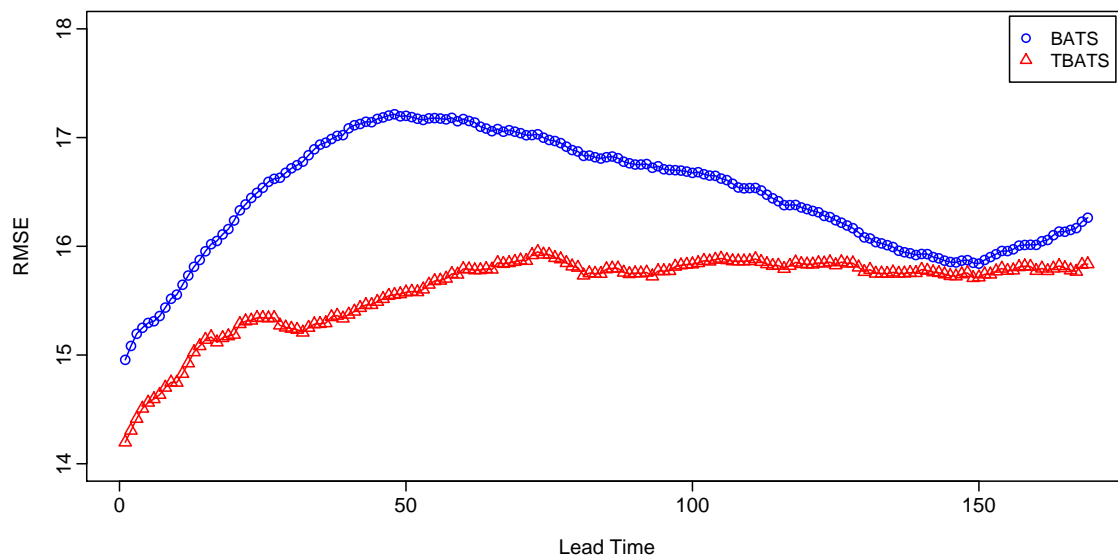


Figure 4: *Out-of-sample results for the call center data using $BATS(0.4306, NA, 3, 0, 169, 845)$ and $TBATS(1, NA, 3, 1, \{169, 29\}, \{845, 15\})$.*

397 The post-sample forecasting accuracies of the selected BATS and TBATS models are
 398 compared in Figure 4. Again TBATS, which requires fewer parameters to be estimated, is
 399 more accurate than BATS.

400 The estimated parameters for the TBATS model shown in Table 1 imply that no Box-Cox
 401 transformation is necessary for this time series, and that the weekly seasonal component is
 402 more variable than the daily seasonal component. The irregular component is modeled by
 403 an ARMA(3,1) process.

404 The decomposition obtained from TBATS, as shown in Figure 5, clearly exhibits strong daily
 405 and weekly seasonal components. The weekly seasonal pattern evolves considerably over
 406 time but the daily seasonal pattern is relatively stable. As is seen from the time series plot
 407 itself, the trend component is very small in magnitude compared to the seasonal components.

