

*Special Issue in Games and Decisions  
in Reliability and Risk, GDRR 2017.*

## Assessing the effect of advertising expenditures upon sales: a Bayesian structural time series model

Víctor Gallego<sup>1</sup>, Pablo Suárez-García<sup>2</sup>, Pablo Angulo<sup>3</sup> and David  
Gómez-Ullate<sup>1,2,4</sup>

<sup>1</sup> *Institute of Mathematical Sciences, ICMAT-CSIC*

<sup>2</sup> *Facultad de Ciencias Físicas, Universidad Complutense de Madrid*

<sup>3</sup> *ETSIN, Universidad Politécnica de Madrid*

<sup>4</sup> *Department of Computer Science, School of Engineering, Universidad de Cádiz*

emails: victor.gallego@icmat.es, pasuarez@ucm.es, pablo.angulo@upm.es,  
david.gomez-ullate@icmat.es

### Abstract

We propose a robust implementation of the Nerlove–Arrow model using a Bayesian structural time series model to explain the relationship between advertising expenditures of a country-wide fast-food franchise network with its weekly sales. Thanks to the flexibility and modularity of the model, it is well suited to generalization to other markets or situations. Its Bayesian nature facilitates incorporating *a priori* information reflecting the manager’s views, which can be updated with relevant data. This aspect of the model will be used to support the decision of the manager on the budget scheduling of the advertising firm across time and channels.

*Key words: Bayesian structural time series, Sales forecasting, Risk management, Budget allocation*

## 1 Introduction

It is widely acknowledged that a firm’s expenditure on advertising has a positive effect on sales [1, 2, 3, 4]. However, the exact relationship between them remains a moot point, see *e.g.* [5] for a broad survey. Since Dorfman and Steiner’s [6] seminal work—one of the first formal theories of optimal monopoly advertising—several models have been proposed to pinpoint this relationship, although consensus on the best approach has not been reached yet. Two diverging model-building schools seem to dominate the marketing literature [7]:

*a priori* models that rely heavily on intuition and are derived from general principles, although usually with a practical implementation on mind ([8], or [9] and [10] *inter alia*) and *statistical* or *econometric* models, which usually start from a specific dataset to be modelled (see [1] for a review). In this work we will mostly rely on the first type of models, *viz.* that of Nerlove-Arrow [8], which extends Dorfman and Steiner to a dynamic setting [11] and adapts seamlessly to the *state-space* or *structural* time series approach.

Bayesian structural time series models [12], in turn, have positioned themselves in the past few years as very effective tools not only for analysing marketing time-series, but also to throw light into more uncertain terrains like causal impacts, incorporating *a priori* information into the model, accommodating multiple sources of variations or supporting variable selection. Although the origins of this formalism can be traced back to the 1950's in engineering problems of filtering, smoothing and forecasting, first with Wiener [13] and specially with Kalman [14], these problems can also be understood from the perspective of estimation in which a vector valued time series  $\{X_0, X_1, X_2, \dots\}$  that we wish to estimate (the *latent* or *hidden* states) is observed through a series of noisy measurements  $\{Y_0, Y_1, Y_2, \dots\}$ . This means that, in the Bayesian sense, we want to compute the joint posterior distribution of all the states given all the measurements [15]. The ever-growing computing power and release of several programming libraries in the last few years like [16, 17] have in part alleviated the difficulties in the mathematical underpinning and computer implementation that this formalism suffers, making these methods broadly known and used. This family of models have been used successfully to model financial time series data [18], infer causal impact of marketing campaigns [19], select variables and nowcast consumer sentiment [20], or for predicting other economic time series models like unemployment [12].

In this paper we use the formalism of Bayesian structural time-series models to formulate a robust model that links advertising expenditures with weekly sales. Due to the flexibility and modularity of the model, it will be well suited to generalization to various markets or situations. Its Bayesian nature also adapts smoothly to the issue of introducing outside or *a priori* information, which can be updated to the posterior distribution of the estimated parameters. The formulation of the model allows for non-gaussian innovations of the process, which will take care of the heavy-tailedness of the distribution of sales increments. We also discuss how the forecasts produced by this model can help the manager into allocating the advertising budget. The decision space is reduced to a one dimensional curve of Pareto optimal strategies for the two moments of the forecast distribution: expected return and variance.

