Comparison of classical and Bayesian approaches for intervention analysis in structural models

Thiago R. Santos[a] Glaura C. Franco[a] Dani Gamerman[b]

[a]*Department of Statistics, Federal University of Minas Gerais, Av. Antonio Carlos 6627, Belo Horizonte, MG 31270-901, Brazil*

[b]*Department of Statistical Methods, Federal University of Rio de Janeiro, Caixa Postal 68530, Rio de Janeiro, RJ 21945-970, Brazil*

E-mails: thiagors@ufmg.br, glauraf@ufmg.br, dani@im.ufrj.br

Abstract

This paper provides comparisons of the classical and Bayesian approaches to estimation and prediction on intervention models. Intervention analysis has been recently the subject of several studies, mainly because real time series present a wide variety of phenomena that are caused by external and/or unexpected events. In this work, transfer functions are used to model different forms of intervention to the mean level of a time series. This is performed into the framework of state-space or structural models. Two canonical forms of intervention are considered: pulse and step functions. The former affects temporarily the level of the series and the latter affects permanently the level of the series. Also, the models considered allow for static and dynamic explanation of the intervention effects. Classical inference for these models is introduced and comparison between the two approaches, classical and Bayesian, is performed through Monte Carlo simulation. Bootstrap and MCMC methods are used for approximate classical and Bayesian inferences, respectively. Results are compared in terms of point and interval estimation. Point estimation shows that the maximum likelihood and the posterior mode estimators perform better than posterior mean and median. Interval estimation shows that Bayesian credibility intervals perform better than the respective classical confidence intervals. The methodology was applied to two Brazilian economic indexes and showed satisfactory results for the intervention models used.

Keywords: state space models, dynamic linear models, prediction, transfer function, bootstrap, MCMC.

1 Introduction

Time series are frequently affected by external events, known as interventions, that can change the structure of the series (such as strikes, policy changes, etc.). The first proposals of intervention analysis seem to have arisen in the Social Sciences with the work of Campbell & Stanley (1968), but the term intervention was first introduced by Glass (1972). However, it was only in 1975 that Box and Tiao developed the theory of intervention analysis to study structural changes in time series (Box & Tiao, 1975).

At the same time, Box & Jenkins (1976) introduced the transfer function (TF) models in the context of the Autoregressive Integrated Moving Average (ARIMA) process. Transfer function models were designed to measure the relationship between an output series and one or more input series. For example, in the case of an output series y_t and an input series x_t, the transfer function relates the variables through a linear filter of the form

$$y_t = \vartheta(B)x_t + \epsilon_t$$

where $\vartheta(B) = \sum_{j=-\infty}^{\infty} \vartheta_j B^j$, B is the backshift operator $B^k y_t = y_{t-k}$ and the error series ϵ_t is the possibly time-correlated. The coefficients ϑ_j in the transfer function model are called impulse response function.

If the input series x_t is a deterministic function that accounts for the structural changes in the output series y_t, TF can be used to model structural breaks following the same idea of Box & Tiao (1975).

There exists a wide variety of approaches to model a time series. Structural models, developed by Harvey (1989) and West & Harrison (1997), are a rich class of models formulated directly in terms of non-observable components, such as trend, seasonality, cycles and noise. The model is generally written in a state-space representation, to allow the use of the Kalman filter (Kalman, 1960) to estimation and prediction (see more in Durbin & Koopman (2001)). In the Bayesian context, these models are also known as dynamic linear models.

Due to the fact that ARIMA models can be written in the state space form, adapting the idea of transfer functions to structural models is quite natural. The flexibility of the state space representation allows the insertion of covariates in the observation or state equations. Some work on this subject include Penzer (2007), de Jong & Penzer (1998), Harvey & Koopman (1992), Salvador & Gargallo (2004), Ravines et al. (2008) and Alves et al. (2009).

In this work, two kinds of intervention procedures are described: pulse and step functions. These are built in addition to the Local Level Model (LLM), the simplest structural model for the trend component. Also, the gain factor associated with the intervention is either fixed in time or is allowed to vary in time. Inference about the parameters of the TF and the variances of the errors in the observation and state equations are performed using classical or Bayesian approaches. Classical inference for these time-varying intervention models was not done before, to the knowledge of the authors.

Confidence intervals under the classical paradigm can be built using the bootstrap (Efron, 1979). The residuals of the fitted model are used here to generate the bootstrap series, under a parametric approach. Previous application of the bootstrap in structural models include Stoffer & Wall (1991), Franco & Souza (2002), Pfeffermann & Tiller (2004) and Franco et al. (2008).

The Bayesian approach for these models does not lead to analytically tractable posterior distributions. Thus, approximate methods are required and Markov chain Monte Carlo (MCMC) methods are employed (Reis et al., 2006). Point estimates are obtained and credibility intervals are also built. References about MCMC in dynamic linear models include Carter & Kohn (1994), Fruhwirth-Schnatter (1994), Lopes & Moreira (1999), Schmidt et al. (1999) and Santos & Franco (2008).

The objective of this paper is to compare the efficiency of the classical and Bayesian paradigms for inference about the parameters of the intervention models considered. Also, the efficiency of approximating methods is empirically evaluated.

Monte Carlo experiments are conducted for this purpose in a variety of settings and the classical and Bayesian estimates are compared through the bias and mean square error (MSE). Confidence and credibility intervals are built for the parameters and the one-step ahead forecasts, and they are compared using the width and the coverage rate.

Applications of the methodologies described are performed on two real monthly time series of Brazilian indexes, an inflation index for the city of Belo Horizonte and the São Paulo stock market index.

This paper is organized as follows. Section 2 shows how the intervention analysis can be considered in structural models using transfer functions. Section 3 presents the estimation procedures considered, along with confidence and credibility intervals. Monte Carlo simulations are performed in Section 4 and Section 5 presents two applications on real time series. Section 6 concludes the work.

2 Intervention analysis in structural models

The state space model for a univariate time series $y_t, t = 1, ..., n$ is defined by the equations

$$y_t = \boldsymbol{Z}'_t \boldsymbol{\alpha}_t + d_t + \epsilon_t \quad (observation \quad equation) \tag{1}$$

$$\boldsymbol{\alpha}_t = \boldsymbol{T}_t \boldsymbol{\alpha}_{t-1} + \boldsymbol{c}_t + \boldsymbol{R}_t \boldsymbol{\omega}_t, \quad (state \quad equation) \tag{2}$$

where $\boldsymbol{\alpha}_t$ is a $m \times 1$ vector of unobserved state variables, d_t and $\boldsymbol{c}_t$ are the effect of exogenous covariates and $\mathbf{Z_t}$, $\boldsymbol{T}_t$ and $\boldsymbol{R}_t$ are system matrices. The model becomes non-linear when some of these matrices involve unknown quantities. In this case, calculations become more complicated as updating equations for inference in known analytic form no longer exist. The terms ϵ_t and $\boldsymbol{\omega}_t$ represent zero-mean random processes, independent and identically distributed with variances h_t and $\boldsymbol{Q}_t$, respectively. Furthermore, the disturbances are serially uncorrelated.

The simplest structural model to describe series that present only the trend component is the local level model (LLM), defined by

$$y_t = \mu_t + \varepsilon_t, \quad \varepsilon_t \sim (0, \sigma_\varepsilon^2),$$

$$\mu_t = \mu_{t-1} + \eta_t, \quad \eta_t \sim (0, \sigma_\eta^2).$$

In this case, $\alpha_t = \mu_t$, $d_t = c_t = 0$, $\omega_t = \eta_t$, $\epsilon_t = \varepsilon_t$, $h_t = \sigma_\varepsilon^2$, $Q_t = \sigma_\eta^2$ and $Z'_t = T_t = R_t = 1$.

If a change of level is observed in the series, its effect can be modelled inserting a component E_t in the state equation. Let the component E_t be written as the transfer function filter of Box et al. (1994),

$$E_t = \vartheta(B) x_t \tag{3}$$

where x_t is an exogenous variable. It will be assumed that $\vartheta(B) = \beta(B)B^b/\rho(B)$ where $\beta(B) = \beta_0 - \beta_1 B - ... - \beta_s B^s$, $\rho(B) = \rho_0 - \rho_1 B - ... - \rho_r B^r$ and b is a delay parameter. This model will be denoted here by TF(r, b, s). More details can be found in Wei (1991).

In this work, the exogenous variable x_t in (3) will represent intervention effects. In practice, there are many possibilities for the occurrence of interventions. For example, the impact of an external event can be felt b periods after the intervention, with an effect only at that moment. In this case, the transfer function model is a TF(0,0,b):

$$E_t = \beta_0 B^b x_t = \beta_0 x_{t-b}.$$

Or it can happen that the impact is felt on the time of the intervention, but the response is gradual. Thus, a suitable model is a TF(1,0,0),

$$E_t = \frac{\beta_0 B^0}{\rho_0 - \rho_1 B} x_t \Rightarrow (\rho_0 - \rho_1 B)E_t = \beta_0 x_t \Rightarrow E_t = \rho E_{t-1} + \beta x_t. \tag{4}$$

where $\rho = \rho_1/\rho_0$ and $\beta = \beta_0/\rho_0$.