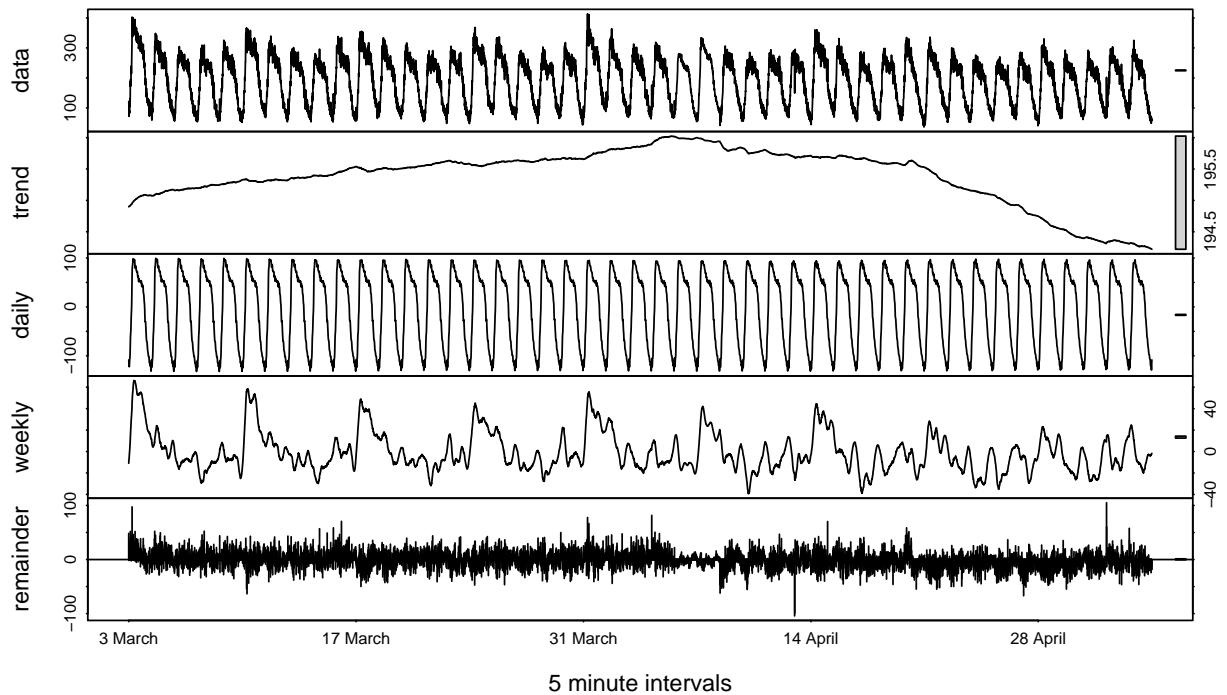


Figure 5: Trigonometric decomposition of the call center data. The within sample RMSE was 14.7.

4.3 Application to the Turkey electricity demand data

The Turkey electricity demand series shown in Figure 1(c) has a number of important features that should be reflected in the model structure. Three seasonal components with frequencies $m_1 = 7$, $m_2 = 354.37$ and $m_3 = 365.25$ exist in the series. The sharp drops seen in the seasonal component with period 354.37 are due to the Seker (also known as *Eid ul-Fitr*) and Kurban (also known as *Eid al-Adha*) religious holidays, which follow the Hijri calendar, while those seen in the seasonal component with frequency 365.25 are due to national holidays which follow the Gregorian calendar. Table 3 gives the dates of the holidays from the Hijri and Gregorian calendars. Seker is a three-day festival when sweets are eaten to celebrate the end of the fast of Ramadan. Kurban is a four-day festival when sacrificial sheep are slaughtered and their meat distributed to the poor. In addition, there are national holidays which follow the Gregorian calendar as shown in the table.

Table 3: *The dates of Turkish holidays between 1 January 2000 to 31 December 2006.*

Year	Religious holidays		National holidays
	Seker holiday	Kurban holiday	
2000	08 Jan–10 Jan 27 Dec–29 Dec	16 Mar–19 Mar	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2001	16 Dec–18 Dec	05 Mar–08 Mar	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2002	05 Dec–07 Dec	22 Feb–25 Feb	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2003	25 Nov–27 Nov	11 Feb–14 Feb	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2004	14 Nov–16 Nov	01 Feb–04 Feb	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2005	03 Nov–05 Nov	20 Jan–23 Jan	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct
2006	23 Oct–25 Oct	10 Jan–13 Jan 31 Dec	01 Jan, 23 Apr, 19 May, 30 Aug, 29 Oct

Table 4: *Number of estimated parameters for each model in each application.*

Data	Model	No.parameters
Gasoline	BATS(0.9875, 1, 0, 1, 52)	60
	TBATS(0.9922, 1, 0, 1, {365.25/7, 7})	23
Call center	BATS(0.4306, NA, 3, 0, 169, 845)	1026
	TBATS(1, NA, 3, 1, {169, 29}, {845, 15})	102
Electricity	BATS(0.0013, 1, 0, 7, 354, 365)	735
	TBATS(0.1393, 1, 3, 2, {7, 3}, {354.37, 23}, {365.25, 3})	79

420 In this study, the series, which covers a period of 9 years, is split into two parts:
 421 a fitting sample of $n = 2191$ observations (6 years) and a post-sample period of
 422 $p = 1096$ observations (3 years). The model selection procedure was followed to
 423 give the TBATS(0.1393, 1, 3, 2, {7, 3}, {354.37, 23}, {365.25, 3}) and BATS(0.0013, 1, 0, 0,
 424 7, 354, 365) models.

425 Figure 6 shows that a better post-sample forecasting performance is again obtained from
 426 the TBATS model. The poor performance of the BATS model may be explained by its
 427 inability to capture the dual seasonal calendar effects and the large number of values that is
 428 required to be estimated. The estimated zero values for the smoothing parameters shown
 429 in Table 2 for the BATS solution suggest stable seasonal components. The Hijri seasonal
 430 component based on the TBATS solution displays a similar level of stability. However,
 431 moderate change is implied by the TBATS model in the weekly and Gregorian seasonal
 432 components. Both models required strong Box-Cox transformations in order to handle the
 433 obvious non-linearity in the time series plot.

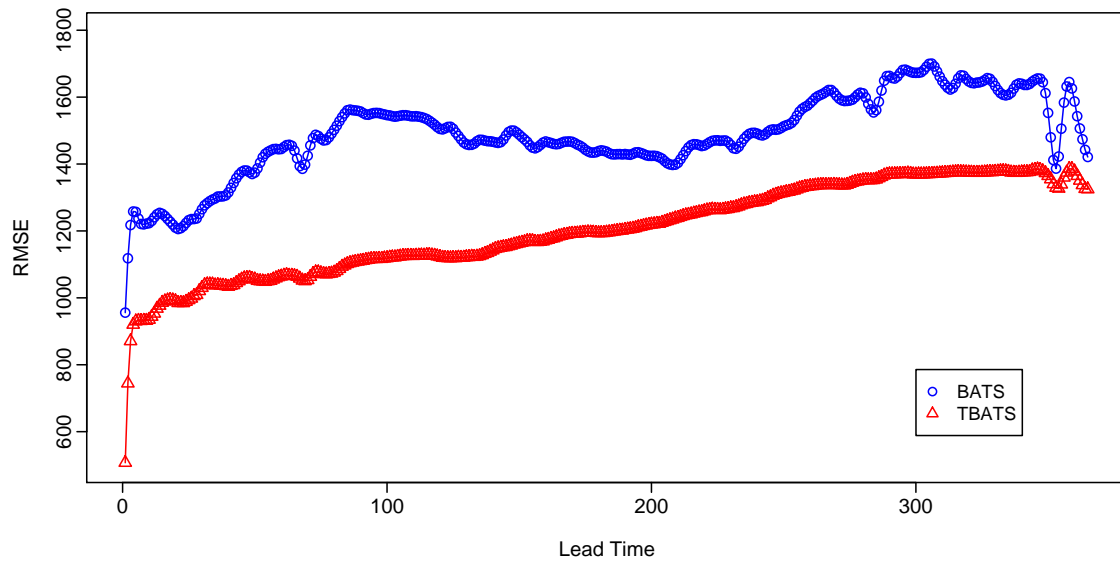


Figure 6: Out-of-sample results for the Turkey electricity demand data using $BATS(0.0013, 1, 0, 0, 7, 354, 365)$ and $TBATS(0.1393, 1, 3, 2, \{7, 3\}, \{354.37, 23\}, \{365.25, 3\})$.