Box & Tiao (1975) define two common types of intervention variables, step and pulse functions, that can be represented by dummy variables as follows:

(1) Step function: If the intervention takes place at some fixed time T and remains in effect thereafter:

$$S_t^T = \begin{cases} 0, & t < T \\ 1, & t \geq T \end{cases} ;$$

(2) Pulse function: If the intervention takes place at some fixed time T and has an effect only in that period:

$$I_t^T = \begin{cases} 1, & t = T \\ 0, & t \neq T \end{cases} .$$

Figure 1 presents the behavior of the structural block E_t for a series of size $n = 100$ with a change of level at time $T = 50$, under Model TF(1,0,0), for pulse and step functions. For the pulse function it can be seen that, increasing the value of ρ causes a very slow return of the series to the mean level presented before the impact. For the step function, if ρ increases the series takes more time to attain a

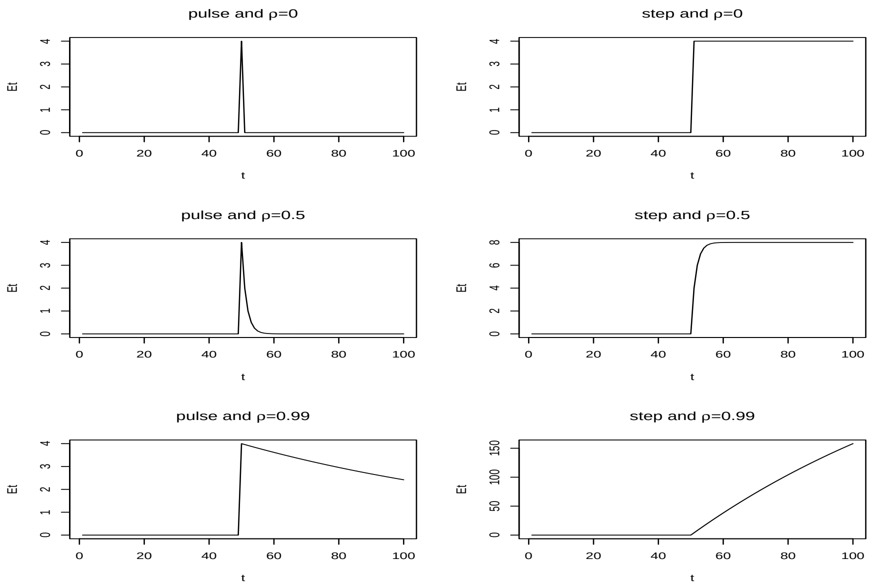

Figure 1: Intervention effect of pulse and step functions.

new mean level.

In this work, Model TF(1,0,0) will be considered with two kinds of impulse response function. In the first one (Model 1) the gain factor β in equation (4) is fixed in time, while Model 2 allows β to vary in time. This later model is known as a transfer function model with dynamic gain factor (see Alves et al. (2009)).

The local level model (LLM), is considered along with an intervention component. The structure of Models 1 and 2 under the local level approach are described below, with examples of series generated from these models.

Model 1

The local level model with an intervention component assuming the β coefficient fixed in time has the following structure:

$$\begin{cases} y_t = \mu_t + E_t + \varepsilon_t, & \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \\ \mu_t = \mu_{t-1} + \eta_t, & \eta_t \sim N(0, \sigma_\eta^2) \\ E_t = \rho E_{t-1} + \beta X_t \end{cases}, \tag{5}$$

where $t = 1, 2, \ldots, n$, ε_t and η_t are independent and $0 \leq \rho \leq 1$. Thus, the impulse response function has a geometric decay. For this reason, ρ is usually called persistence or memory of the effect. Note also that it induces a non-linearity in the model. The observation mean response is given by $\theta_t = \mu_t + E_t$.

The model in (5) can be written in the state space form with matrices given by

$$\boldsymbol{Z}_t' = \begin{bmatrix} 1 & 1 \end{bmatrix}, \ \boldsymbol{R}_t = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \ \boldsymbol{T}_t = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}, \ \boldsymbol{\omega}_t = \begin{bmatrix} \eta_t \\ 0 \end{bmatrix}, \ h_t = [\sigma_\varepsilon^2],$$

$$\boldsymbol{Q}_t = \begin{bmatrix} \sigma_\eta^2 & 0 \\ 0 & 0 \end{bmatrix}, \ d_t = 0, \ \boldsymbol{c}_t = \begin{bmatrix} 0 \\ \beta X_t \end{bmatrix} \text{ and } \boldsymbol{\alpha}_t = \begin{bmatrix} \mu_t \\ E_t \end{bmatrix}.$$

Figure 2 shows some simulated series under Model 1. As expected, the behavior of the series is very similar to the behavior of block E_t showed in Figure 1. For the pulse response, it can be seen that small values of ρ causes in the series an effect similar to the presence of an outlier at the point of intervention. For large values of ρ, the series presents a jump at the point of the intervention, but with a gradual return to the mean value. For the step response, when $\rho = 0.00$ or 0.50, there is a sudden jump at the time of intervention, with a change of level and the series remaining at the new level. When $\rho = 0.99$ there is a gradual change of level and the series takes a long time to attain a new level.

Model 2

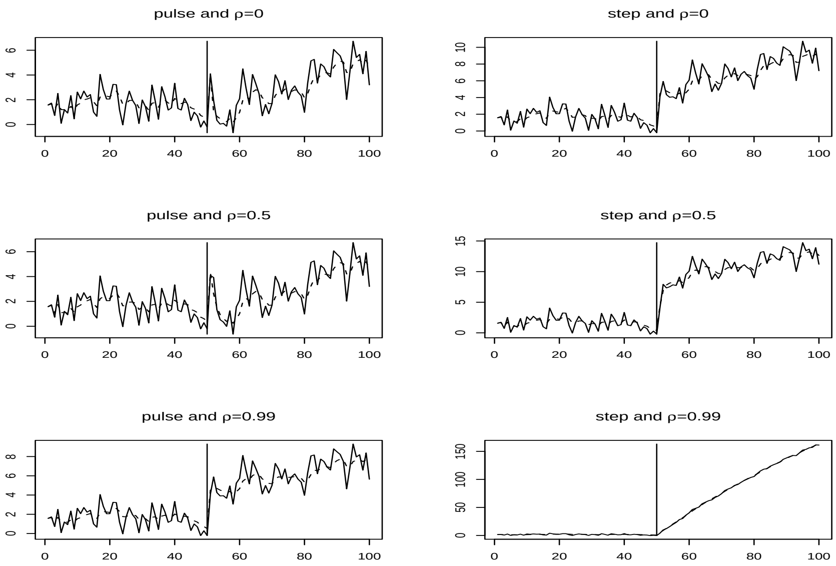

Figure 2: Examples of pulse and step functions for Model 1. In this case, $n = 100$, $\beta = 4$ and $\rho = 0.00$, 0.50 and 0.99. The vertical line indicates the time of intervention, $T = 50$, the solid line indicates the series, y_t, and the dashed line indicates the mean response θ_t.

The local level model with a dynamic gain factor (Alves et al., 2009) is given by

$$\begin{cases} y_t = \mu_t + E_t + \varepsilon_t, & \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \\ \mu_t = \mu_{t-1} + \eta_t, & \eta_t \sim N(0, \sigma_\eta^2) \\ E_t = \rho E_{t-1} + \beta_{t-1} X_t & \\ \beta_t = \beta_{t-1} + \xi_t, & \xi_t \sim N(0, \sigma_\xi^2) \end{cases},$$

where $t = 1, 2, \ldots, n$, ε_t, η_t and ξ_t are jointly independent.

Unlike Model 1, in Model 2 the parameter β_t is stochastic and is obtained dynamically through time by a stochastic law (here a random walk).

Model 2 can be written in the state space form with matrices given by

$$\mathbf{Z}_t' = [\ 1 \quad 1 \quad 0\ ],\ \boldsymbol{R}_t = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},\ \boldsymbol{T}_t = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \rho & X_t \\ 0 & 0 & 1 \end{bmatrix},\ \boldsymbol{\omega}_t = \begin{bmatrix} \eta_t \\ 0 \\ \xi_t \end{bmatrix},\ h_t =$$

$$[\sigma_\varepsilon^2],\ \boldsymbol{Q}_t = \begin{bmatrix} \sigma_\eta^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sigma_\xi^2 \end{bmatrix},\ d_t = 0,\ \boldsymbol{c}_t = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \text{ and } \boldsymbol{\alpha}_t = \begin{bmatrix} \mu_t \\ E_t \\ \beta_t \end{bmatrix}.$$

Figure 3 shows some simulated series under Model 2. The behavior is very similar to Model 1, but in this case the dynamic β causes more noise in the series. If a pulse function is used, the estimation of the β_t's is based entirely on one observation, and this compromises their estimation. For this reason, only the step function will be used in Model 2.

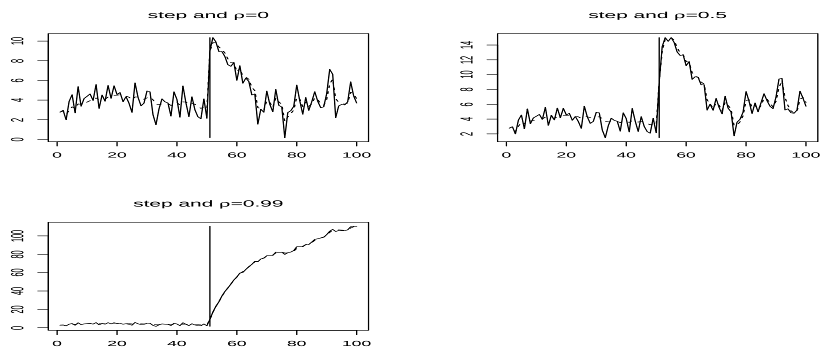

Figure 3: Examples of step functions for Model 2. In this case, $n = 100$ and $\rho = 0.00$, 0.50 and 0.99. The vertical line indicates the time of intervention, $T = 50$, the solid line indicates the series, y_t, and the dashed line indicates the mean response θ_t.