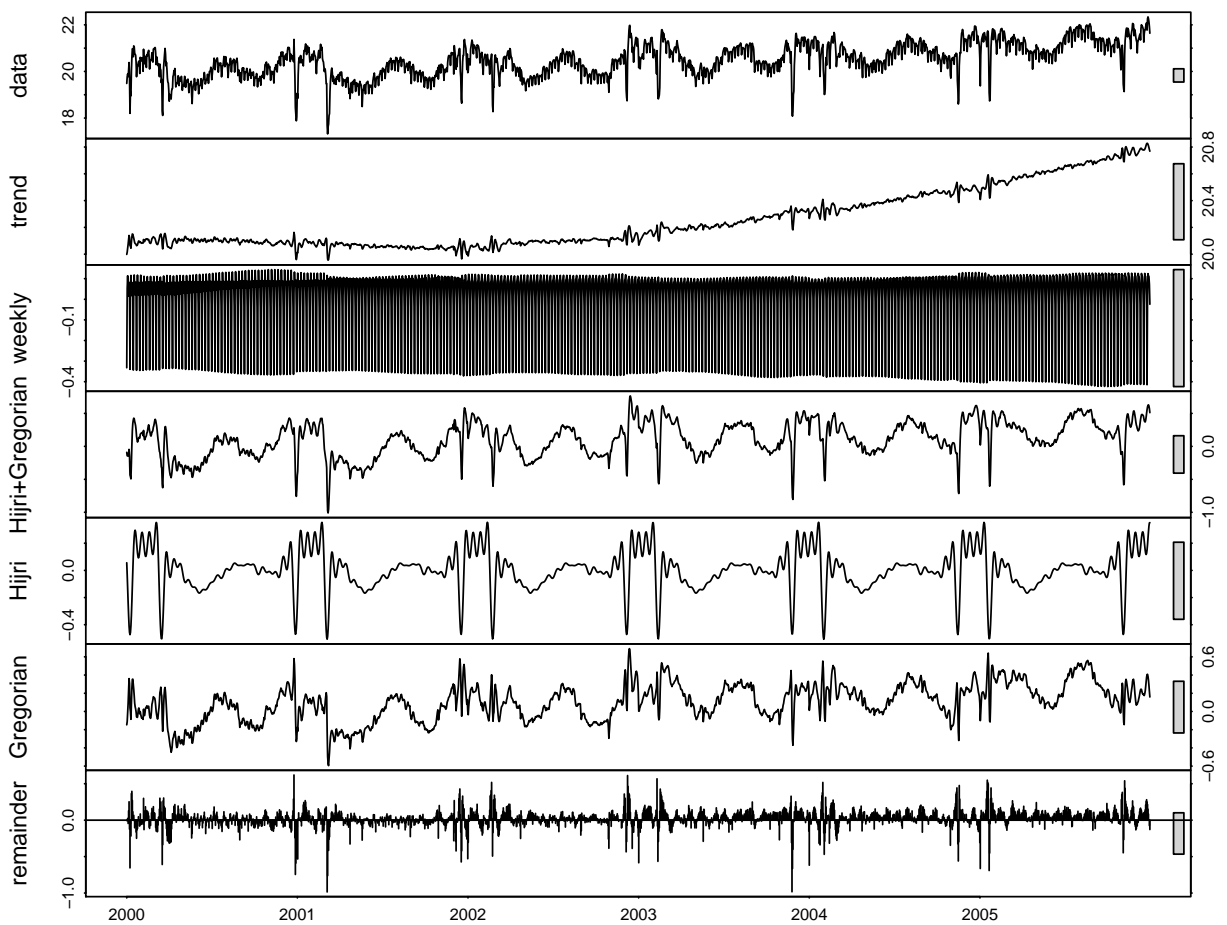


Figure 7: Trigonometric decomposition of the Turkey electricity demand data. The within sample RMSE was 0.1346.

434 The decomposition of the series obtained by using the chosen TBATS model, is shown in
435 Figure 7. The first panel shows the transformed observations and the second shows the trend
436 component. The third panel shows the weekly seasonal component with period 7, and the
437 fifth and the sixth panels show the seasonal component based on the Hijri calendar with
438 period 354.37 and the seasonal component based on the Gregorian calendar with period
439 365.25 respectively. The seasonal components shown in the fifth and sixth panels may
440 initially appear to be mirror images. However, their combined effect, shown in the fourth
441 panel, indicates that this is not the case. Interpreting their combined effect as *the annual*
442 *seasonal component* would be misleading as there is no unique annual calendar in this
443 situation: both constituent calendars are of different lengths.

444 The rather wiggly components of the decomposition are probably due to the use of a large
445 number of harmonics in each seasonal component. This is necessary to capture the sharp
446 drops seen in the time series plot. If we were to augment the stochastic seasonal component
447 by deterministic holiday effects (given in Table 3) represented by dummy variables, the
448 number of harmonics required might be reduced. Using a trend component, a seasonal
449 component and holiday dummy variables, regression was performed on the transformed $y_t^{(\omega)}$
450 values. The term $\sum_{i=1}^3 \sum_{j=1}^{k_i} a_j^{(i)} \cos(\lambda_j^{(i)} t) + b_j^{(i)} \sin(\lambda_j^{(i)} t)$ was used to capture the multiple
451 seasonality with $k_1 = 3$ and $k_2 = k_3 = 1$. The estimated holiday effect was then removed
452 from the series and the remainder was decomposed using TBATS to achieve the result shown
453 in Figure 8, which provides a much smoother seasonal decomposition. Again, the sum of
454 the Hijri seasonal component and the Gregorian seasonal component shown in the fourth
455 panel illustrates that the Hijri and Gregorian seasonal components do not offset each other.

456 This analysis demonstrates the capability of our trigonometric decomposition in extracting
457 those seasonal components which are otherwise not apparent in graphical displays.

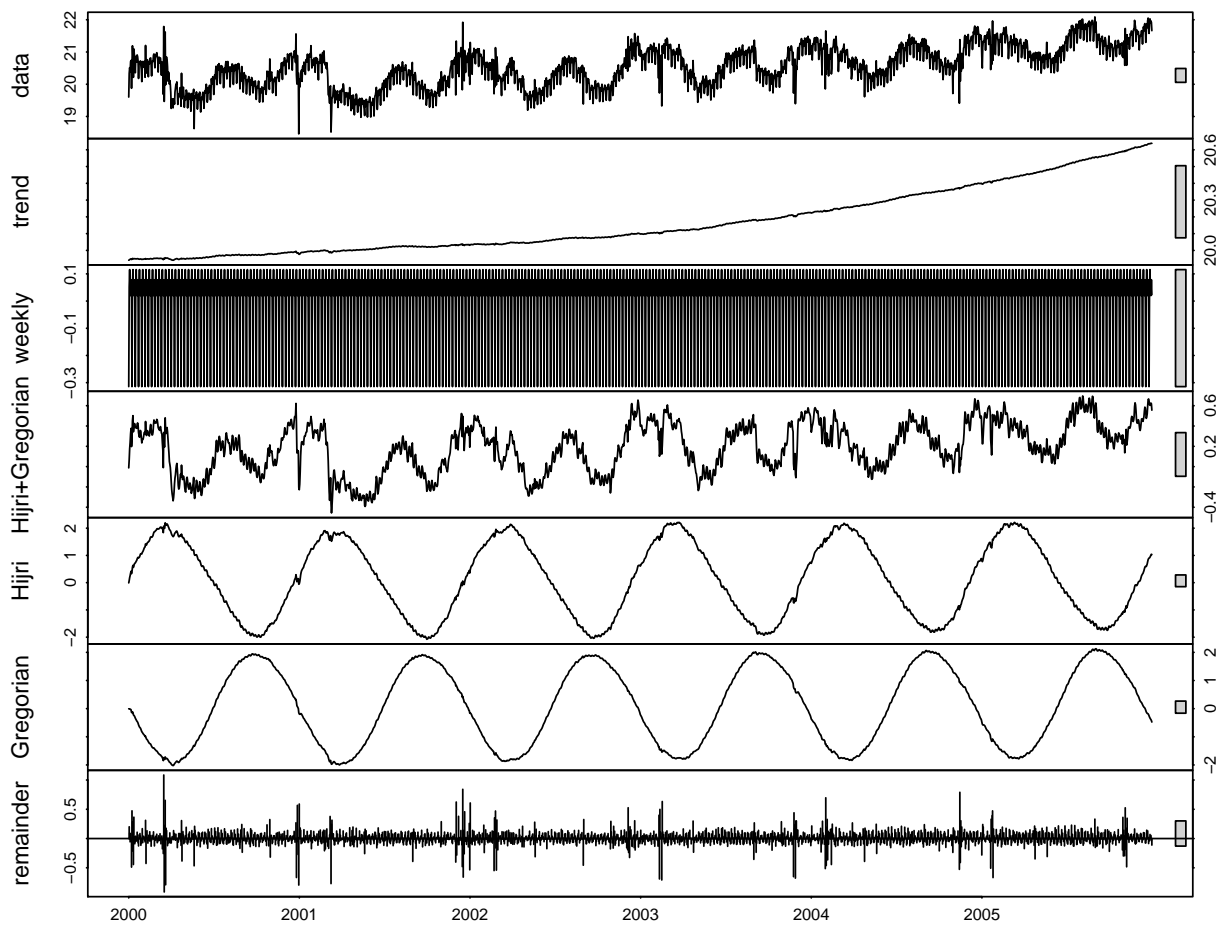


Figure 8: *Trigonometric decomposition of the regressed Turkey electricity demand data. The within sample RMSE was 0.1296.*

458 In forecasting complex seasonal time series with such deterministic effects, both BATS and
 459 TBATS models may be extended to accommodate regressor variables, allowing additional
 460 information to be included in the models (see [De Livera 2010a](#), Chap 7 for a detailed
 461 discussion of the BATS and TBATS models with regressor variables.)