3 Estimation and prediction procedures

The parameters in Models 1 and 2 can be estimated using classical methods (i.e, maximizing the likelihood function) or Bayesian methods. Parameters can be divided into two groups: state parameters α_t and hyperparameters, denoted by $\boldsymbol{\varphi}$. The components of $\boldsymbol{\varphi}$ are $(\sigma_\varepsilon^2, \sigma_\eta^2, \rho, \beta)$ for Model 1 and $(\sigma_\varepsilon^2, \sigma_\eta^2, \sigma_\xi^2, \rho)$ for Model 2.

In both cases, the Kalman filter (Kalman, 1960) algorithm will be used to estimate the state components, α_t, given the observations $\boldsymbol{Y}_t = \{y_1, ..., y_t\}$ and $\boldsymbol{\varphi}$. Denoting the linear estimator and its variance matrix by

$$\boldsymbol{a}_t = E(\boldsymbol{\alpha}_t | Y_t) \quad \text{and} \quad \boldsymbol{V}_t = Var(\boldsymbol{\alpha}_t | Y_t),$$

respectively, the Kalman filter can be used to calculate the one-step ahead forecast error, ν_t, and its variance

$$\nu_t = y_t - E(y_t | Y_{t-1}) = y_t - \widetilde{y}_{t|t-1} \quad \text{and} \quad F_t = \boldsymbol{Z}_t \boldsymbol{V}_{t|t-1} \boldsymbol{Z}_t^{'} + h_t,$$

where $\boldsymbol{V}_{t|t-1} = Var(\boldsymbol{\alpha}_t | Y_{t-1})$. Note that both sequences of ν_t's and F_t's depend on $\boldsymbol{\varphi}$.

Assuming that the disturbances ϵ_t and ω_t are normally distributed, the likelihood function $L(\boldsymbol{\varphi}; \boldsymbol{Y}_n)$ can be computed after integrating out state variables. For a univariate time series of size n, the logarithm of the likelihood function is given by

$$\ln L(\boldsymbol{\varphi}; \boldsymbol{Y}_n) = \ln \prod_{t=1}^{n} p(y_t | \boldsymbol{Y}_{t-1}, \boldsymbol{\varphi}) = -\frac{n}{2} \ln(2\pi) - \frac{1}{2} \sum_{t=1}^{n} \ln |F_t| - \frac{1}{2} \sum_{t=1}^{n} \nu_t' F_t^{-1} \nu_t. \quad (6)$$

The range of possible values considered for the components of $\boldsymbol{\varphi}$ are $\Re^+$ for variances, $[0, 1]$ for ρ and $\Re$ for β.

In the next sections, the classical and Bayesian estimation procedures are detailed.

3.1 Classical inference

Maximizing the logarithm of the likelihood function, given in equation (6), with respect to $\boldsymbol{\varphi}$ yields the maximum likelihood estimator of the parameters. As the likelihood is a nonlinear function of $\boldsymbol{\varphi}$, numerical procedures should be used. In this work, the well-known BFGS optimization algorithm is employed (see Franco et al. (2008) for details).

Inferences for the parameters under the classical approach will be done using the bootstrap (Efron (1979)). Bootstrap methods are resampling techniques designed to approximate the probability distribution function of the data by an empirical function of a finite sample. Their use in time series must be performed with caution because the observations are not independent and the time series structure may be lost in a careless resampling. Thus, the time series must be resampled indirectly.

The most common way of performing the bootstrap in time series is resampling the residuals of the fitted model. In this work, a parametric bootstrap method in the residuals is used. The method consists in the following. Initially, the parameters of the model are estimated based on the original time series, y_t. As the noises in the observation and state equations follow the normal distribution with zero mean and variances h_t and $\boldsymbol{Q}_t$, respectively, the parametric bootstrap is easily performed replacing the variances by their estimated values and sampling observations from this distribution, thus obtaining the bootstrap residuals, $\boldsymbol{\epsilon}_t^*$ and $\boldsymbol{\omega}_t^*$. The bootstrap series y_t^*, $t = 1, \ldots, n$ is calculated using the bootstrap residuals as

$$
\begin{aligned}
y_t^* &= \boldsymbol{Z}_t^{'}\boldsymbol{\alpha}_t^* + d_t + \varepsilon_t^* \\
\boldsymbol{\alpha}_t^* &= \boldsymbol{T}_t\boldsymbol{\alpha}_{t-1}^* + \boldsymbol{c}_t + \boldsymbol{R}_t\boldsymbol{\omega}_t^*.
\end{aligned}
$$

The percentile bootstrap confidence interval (Efron & Tibshirani, 1993) will be employed to build intervals for the parameters. In this case, R independent bootstrap samples $y_t^{*1}, y_t^{*2}, ..., y_t^{*R}$ are generated and the parameter vector $\hat{\boldsymbol{\varphi}}^*$ is estimated for each bootstrap series. The lower and upper bounds of the $100(1-\kappa)\%$ percentile bootstrap confidence interval will be given by

$$
[\hat{\varphi}_i^{*(\kappa/2)} \quad ; \quad \hat{\varphi}_i^{*(1-\kappa/2)}]
$$

in which $\hat{\varphi}_i^{*(\kappa)}$ is the $R.(\kappa)th$ ordered value of the bootstrap replication for each component φ_i of vector $\boldsymbol{\varphi}$, $i = 1, ..., p$.

3.2 Bayesian inference

Using the Bayesian approach, initially it requires the specification of a prior distribution of the parameter vector, $\boldsymbol{\varphi}$, which is denoted by $\pi(\boldsymbol{\varphi})$. The prior distributions obtained from $\pi(\boldsymbol{\varphi})$ should reflect the previous knowledge about the parameters and, most of the times, this is not feasible or very difficult to obtain. Whereas the state parameters are easily understood, the hyperparameters refer to characteristic of the models that are not readily understandable by users. Thus, lack of prior information seems to be the norm for them. In these cases, it seems advisable to use non-informative or reference prior distributions. West & Harrison (1997) details the use of reference prior distributions but only for the state parameters. Gomes (2006) discusses the specification of reference priors for the hyperparameters in a few, specific structural models but no general results are available to the knowledge of the authors. A simpler alternative is the uniform prior, given by $\pi(\boldsymbol{\varphi}) \propto c$, for all possible values of $\boldsymbol{\varphi}$, and 0, otherwise.

The likelihood function $L(\boldsymbol{\varphi}; \boldsymbol{Y}_n)$ that summarizes the sample information about $\boldsymbol{\varphi}$, is used to update the prior distribution, thus generating the posterior distribution of $\boldsymbol{\varphi}$. It is obtained via the Bayes' theorem as follows

$$\pi(\boldsymbol{\varphi} \mid \boldsymbol{Y}_n) = \frac{L(\boldsymbol{\varphi}; \boldsymbol{Y}_n)\pi(\boldsymbol{\varphi})}{\int L(\boldsymbol{\lambda}; \boldsymbol{Y}_n)\pi(\boldsymbol{\lambda})d\boldsymbol{\lambda}}.$$

The posterior distribution reflects all the uncertainty about $\boldsymbol{\varphi}$ after the data has been observed.

The Bayes estimators considered here are the posterior mean, obtained when the quadratic loss function is used, the posterior median, obtained when the absolute loss function is used and the posterior mode, obtained when the 0-1 loss function is used (Migon & Gamerman, 1999).

As in this case the posterior distribution does not have a closed form, numerical methods such as MCMC should be used to obtain the estimates of the parameters (see Gamerman & Lopes (2006) for details).

In this work, a hybrid version of the Metropolis-Hastings algorithm, a known MCMC method, is adopted. In this version, the components of $\boldsymbol{\varphi}$ are updated separately, with different proposed densities. This approach considers the marginal posterior distribution of the hyperparameters once the state parameters are integrated out. This is equivalent to sampling all parameters jointly, since the full conditional for the state parameters given the hyperparameters is analytically available. Reis et al. (2006) provide substantial empirical evidence in favor of this joint sampling scheme of all model parameters against other blocking schemes. There are several possibilities of candidate-generating densities in the literature. Here, a candidate-generating density proposed by Chib & Greenberg (1995), given by a random walk centered in the last value of the parameter is used.

Credibility intervals for $\varphi_i,\ i = 1, ..., p$ are built as follows. Given a value $0 < \kappa < 1$, any interval (t_1, t_2) satisfying

$$\int_{t_1}^{t_2} \pi(\varphi_i \mid \boldsymbol{Y}_n)\ d\varphi_i = 1 - \kappa$$

is a credibility interval for φ_i with level $100(1-\kappa)\%$. These intervals are approximated by the corresponding order statistics of the MCMC generated sample from $\varphi_i, i = 1, ..., p$.