462 5 Concluding remarks

463 A new state space modeling framework, based on the innovations approach, was developed
 464 for forecasting time series with complex seasonal patterns. The new approaches offer
 465 alternatives to traditional counterparts, providing several advantages and additional options.

466 A key feature of the proposed trigonometric framework is its ability to model both linear and
467 non-linear time series with single seasonality, multiple seasonality, high period seasonality,
468 non-integer seasonality and dual calendar effects. We are not aware of another modeling
469 procedure that is able to forecast and decompose all these complex seasonal time series
470 features within a single framework.

471 In addition, the framework consists of a new estimation procedure which is sufficiently
472 general to be applied to *any* innovations state space model. By relying on maximum
473 likelihood estimation, it avoids the ad hoc startup choices with unknown statistical properties
474 commonly used with exponential smoothing. By incorporating the least-squares criterion, it
475 streamlines the process of obtaining the maximum likelihood estimates.

476 The applications of the proposed modeling framework to three complex seasonal time series
477 demonstrated that the trigonometric models led to a better out of sample performance with
478 substantially fewer values to be estimated than traditional seasonal exponential smoothing
479 approaches (see Table 4). The trigonometric approach was also illustrated as a means of
480 decomposing complex seasonal time series.

481 A further advantage of the proposed framework is its adaptability. It can be altered to
482 encompass various deterministic effects that are often seen in real life time series. For
483 instance, the moving holidays such as Easter and irregular holidays can be handled by
484 incorporating dummy variables in the models (see [De Livera 2010a](#), Chap 7), and the
485 varying length of months can be managed by adjusting the data for trading days before
486 modeling (see [Makridakis et al. 1998](#)).

487 The framework can also be adapted to handle data with zero and negative values. The
488 use of a Box-Cox transformation limits our approach to positive time series, as is
489 often encountered in complex seasonal time series. However, the inverse hyperbolic sine
490 transformation ([Johnson 1949](#)) can be used in its place should the need arise.

491 The derivation of likelihood estimates of the proposed approach relies on the assumption of
492 a Gaussian distribution for the errors, something that is often a reasonable approximation
493 when the level of the process is sufficiently far from the origin (Hyndman et al. 2008).
494 In cases where such an assumption may conflict with the underlying structure of the data
495 generating process, our approach can be readily adapted to non-Gaussian situations. Being
496 based on exponential smoothing where the conditioning process ensures that successive
497 state vectors become fixed quantities, any suitable distribution can substitute for the role
498 of the Gaussian distribution. Thus, if the innovations have a t -distribution, the prediction
499 error form of the likelihood can be formed directly from the product of t -distributions. The
500 analytical form of successive prediction distributions is unknown, but they can be simulated
501 from successive t -distributions using means obtained from the application of the equations
502 of the innovations state space model. This can be contrasted with a Kalman filter approach,
503 which must usually be adapted in the presence of non-Gaussian distributions, to a form
504 which necessitates the use of computationally intensive simulation methods.

505 The proposed frameworks can also be extended to exploit the potential inter-series dependen-
506 cies of a set of related time series, providing an alternative to the existing vector exponential
507 smoothing framework (De Silva et al. 2007), but with several advantages (see De Livera
508 2010a, Chap 7).

509 The R code for the methods implemented in this paper will be available in the *forecast*
510 package for R (Hyndman 2010).

511 References

- 512 Akram, M., Hyndman, R. & Ord, J. (2009), 'Exponential smoothing and non-negative data',
513 *Australian and New Zealand Journal of Statistics* **51**(4), 415–432.
- 514 Anderson, B. D. O. & Moore, J. B. (1979), *Optimal filtering*, Prentice-Hall, Englewood
515 Cliffs.
- 516 Ansley, C. F. & Kohn, R. (1985), 'Estimation, filtering, and smoothing in state space models
517 with incompletely specified initial conditions', *The Annals of Statistics* **13**, 1286–1316.
- 518 Archibald, B. C. & Koehler, A. B. (2003), 'Normalization of seasonal factors in Winters'
519 methods', *International Journal of Forecasting* **19**, 143–148.
- 520 Billah, B., Hyndman, R. J. & Koehler, A. B. (2005), 'Empirical information criteria for time
521 series forecasting model selection', *Journal of Statistical Computation and Simulation*
522 **75**, 831–840.
- 523 Box, G. E. P. & Cox, D. R. (1964), 'An analysis of transformations', *Journal of the Royal*
524 *Statistical Society, Series B* **26**(2), 211–252.
- 525 Box, G. E. P. & Jenkins, G. M. (1970), *Time series analysis: forecasting and control*, 1st
526 edn, Holden-Day, San Francisco.
- 527 Burnham, K. P. & Anderson, D. R. (2002), *Model selection and multimodel inference: a*
528 *practical information-theoretic approach*, 2nd edn, Springer-Verlag.
- 529 Chatfield, C. (1978), 'The Holt-Winters forecasting procedures', *Applied Statistics* **27**, 264–
530 279.
- 531 Chatfield, C. & Yar, M. (1991), 'Prediction intervals for multiplicative Holt-Winters',
532 *International Journal of Forecasting* **7**, 31–37.
- 533 De Jong, P. (1991), 'The diffuse Kalman filter', *The Annals of Statistics* **19**, 1073–1083.