3.3 Predictions

An important issue in time series analysis is to predict future values, based on the fitted model. The prediction of a future observation y_{n+k} based on $\boldsymbol{Y}_n$ can be obtained by the combination of the observation equation (1) at time $n+k$ with the k-step ahead state equation given by

$$\begin{aligned} \boldsymbol{\alpha}_{n+k} &= \left(\left(\prod_{i=1}^{k} \boldsymbol{T}_{n+i}\right) \boldsymbol{\alpha}_n + \sum_{i=1}^{k} \left(\prod_{j=1}^{k-i} \boldsymbol{T}_{n+k-j+1}\right) \boldsymbol{c}_{n+i}\right) \\ &+ \sum_{i=1}^{k} \left(\prod_{j=1}^{k-i} \boldsymbol{T}'_{n+k-j+1}\right) \boldsymbol{R}'_{n+i}\boldsymbol{\omega}_{n+i}, \end{aligned} \tag{7}$$

where $\prod_{i=1}^{0} A_i = I$, for any matrices $A_1, A_2, ...$

In order to obtain k-step ahead predictions, consider the predictive distribution of $(y_{n+k}|\boldsymbol{Y}_n, \boldsymbol{\varphi})$. For the structural model (1)-(2), its mean and variance are given by

$$E(y_{n+k}|\boldsymbol{Y}_n, \boldsymbol{\varphi}) = \boldsymbol{Z}'_{n+k}\left(\left(\prod_{i=1}^{k}\boldsymbol{T}_{n+i}\right)\boldsymbol{a}_n + \sum_{i=1}^{k}\left(\prod_{j=1}^{k-i}\boldsymbol{T}_{n+k-j+1}\right)\boldsymbol{c}_{n+i}\right) + d_{n+k} \quad (8)$$

and

$$\begin{aligned} V(y_{n+k}|\boldsymbol{Y}_n, \boldsymbol{\varphi}) = \boldsymbol{Z}'_{n+k}\left(\textstyle\prod_{i=1}^{k} T'_{n+i}\right) V_n \left(\textstyle\prod_{i=1}^{k} T_{n+i}\right)\boldsymbol{Z}_{n+k} + \\ \boldsymbol{Z}'_{n+k}\textstyle\sum_{i=1}^{k}\left(\prod_{j=1}^{k-i}\boldsymbol{T}'_{n+k-j+1}\right)\boldsymbol{R}'_{n+i}\boldsymbol{Q}_{n+i}\boldsymbol{R}_{n+i}\left(\prod_{j=1}^{k-i}\boldsymbol{T}_{n+k-j+1}\right)\boldsymbol{Z}_{n+k} + h_{n+k}. \end{aligned} \quad (9)$$

Classical approach

The prediction function is given by $E(y_{n+k}|\boldsymbol{Y}_n, \hat{\boldsymbol{\varphi}})$, i.e., it is obtained by replacing the parameter vector $\boldsymbol{\varphi}$ by its maximum likelihood estimator $\hat{\boldsymbol{\varphi}}$ (Brockwell & Davis, 1996) and is denoted by $\tilde{y}^{(c)}_{n+k|n}$. Similarly, the mean square error (MSE) of the predictions are obtained from (9) after replacing $\boldsymbol{\varphi}$ by $\hat{\boldsymbol{\varphi}}$.

The prediction function can be obtained from (8) for Models 1 and 2. They are respectively given by

$$\tilde{y}^{(c)}_{n+k|n} = a_n^{(\mu)} + \hat{\rho}^k a_n^{(E)} + \hat{\beta}\sum_{i=1}^{k}\hat{\rho}^{k-i}x_{n+i}$$

and

$$\tilde{y}^{(c)}_{n+k|n} = a_n^{(\mu)} + \hat{\rho}^k a_n^{(E)} + \sum_{i=1}^{k}\hat{\rho}^{k-i}a_n^{(\beta)}x_{n+i}$$

where $a_n^{(\mu)}$ is the component of the state vector that estimates μ_t, a_n^E the component that estimates E_t and $a_n^{(\beta)}$ the component that estimates β_n, using the maximum likelihood vector $\hat{\boldsymbol{\varphi}} = (\hat{\sigma}^2_\varepsilon, \hat{\sigma}^2_\eta, \hat{\rho}, \hat{\beta})$ for Model 1 and $\hat{\boldsymbol{\varphi}} = (\hat{\sigma}^2_\varepsilon, \hat{\sigma}^2_\eta, \hat{\sigma}^2_\xi, \hat{\rho})$ for Model 2.

A percentile bootstrap confidence interval of level $100(1-\kappa)\%$ for y_{n+k} is given by

$$[\tilde{y}^{*(\kappa/2)}_{n+k|n} \quad ; \quad \tilde{y}^{*(1-\kappa/2)}_{n+k|n}]$$

where $\tilde{y}^{*(\kappa)}_{n+k|n}$ is the $R.(\kappa)th$ ordered value of the bootstrap replication for the forecasting. The bootstrap forecast $\tilde{y}^{*}_{n+k|n}$ is calculated based on the work of Thombs & Schuncany (1990).

Bayesian approach

The prediction function is given by the mean of the predictive distribution of $y_{n+k}|\boldsymbol{Y}_n$ and is denoted by $\tilde{y}^{(b)}_{n+k|n}$. It is obtained by solving the following integral

$$\tilde{y}^{(b)}_{n+k|n} = \int E(y_{n+k}|\boldsymbol{Y}_n, \boldsymbol{\varphi})\pi(\boldsymbol{\varphi}|\boldsymbol{Y}_n)d\boldsymbol{\varphi}$$

and a $100(1-\kappa)\%$ credibility interval for y_{n+k} is given by

$$\int_{li}^{lu} p(y_{n+k} \mid \boldsymbol{Y}_n)d\boldsymbol{y} = 1-\kappa.$$

The limits li and lu can be approximately obtained by MCMC simulation and the steps of the algorithm are described below. Once a sample $\boldsymbol{\varphi}^{(\mathbf{1})}, ..., \boldsymbol{\varphi}^{(\boldsymbol{m})}$ is obtained from $\pi(\boldsymbol{\varphi}|\boldsymbol{Y}_n)$ for each j, $j = 1, ..., m$:

1. $\boldsymbol{\alpha}^{(j)}_n$ is generated from the distribution $p(\boldsymbol{\alpha_n}|\boldsymbol{\varphi}^{(j)}, \boldsymbol{Y}_n)$, obtained through Kalman filter;

2. $\boldsymbol{\alpha}^{(j)}_{n+k}$ is generated from the distribution $p(\boldsymbol{\alpha}_{n+k}|\boldsymbol{\alpha}^{(j)}_n, \boldsymbol{\varphi}^{(j)}, \boldsymbol{Y}_n)$, obtained through (7);

3. Then, $y^{(j)}_{n+k} = \boldsymbol{Z'}_{n+k}\boldsymbol{\alpha}^{(j)}_{n+k} + \epsilon^{(j)}_{n+k}$, where $\epsilon^{(j)}_{n+k}$ is generated from a Normal distribution with zero mean and variance $\sigma^{2(j)}_{\epsilon}$.

Finally, the values of $y^{(1)}_{n+k}, ..., y^{(m)}_{n+k}$ are ordered and the $\kappa/2$ and $1-\kappa/2$ quantiles are taken as the lower and upper limits of the interval, respectively.

4 Simulation results

The procedures described in the previous sections are now investigated through Monte Carlo (MC) experiments. Series with pulse and step interventions for the local level model were simulated, based on Models 1 and 2. The performances of the maximum likelihood (MLE), bootstrap (Boot) and Bayes estimators - mean (Mean), median (Med) and mode (Mode) - were evaluated for series of size $n = 100$ and $\rho = 0.00$, 0.50 and 0.99. In all cases, $\sigma^2_\epsilon = 1.00$, $\sigma^2_\eta = 0.10$, $\sigma^2_\xi = 0.50$ and $\beta = 4$. For the Bayes estimators, the uniform prior is used to allow for a fair comparison with the classical approach. Two MCMC chains with 2000 samples were generated from which the first 1000 were excluded. The number of MC and bootstrap replications were fixed at 500. The level and probability of the confidence and credibility intervals, respectively, were fixed at 0.95.

The codes, implemented by the authors in the Ox language (Doornik, 1999), are available on the site ftp://ftp.est.ufmg.br/pub/glaura/strucmod. Convergence diagnostic for the MCMC methods were based on Plummer et al. (2005).

4.1 Model 1

Figures below present the average of estimated bias and MSE over 500 replications for Model 1 with pulse and step functions.

For the pulse function (Figures 4, 5 and 6), the first conclusion that can be drawn is that ρ is always underestimated, except in the case $\rho = 0.00$. If ρ is large ($\rho = 0.99$), the best estimator is the Mode, with a very small bias and MSE, but for $\rho = 0.50$ the best estimators are the Med and Mean and for $\rho = 0.00$ the best one is the MLE. The other parameters show a very satisfactory performance for all methods. The best estimators for σ^2_η are the MLE and the Mode, regardless of the value of ρ. For σ^2_ε, all the procedures show the same behavior. In the case of parameter β, if ρ is large all the procedures present approximately the same performance, but the Mean and Med are slightly better for $\rho = 0.50$ and 0.00. With regard to the bootstrap, it can be seen that it mimics well the behavior of the MLE, thus allowing this technique to be used in the classical approach to build confidence intervals for the parameters.

For the step function (Figures 7, 8 and 9), ρ is underestimated in cases $\rho = 0.99$, with the Bayesian estimators, and $\rho = 0.50$. All the procedures show an excellent performance to estimate ρ when this parameter assumes a large value, but for $\rho = 0.50$ and 0.00, the MLE and the posterior mode are slightly better. The performances of the methods for estimating the other parameters in the step function are very similar to the pulse function, except in the case $\rho = 0.99$. It seems that in this case the MLE is not able to estimate σ_η^2, σ_ε^2 and β very accurately, presenting larger bias and MSE compared to the Bayesian estimators. Once again, the behavior of the bootstrap is very close to the MLE.

Table 1: Confidence and credibility intervals for Model 1 with pulse function.

	True	Confidence Interval			Credibility Interval		
	value	mean limits	width	coverage	mean limits	width	coverage
ρ	0.99	[0.177; 0.983]	0.806	0.87	[0.536; 0.994]	0.458	0.95
β	4.00	[2.167; 5.681]	3.514	0.91	[1.921; 5.845]	3.924	0.97
σ_ε^2	1.00	[0.687; 1.371]	0.684	0.94	[0.683; 1.427]	0.744	0.95
σ_η^2	0.10	[0.007; 0.256]	0.249	0.93	[0.043; 0.449]	0.406	0.93
y_{n+1}	*	[0.221; 4.590]	4.369	0.93	[0.013; 4.275]	4.262	0.93
ρ	0.50	[0.011; 0.757]	0.746	0.91	[0.040; 0.724]	0.684	0.89
β	4.00	[1.985; 6.144]	4.159	0.92	[1.724; 6.197]	4.473	0.96
σ_ε^2	1.00	[0.668; 1.328]	0.660	0.93	[0.700; 1.445]	0.745	0.95
σ_η^2	0.10	[0.011; 0.213]	0.202	0.87	[0.046; 0.390]	0.344	0.94
y_{n+1}	*	[-2.179; 1.960]	4.139	0.92	[-2.427; 2.024]	4.451	0.95
ρ	0.00	-	-	-	-	-	-
β	4.00	[1.944; 6.148]	4.204	0.94	[1.802; 6.231]	4.429	0.96
σ_ε^2	1.00	[0.663; 1.326]	0.663	0.93	[0.702; 1.441]	0.739	0.95
σ_η^2	0.10	[0.015; 0.221]	0.206	0.90	[0.046; 0.377]	0.331	0.94
y_{n+1}	*	[-2.149; 1.944]	4.093	0.92	[-2.424; 2.025]	4.449	0.95

* - there is not a single true value for y_{n+1} as each generated series has its own value.

Tables 1 and 2 present bootstrap and credibility intervals for Model 1 with

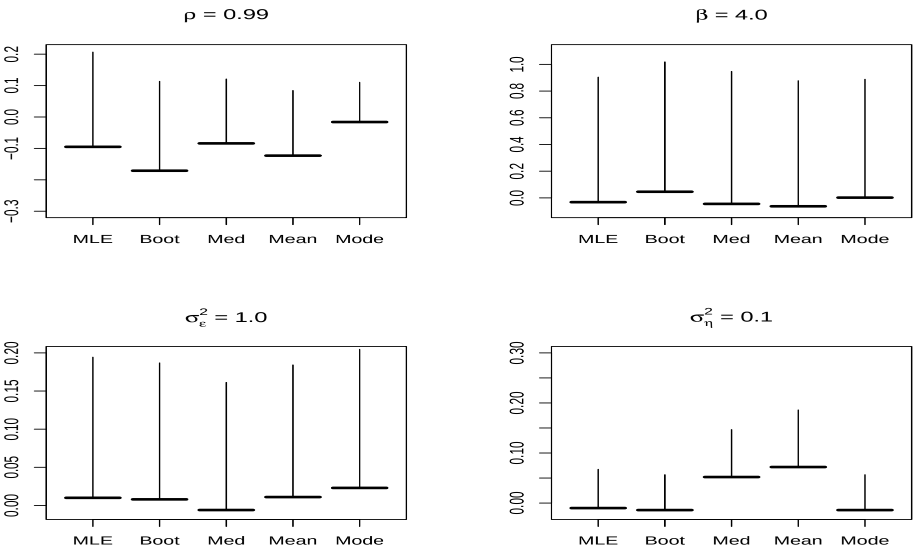

Figure 4: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with pulse function and $\rho = 0.99$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

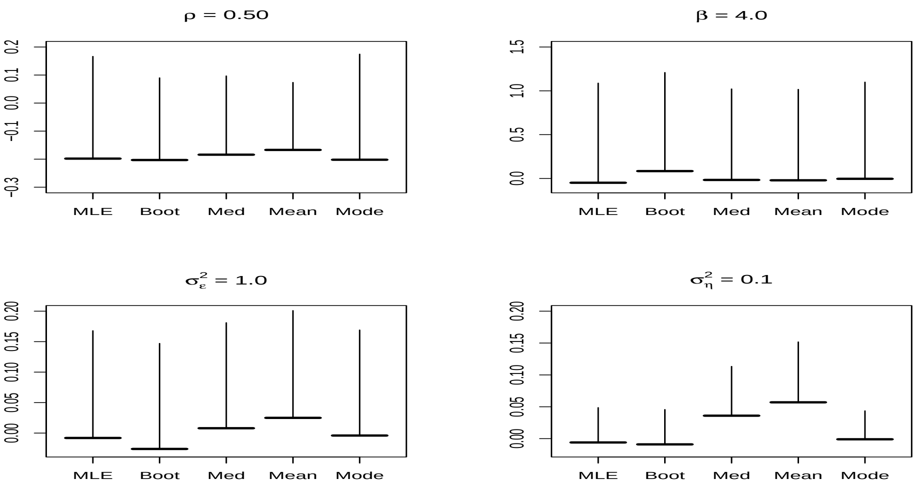

Figure 5: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with pulse function and $\rho = 0.50$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

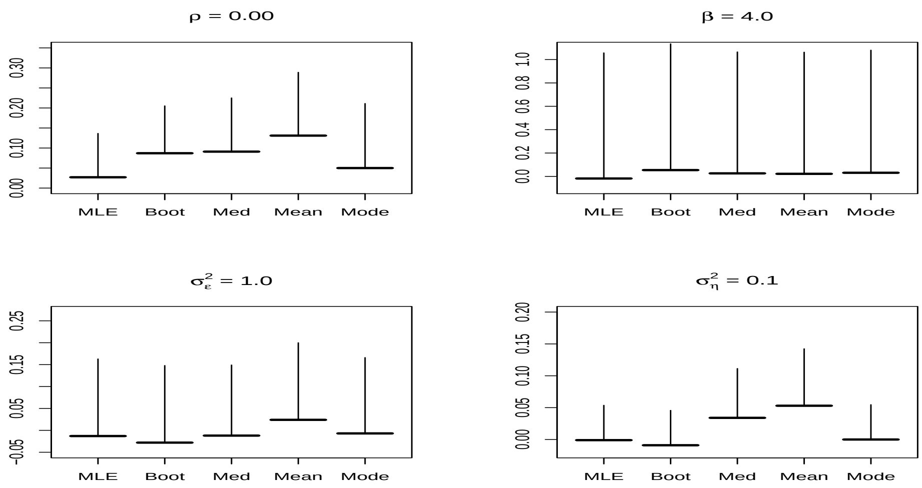

Figure 6: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with pulse function and $\rho = 0.00$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

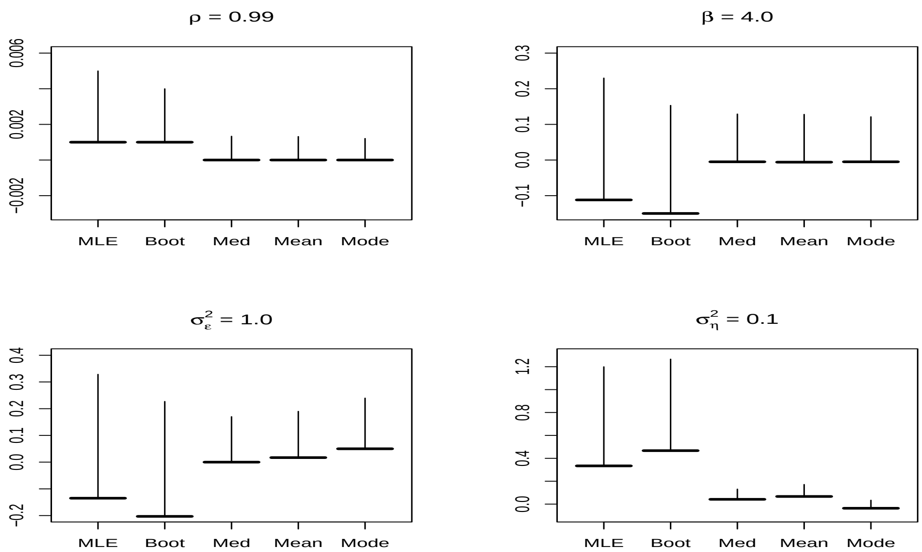

Figure 7: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with step function and $\rho = 0.99$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

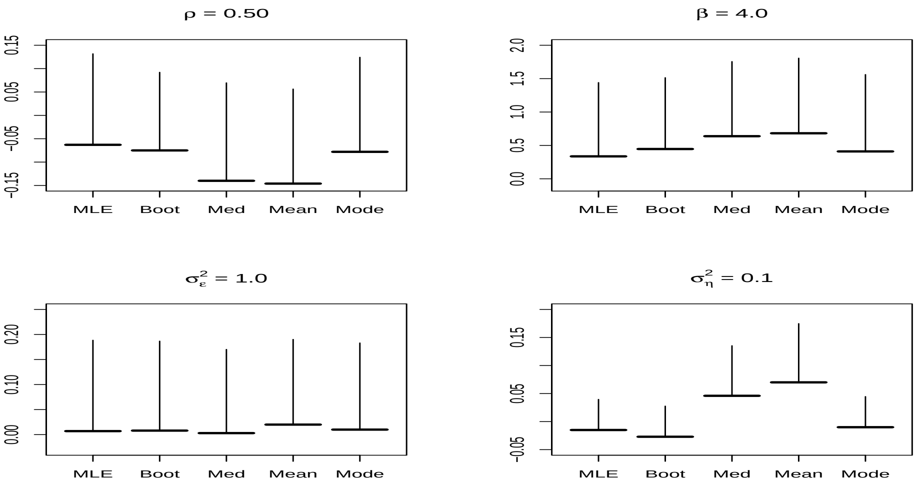

Figure 8: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with step function and $\rho = 0.50$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