- 534 De Livera, A. M. (2010a), Automatic forecasting with a modified exponential smoothing
535 state space framework, Working paper 10/10, Department of Econometrics & Business
536 Statistics, Monash University.
- 537 De Livera, A. M. (2010b), Forecasting time series with complex seasonal patterns using
538 exponential smoothing, PhD thesis, Monash University.
- 539 De Silva, A., Hyndman, R. & Snyder, R. (2007), ‘The vector innovation structural time
540 series framework: a simple approach to multivariate forecasting’.
- 541 Gardner, Jr, E. S. (1985), ‘Exponential smoothing: The state of the art’, *Journal of Forecast-*
542 *ing* **4**, 1–28.
- 543 Gardner, Jr, E. S. & McKenzie, E. (1985), ‘Forecasting trends in time series’, *Management*
544 *Science* **31**(10), 1237–1246.
- 545 Gilchrist, W. (1976), *Statistical Forecasting*, Wiley: Chichester.
- 546 Gould, P. G., Koehler, A. B., Ord, J. K., Snyder, R. D., Hyndman, R. J. & Vahid-Araghi, F.
547 (2008), ‘Forecasting time series with multiple seasonal patterns’, *European Journal of*
548 *Operational Research* **191**(1), 207–222.
- 549 Hannan, E. J., Terrel, R. & Tuckwell, N. (1970), ‘The seasonal adjustment of economic time
550 series’, *International Economic Review* **11**, 24–52.
- 551 Harrison, P. J. & Stevens, C. F. (1976), ‘Bayesian forecasting’, *Journal of the Royal*
552 *Statistical Society, Series B* **38**(3), 205–247.
- 553 Harvey, A. (1989), *Forecasting structural time series models and the Kalman filter*, Cam-
554 bridge University Press.
- 555 Harvey, A. & Koopman, S. J. (1993), ‘Forecasting hourly electricity demand using time-
556 varying splines.’, *Journal of the American Statistical Association* **88**, 1228–1236.

- 557 Harvey, A., Koopman, S. J. & Riani, M. (1997), ‘The modeling and seasonal adjustment of
558 weekly observations’, *Journal of Business and Economic Statistics* **15**, 354–368.
- 559 Hyndman, R. J. (2010), *Forecasting functions for time series*. R version 2.09.
560 **URL:** <http://robjhyndman.com/software/forecast/>
- 561 Hyndman, R. J., Akram, M. & Archibald, B. C. (2007), ‘The admissible parameter space
562 for exponential smoothing models’, *Annals of the Institute of Statistical Mathematics*
563 **60**, 407–426.
- 564 Hyndman, R. J., Koehler, A. B., Ord, J. K. & Snyder, R. D. (2005), ‘Prediction intervals
565 for exponential smoothing using two new classes of state space models’, *Journal of*
566 *Forecasting* **24**, 17–37.
- 567 Hyndman, R. J., Koehler, A. B., Ord, J. K. & Snyder, R. D. (2008), *Forecasting with*
568 *exponential smoothing: the state space approach*, Springer-Verlag, Berlin.
569 **URL:** www.exponentialsMOOTHING.net
- 570 Hyndman, R. J., Koehler, A. B., Snyder, R. D. & Grose, S. (2002), ‘A state space framework
571 for automatic forecasting using exponential smoothing methods’, *International Journal of*
572 *Forecasting* **18**(3), 439–454.
- 573 Hyndman, R. & Khandakar, Y. (2008), ‘Automatic time series forecasting: The forecast
574 package for R’, *Journal of Statistical Software* **26**(3).
- 575 Johnson, N. L. (1949), ‘Systems of frequency curves generated by methods of translation’,
576 *Biometrika* **36**, 149–176.
- 577 Lin, J. & Liu, T. (2002), ‘Modeling lunar calendar holiday effects in Taiwan’, *Taiwan*
578 *Economic Forecast and Policy* **33**(1), 1–37.
- 579 Makridakis, S., Anderson, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton,
580 J., Parzen, E. & Winkler, R. (1982), ‘The accuracy of extrapolation (time series) methods:
581 results of a forecasting competition’, *International Journal of Forecasting* **1**, 111–153.