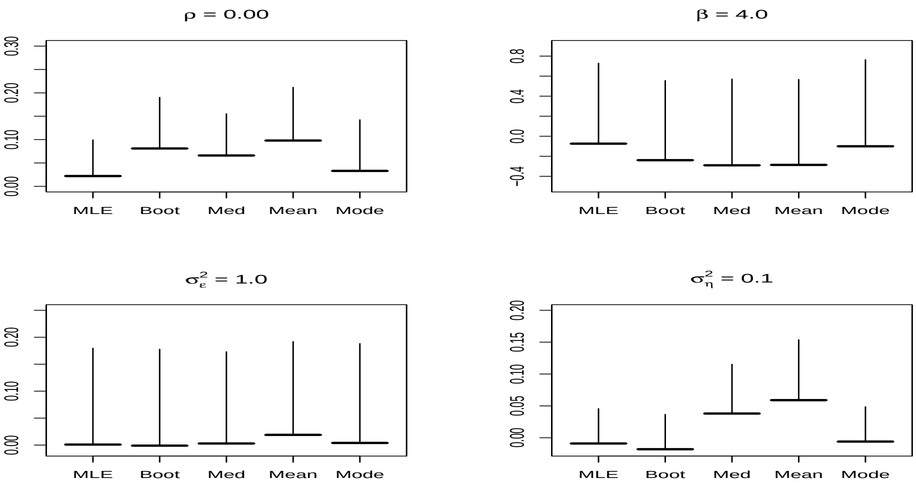

Figure 9: Maximum likelihood, bootstrap and Bayesian estimation for Model 1 with step function and $\rho = 0.00$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

pulse and step functions. In these tables, besides the estimators, the one-step ahead forecast is also included. When $\rho = 0.99$, both approaches have intervals with coverage close to 0.95 in the pulse function for all the parameters considered, except the bootstrap confidence interval for ρ. By the other hand, the performance of the credibility interval for the step function is much better than the bootstrap interval, except for the one-step ahead forecast, where both are very far from the 0.95 point. If $\rho = 0.50$ or 0.00, both intervals show approximately the same behavior for all parameters, either for the pulse and step functions, with coverage rates very close to 0.95. The only exceptions are the bootstrap interval for σ^2_η, which is slightly worse than the credibility interval, and the intervals for ρ in the step function, where both methods present coverage rate below 0.90. It should be noted that intervals for $\rho = 0.00$ are not built because, being at the boundary of parameter space, the intervals would not cover the true value of this parameter.

Table 2: Confidence and credibility intervals for Model 1 with step function.

	True	Confidence Interval			Credibility Interval		
	value	mean limits	width	coverage	mean limits	width	coverage
ρ	0.99	[0.988; 0.998]	0.010	0.75	[0.987; 0.992]	0.005	0.88
β	4.00	[3.165; 4.169]	1.004	0.78	[3.726; 4.257]	0.531	0.90
σ^2_ε	1.00	[0.027; 1.255]	1.228	0.79	[0.688; 1.434]	0.746	0.94
σ^2_η	0.10	[0.201; 2.238]	2.037	0.81	[0.035; 0.443]	0.407	0.95
y_{n+1}	*	[155.64; 160.35]	4.71	0.51	[155.45; 159.83]	4.38	0.38
ρ	0.50	[0.139; 0.655]	0.516	0.89	[0.090; 0.601]	0.511	0.81
β	4.00	[2.839; 6.330]	3.491	0.92	[2.964; 6.634]	3.670	0.91
σ^2_ε	1.00	[0.695; 1.365]	0.670	0.94	[0.692; 1.441]	0.749	0.96
σ^2_η	0.10	[0.005; 0.187]	0.182	0.80	[0.046; 0.436]	0.390	0.94
y_{n+1}	*	[5.761; 10.052]	4.291	0.93	[5.536; 9.819]	4.283	0.94
ρ	0.00	-	-	-	-	-	-
β	4.00	[2.135; 5.329]	3.194	0.94	[0.729; 5.465]	4.736	0.95
σ^2_ε	1.00	[0.686; 1.360]	0.674	0.93	[0.697; 1.434]	0.737	0.95
σ^2_η	0.10	[0.007; 0.201]	0.194	0.84	[0.046; 0.395]	0.349	0.93
y_{n+1}	*	[1.792; 5.977]	4.185	0.92	[1.543; 5.816]	4.273	0.94

* - there is not a single true value for y_{n+1} as each generated series has its own value.

4.2 Model 2

Figures 10, 11 and 12 present the average of estimated bias and MSE over 500 MC replications for Model 2. As was explained in Section 2, in this case the model is not able to estimate the influence of a pulse function. Therefore, only the results for a step function are presented below.

It can be seen that, in general, the MLE and the posterior mode are slightly better than the other procedures for estimating all parameters, with smaller bias and MSE. Once again, ρ is always underestimated, except in the case $\rho = 0.00$.

The results for the bootstrap and credibility intervals with the step function are presented in Table 3. The performance of the credibility interval is again better than the bootstrap interval, in the case $\rho = 0.99$, like Model 1 with the step function. If $\rho = 0.50$ or 0.00, the credibility interval shows, in general, coverage rates closer to 0.95, except for σ^2_ξ, where the bootstrap interval is better. Once again, intervals for $\rho = 0.00$ are not built.

Table 3: Confidence and credibility intervals for Model 2 with step function.

	True value	Confidence Interval			Credibility Interval		
		mean limits	width	coverage	mean limits	width	coverage
ρ	0.99	[0.671; 0.998]	0.327	0.98	[0.887;0.997]	0.110	0.96
σ^2_ξ	0.50	[0.128; 1.243]	1.115	0.99	[0.255; 1.815]	1.125	0.93
σ^2_ε	1.00	[0.656; 1.407]	0.751	0.96	[0.604; 1.436]	0.832	0.97
σ^2_η	0.10	[0.001; 0.280]	0.279	0.74	[0.032; 0.722]	0.690	0.95
y_{n+1}	*	[-27.19; -22.71]	4.48	0.27	[-29.49; -20.53]	6.956	0.81
ρ	0.50	[0.001; 0.608]	0.607	0.67	[0.103; 0.613]	0.510	0.81
σ^2_ξ	0.50	[0.169; 2.058]	1.889	0.93	[0.240; 2.174]	1.934	0.90
σ^2_ε	1.00	[0.562; 1.343]	0.781	0.90	[0.559; 1.420]	0.861	0.95
σ^2_η	0.10	[0.003; 0.309]	0.306	0.85	[0.041; 0.762]	0.721	0.92
y_{n+1}	*	[-3.434; 0.855]	4.289	0.77	[-3.258; 2.660]	5.920	0.84
ρ	0.00	-	-	-	-	-	-
σ^2_ξ	0.50	[0.052; 1.141]	1.089	0.92	[0.102; 1.526]	1.424	0.97
σ^2_ε	1.00	[0.636; 1.417]	0.781	0.93	[0.623; 1.471]	0.848	0.97
σ^2_η	0.10	[0.003; 0.270]	0.267	0.85	[0.037; 0.588]	0.551	0.93
y_{n+1}	*	[-2.759; 1.368]	4.127	0.79	[-2.662; 3.329]	5.991	0.87

Obs.: There is not a single true value for y_{n+1} as each generated series has its own value.

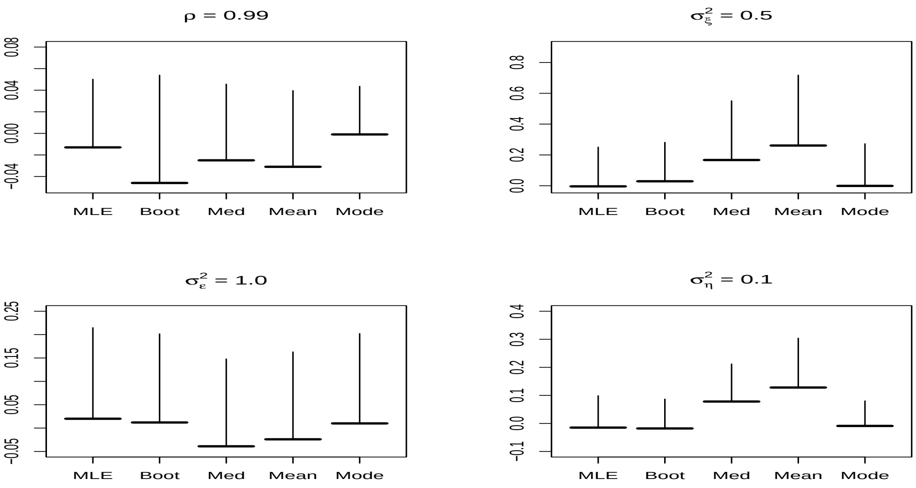

Figure 10: Maximum likelihood, bootstrap and Bayesian estimation for Model 2 with step function and $\rho = 0.99$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

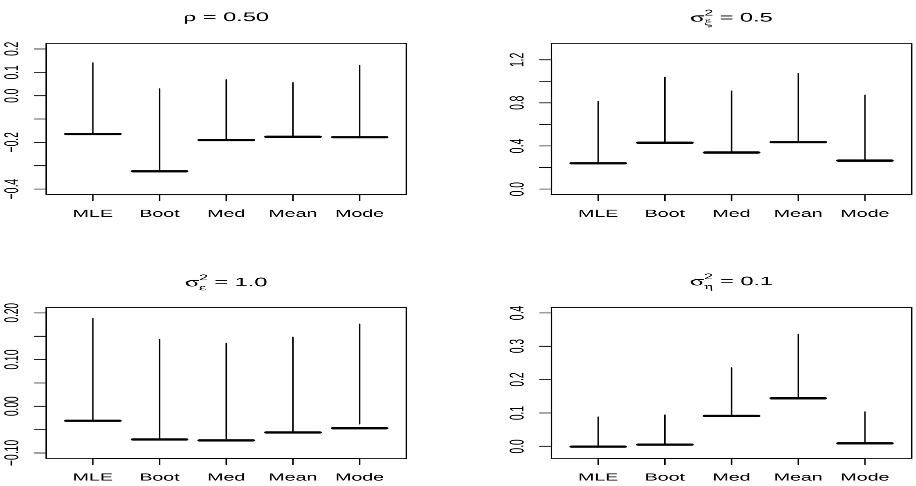

Figure 11: Maximum likelihood, bootstrap and Bayesian estimation for Model 2 with step function and $\rho = 0.50$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

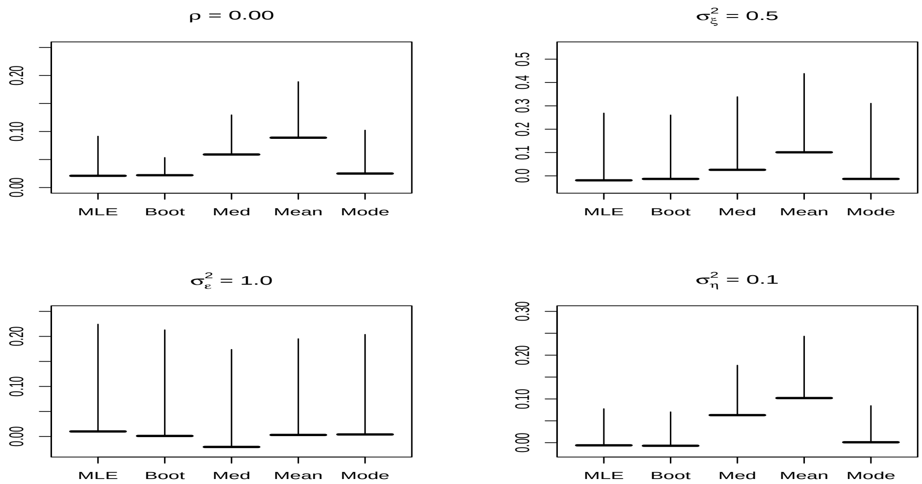

Figure 12: Maximum likelihood, bootstrap and Bayesian estimation for Model 2 with step function and $\rho = 0.00$. Horizontal and vertical lines indicate the bias and the MSE root, respectively.

5 Application to real series

In this section, the methodology previously described is applied to two real time series. The first one is the Ample Price to Consumer Index (APCI) in the city of Belo Horizonte, Brazil, from July, 1997 to June, 2008. This series appears to have a pulse intervention around October, 2002. The second series is the monthly index of BOVESPA (stock market of São Paulo, Brazil), in the period from January, 1991 to August, 2008. For this series, the intervention takes the form of a step function, presenting a jump around mid 1994.

5.1 APCI series

The data for the Ample Price to Consumer Index in Belo Horizonte are collected by IPEAD - *Fundação Instituto de Pesquisas Econômicas, Administrativas e Contábeis de Minas Gerais - Brazil.* This index measures the evolution of the incomes in families spending from 1 to 40 minimum salaries per month. The APCI series is composed of 132 monthly observations in the period July, 1997 to June, 2008.

From Figure 13, the APCI series seems to present a behavior of a local level model (LLM), with an intervention around October, 2002. This shift was due to the concerns in the economy after the election of the leftist President Lula. Therefore, the LLM was fitted with and without the intervention component, in order to assess the relevance of the intervention. As it was seen in Section 2, Model 2 is not appropriate in this case, as the intervention is of a pulse form.

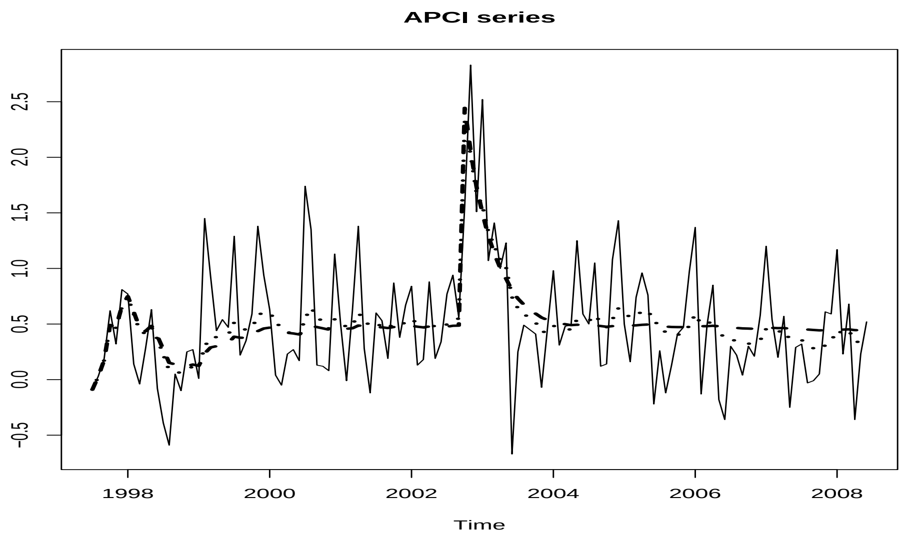

Figure 13: The solid line indicates the series, y_t, and the dotted line and dashed line indicate the mean responses of Bayesian Model 1 and Classical Model 1, respectively.

Table 4: AIC and one-step ahead forecast for the APCI series.

Models	AIC	Forecast (true value=0.350)	Forecast Int.
LLM Classical	212.650	0.320	[-0.508; 1.456]
LLM Bayesian	212.600	0.330	[-0.302; 0.962]
Model 1 Classical	198.050	0.419	[-0.407; 1.442]
Model 1 Bayesian	198.101	0.349	[-0.443; 1.430]

Table 5: Fit of Model 1 to the APCI series.

	Classical Inference		Bayesian Inference	
	MLE	Interval	Mode	Interval
ρ	0.797	[0.030; 0.883]	0.801	[0.558; 0.884]
β	1.975	[1.411; 2.888]	1.968	[0.851; 2.685]
σ^2_ε	0.225	[0.167; 0.277]	0.222	[0.168; 0.298]
σ^2_η	5.88×10^{-12}	[2.13×10^{-18}; 1.63×10^{-3}]	3.52×10^{-5}	[2.88×10^{-5}; 2.52×10^{-2}]

From Table 4, it is verified that Model 1, in both classical and Bayesian estimation, presents a lower AIC than the LLM, thus confirming that the pulse intervention should be included in the model. One-step ahead forecast was calculated for July, 2008. The true value is 0.35, and it can be seen that the Bayesian prediction is much closer to the true value than the classical one, with smaller interval width. Estimates for the parameters are shown in Table 5. It seems that the level term μ_t is constant in time, as the credibility interval of σ^2_η is tightly concentrated around zero. The most probable value of ρ is 0.80 with a skewed posterior distribution. The posterior mode of the magnitude of the jump is 1.934, compatible with the visual inspection of the series.

5.2 IBOVESPA series

IBOVESPA series is the monthly series of an index of the stock market of São Paulo (Brazil), in the period from January, 1991 to August, 2008. This series is composed of 212 monthly observations. From Figure 14, it seems IBOVESPA follows a LLM with intervention component of a step form. The series has values above zero up to mid 1994 and around zero afterwards. One reason for this behavior is the introduction of the Real Plan, a government strategy that changed currency in Brazil, implemented in July 1994. In order to compare the procedures, the LLM and Models 1 and 2 with intervention of a step form will be fitted to IBOVESPA.

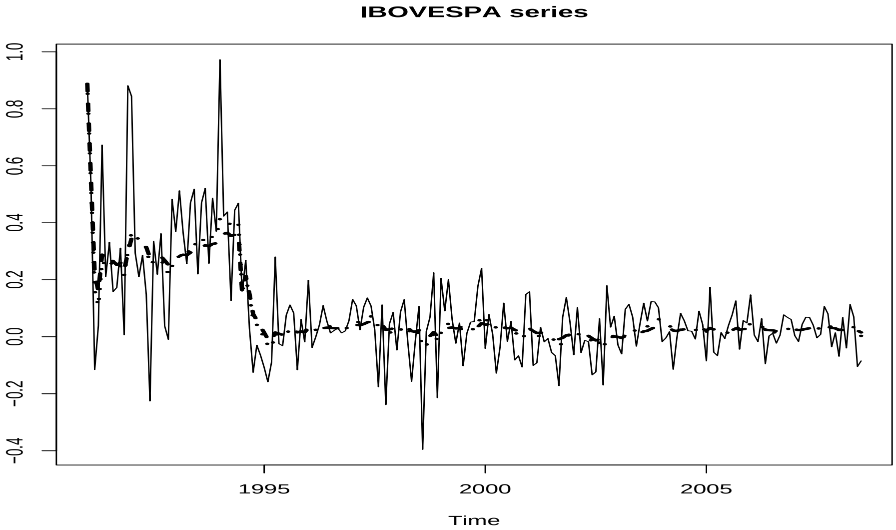

Figure 14: The solid line indicates the series, y_t, and the dotted line and dashed line indicate the mean responses of Classical Model 2 and Bayesian Model 2, respectively.