- 582 Makridakis, S. & Hibon, M. (2000), 'The M3-competition: Results, conclusions and
583 implications', *International Journal of Forecasting* **16**, 451–476.
- 584 Makridakis, S., Wheelwright, S. C. & Hyndman, R. J. (1998), *Forecasting: methods and
585 applications*, 3rd edn, John Wiley & Sons, New York.
- 586 **URL:** www.robjhyndman.com/forecasting/
- 587 McKenzie, E. (1984), 'General exponential smoothing and the equivalent ARMA process',
588 *Journal of Forecasting* **3**, 333–344.
- 589 Ord, J. K., Koehler, A. B. & Snyder, R. D. (1997), 'Estimation and prediction for a class
590 of dynamic nonlinear statistical models', *Journal of the American Statistical Association*
591 **92**, 1621–1629.
- 592 Pankratz, A. & Dudley, U. (1987), 'Forecasts of power-transformed series', *Journal of*
593 *Forecasting* **6**, 239–248.
- 594 Pedregal, D. J. & Young, P. C. (2006), 'Modulated cycles, an approach to modelling
595 periodic components from rapidly sampled data', *International Journal of Forecasting*
596 **22**, 181–194.
- 597 Proietti, T. (2000), 'Comparing seasonal components for structural time series models',
598 *International Journal of Forecasting* **247-260**, 16.
- 599 Proietti, T. & Riani, M. (2009), 'Transformations and seasonal adjustment', *Journal of Time*
600 *Series Analysis* **30**, 47–69.
- 601 Reid, D. J. (1975), A review of short-term projection techniques, in H. A. Gordon, ed.,
602 'Practical aspects of forecasting', Operational Research Society: London, p. 825.
- 603 Riazuddin, R. & Khan, M. (2005), 'Detection and forecasting of Islamic calendar effects in
604 time series data', *State Bank of Pakistan Research Bulletin* **1(1)**, 25–34.
- 605 Snyder, R. (2006), 'Discussion', *International Journal of Forecasting* **22(4)**, 673 – 676.

- 606 Snyder, R. D. (1985a), Estimation of dynamic linear model—another approach, Technical
607 report, Department of Econometrics and Operations Research, Monash University.
- 608 Snyder, R. D. (1985b), ‘Recursive estimation of dynamic linear models’, *Journal of the*
609 *Royal Statistical Society, Series B* **47**(2), 272–276.
- 610 Snyder, R., Koehler, A. & Ord, J. (2002), ‘Forecasting for inventory control with exponential
611 smoothing’, *International Journal of Forecasting* **18**(1), 5–18.
- 612 Taylor, J. (2008), ‘An evaluation of methods for very short-term load forecasting using
613 minute-by-minute British data’, *International Journal of Forecasting* **24**(4), 645–658.
- 614 Taylor, J. (2010), ‘Triple seasonal methods for short-term electricity demand forecasting’,
615 *European Journal of Operational Research* **204**, 139–152.
- 616 Taylor, J. W. (2003), ‘Short-term electricity demand forecasting using double seasonal
617 exponential smoothing’, *Journal of the Operational Research Society* **54**, 799–805.
- 618 Taylor, J. W. & Snyder, R. D. (2009), Forecasting intraday time series with multiple seasonal
619 cycles using parsimonious seasonal exponential smoothing, Technical Report 09/09,
620 Department of Econometrics and Business Statistics, Monash University.
- 621 Weinberg, J., Brown, L. & Stroud, J. (2007), ‘Bayesian forecasting of an inhomogeneous
622 Poisson process with applications to call center data’, *Journal of the American Statistical*
623 *Association* **102**(480), 1185–1198.
- 624 West, M. & Harrison, J. (1997), *Bayesian forecasting and dynamic models*, 2nd edn,
625 Springer-Verlag, New York.