Table 6 shows the AIC and one-step ahead forecast using the three models mentioned above. It can be seen that Model 2, in both classical and Bayesian approaches, presents a lower AIC than the LLM and Model 1, but the Bayesian procedure gives a forecast closer to the true value of -0.064 than the classical procedure, with smaller interval width. Estimates for the parameters are shown in Table 7. Once again, it seems that the level term μ_t is constant in time, as it can be verified from the credibility interval for σ^2_η. The value of ρ is very small.

6 Concluding remarks

Intervention analysis with structural models was investigated using classical and Bayesian approaches for inference. Intervention was only applied for the mean level of the series and in two canonical forms: pulse and step functions. These functions allowed for breaks of abrupt and gradual forms, respectively. Static and dynamic gain factors were considered. Normal observations were assumed and they led to

Table 6: AIC and one step ahead forecast for the IBOVESPA series.

Models	AIC	forecast (true value=-0.064)	Interval forecast
LLM Classical	-174.389	-0.020	[-0.291; 0.229]
LLM Bayesian	-174.042	-0.009	[-0.314; 0.294]
Model 1 Classical	-212.488	0.020	[-0.255; 0.296]
Model 1 Bayesian	-212.050	0.003	[-0.229; 0.348]
Model 2 Classical	-255.019	0.010	[-0.291; 0.306]
Model 2 Bayesian	-255.009	-0.016	[-0.195; 1.001]

Table 7: Fit of Model 2 to the IBOVESPA series.

	Classical Inference		Bayesian Inference	
	MLE	Interval	Mode	Interval
ρ	0	[0; 0.022]	0.002	[0.001; 0.184]
σ^2_ξ	0.079	[0.042; 0.123]	0.090	[0.046; 0.164]
σ^2_ε	0.010	[0.010; 0.012]	0.010	[0.008; 0.013]
σ^2_η	2.30×10^{-13}	$[3.69\times10^{-21}; 3.23\times10^{-5}]$	1.086×10^{-6}	$[1.815\times10^{-7}; 2.427\times10^{-4}]$

an exact expression for the integrated likelihood for the hyperparameters. Under the classical approach, maximum likelihood (ML) was performed and approximate confidence intervals were built using the bootstrap technique. Under the Bayesian approach, the posterior distribution was obtained and MCMC methods were used to approximate it. Other approximating methods could also be used to perform the likelihood-based inference presented. Some of these alternative methods are mentioned in Durbin & Koopman (2001).

Comparisons between the classical and Bayesian procedures were performed through extensive Monte Carlo simulation in a variety of parameter settings. The ML estimator and the posterior mode behaved very similarly and presented smaller bias and MSE than posterior means and medians overall. The only exception was near non-stationary persistence ($\rho = 0.99$) with static gain factor, where all

Bayesian point estimators showed a clear superiority over classical ones.

Confidence and credibility intervals were built for the parameters and forecasts and compared in terms of width and coverage rate. Credibility intervals provided better results with step function intervention, for the models considered. For the pulse function, the coverage rates of both intervals were very close, in general, to the nominal level.

There are a number of other issues related to the topic that were not addressed here. This paper focused more directly towards hyperparameter estimation and prediction of future values. Therefore, smoothed distributions for the state parameters were not discussed but they can be easily obtained. Seasonality and more general trend forms can also be considered as extra components to be added to the model and can be easily included. Results could also be extended to intervention models in the exponential family, considered in Alves et al. (2009), but exact likelihood is no longer analytically available. A more thorough discussion about forecasting with incorporation of the uncertainty about the hyperparameters could also be undertaken but will be left for future work. This task requires more consideration along the lines suggested by Harvey (1989) for classical inference although this issue is automatically taken into account in Bayesian inference.

Acknowledgements

The author G. C. Franco was partially supported by CNPq-Brazil and *Fundação de Amparo à Pesquisa no Estado de Minas Gerais* (FAPEMIG Foundation). The author Dani Gamerman was supported by CNPq-Brazil and *Fundação de Amparo à Pesquisa no Estado do Rio de Janeiro* (FAPERJ Foundation). The author T. R. Santos was supported by CNPq-Brazil.

References

Alves, M., Gamerman, D. & Ferreira, M. A. R. (2009). Transfer functions in dynamic generalized linear models. To appear in *Stat. Model.*

Box, G. E. P., Jenkins, G. M. & Reinsel, G. C. (1994). *Time Series Analysis: Forecasting and Control*, 3rd ed. New Jersey: Prentice-Hall.

Box, G.E.P. & Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *J. Amer. Statist. Assoc.*, 70, 70-79.

Box, G. E. P. & Jenkins, G. M. (1976). *Time Series Analysis: Forecasting and Control*, 2nd ed. San Francisco: Holden Day.

Brockwell, P. J. & Davis, R. A. (1996). *Introduction to Time Series and Forecasting*, Springer text in Statistics. New York: Springer.

Campbell, D. T. & Stanley, J. C. (1968). The Connecticut crackdown on speeding: Time series data in quasi-experimental analysis. *Law and Society Review*, 3, 33-53.

Carter, C. K. & Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81, 541-533.

Chib, S. & Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *Amer. Statist.*, 49, 327-335.

de Jong, P. & Penzer, J. (1998). Diagnosing shocks in time series. *J. Amer. Statist. Assoc.*, 93, 796-806.

Doornik, J. A. (1999). *Ox: An Object-Oriented Matrix Language*, 3rd ed. London: Timberlake Consultants Press.

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

Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Ann. Statist.*, 7, 1-26.

Efron, B. & Tibshirani, R. (1993). *An Introduction to the Bootstrap.* London: Chapman and Hall.

Franco, G. C. & Souza, R. C. (2002). A comparison of methods for bootstrapping in the local level model. *Journal of Forecasting*, 21, 27-38.

Franco, G. C., Santos, T. R., Ribeiro, J. A. & Cruz, F. R. B.(2008). Confidence intervals for hyperparameters in structural models. *Communications in Statistics*, **37**(3), 486-497.

Fruhwirth-Schnatter, S. (1994). Applied state space modelling of non-Gaussian time series using integration-based Kalman filtering. *Statist. Comput.*, 4, 259-269.

Gamerman, D. & Lopes, H. F. (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference.* London: Chapman and Hall.

Glass, G. V. (1972). Estimating the effects of intervention into a nonstationary time series. *American Educational Research Journal*, **9**(3), 463-477.

Gomes, C. H. A. (2006). *Bayesian Estimation with Jeffreys Priors for State Space Models.* Unpublished M.Sc. Thesis. Universidade Federal do Rio de Janerio (in Portuguese).

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter.* Cambridge: University Press.

Harvey, A. C. & Koopman, S. J. (1992). Diagnostic checking of unobserved components time series models. *J. Bus. Econom. Statist.*, 10, 377-389.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Trans. ASME J. Basic Eng.*, 82, 35-45.

Lopes, H. F., Moreira, A. R. B. & Schmidt, A. M. (1999). Hyperparameter estimation in forecast models. *Comput. Statist. Data Anal.*, 29, 387-410.

Migon, H. S. & Gamerman, D. (1999). *Statistical Inference: an Integrated Approach.* London: Arnold.

Penzer, J. (2007). State space models for time series with patches of unusual observations. *J. Time Ser. Anal.*, 28, 629-645.

Pfeffermann, D. & Tiller, R. (2005). Bootstrap approximation to prediction MSE for state-space models with estimated parameters. *J. Time Ser. Anal.*, 26, 893–916.

Plummer, M., Best, N., Cowles, K. & Vines, K. (2005). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, **6**, 7-11. Software avaliable at http://www-fis.iarc.fr/coda/.

Ravines, R. E. R., Schmidt, A. M. & Migon,H. S. (2008). Revisiting distributed lag models through a Bayesian perspective. *Appl. Stoch. Models Bus. Ind.*, 22, 193-210.

Reis, E. A., Salazar, E. & Gamerman, D. (2006). Comparison of sampling schemes for dynamic linear models. *International Statistical Review*, 74, 203-214.

Salvador, M. & Gargallo, P. (2004). Automatic monitoring and intervention in multivariate dynamic linear models. *Comput. Statist. Data Anal.*, 47, 401-431.

Santos, T. R. & Franco, G. C. (2008). Inference for hyperparameters in structural models under Bayesian and classical perspectives (submitted). *Comput. Statist.* (Zeitschrift).

Schmidt, A. M., Gamerman, D., & Moreira, A.R.B. (1999). An adaptive resampling scheme for cycle estimation. *J. Appl. Statist.*, 26, 619-641.

Stoffer, D. S. & Wall, K. D. (1991). Bootstrapping state-space models: Gaussian maximum likelihood estimation and the Kalman filter. *J. Amer. Statist. Assoc.*, 86, 1024-1033.

Thombs, L. A. & Schucany, W. R. (1990). Bootstrap prediction intervals for autoregression. *J. Amer. Statist. Assoc.*, 85, 486-492.

Wei, W. S. (1991). *Time Series Analysis - Univariate and Multivariate Methods.* New York: Addison-Wesley.

West, M. & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models.* New York: Springer.