The paper is organized as follows: after a brief review of the most usual marketing models and the formalism of structural time series in section 2, we define the model to be used to fit the data. The experimental setup will be detailed in section 3. Section 4 provides a discussion in which alternative models will be compared and also possible uses of this model in the industry. A brief summary and ideas for further research are detailed

in section 5.

## 2 Theoretical background and model definition

### 2.1 The Nerlove-Arrow model

Numerous formulations of aggregate advertising response models exist in the marketing literature, e.g. [7]. The model of Nerlove and Arrow [8] extends the Dorfman-Steiner model to cover the situation in which present advertising expenditures affect future demand for products; it is parsimonious and is considered as a standard in the quantitative marketing community. We use it as our starting point.

In this model, advertising expenditures are considered similar in many ways to investments in durable plant and equipment, in the sense that they affect the present and future character of output and, hence, the present and future net revenue of the investing firm. The idea is to define an “advertising stock” called *goodwill*  $A(t)$  which seemingly summarizes the effects of current and past advertising expenditures over demand. Then, the following dynamics is defined for the goodwill

$$\frac{dA}{dt} = qu(t) - \delta A(t), \quad (1)$$

where  $u(t)$  is the advertising spending rate (e.g., euros or gross rating points per week),  $q$  is a parameter that reflects the advertising quality (an effectiveness coefficient) and  $\delta$  is a decay or forgetting rate. The *goodwill* then increases linearly with the advertisement expenditure but decreases also linearly due to forgetting.

Several extensions and modifications have been proposed to this simple model: it can include a limit for potential costumers [9], a non-linear response function to advertise expenditures [10], wear-in and wear-out effects of advertising [21], interactions between different advertising channels [22], among other. Still, for most of the tasks, the Nerlove-Arrow model remains as a simple and solid starting point.

### 2.2 Bayesian structural time series models

*Structural time series models* or *state-space models* provide a general formulation that allows a unified treatment of virtually any linear time series model through the general algorithms of the Kalman filter and the associated smoother. Several handbooks [23, 24, 15, 25] discuss this topic in depth, so we will not develop the corresponding theory. We will however present the most salient features that concern our modelling problem. For further details, the reader may check the aforementioned handbooks.

The state-space formulation of a time series consists of two different equations: the *state* or *evolution equation* which determines the dynamics of the state of the system as a

first-order Markov process — usually parametrized through *state* variables — and an *observation* or *measurement equation* which links the latent state with the observed state. Both equations are also affected by noise. Denoting by  $\theta_t$  the  $m \times 1$  *state vector* describing the inner state of the system, by  $G_t$  the  $m \times m$  matrix that generates the dynamics, by  $H_t$  a  $m \times g$  matrix and by  $\epsilon_t$  a  $g \times 1$  vector of serially uncorrelated disturbances with mean zero and covariance matrix  $W_t$ , the system would evolve according to the equation

$$\theta_t = G_t \theta_{t-1} + H_t \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, W_t). \quad (2)$$

The states ( $\theta_t$ ) are not generally observable, but are linked to the *observation variables*  $Y_t$  through the *observation equation*:

$$Y_t = F_t \theta_t + \epsilon'_t \quad \epsilon'_t \sim \mathcal{N}(0, V_t), \quad (3)$$

where,  $Y_t \in \mathbb{R}$  is the observed value at timestep  $t$ ,  $F_t$  is the  $1 \times m$  matrix that links the inner state to the observable, and  $V_t \in \mathbb{R}^+$  is the variance of  $\epsilon'_t$ , the random disturbances of the observations.

The specification of the state-space system is completed by assuming that the initial state vector  $\theta_0$  has mean  $\mu_0$  and a covariance matrix  $\Sigma_0$  and it is uncorrelated with the noise. The problem then consists of *estimating the sequence of states*  $\{\theta_1, \theta_2, \dots\}$  for a *given series of observations*  $\{y_1, y_2, \dots\}$  and whichever other structural parameters of the transition and observation matrices. State estimation is readily performed via the *Kalman filter*; different alternatives however arise when structural parameters are unknown. In the classical setting, these are estimated using maximum likelihood. In the Bayesian approach, the probability distribution about the unknown parameters is updated via Bayes Theorem. If exact computation through conjugate priors is not possible, the probability distributions before each measurement are updated by approximate procedures such as Markov chain Monte Carlo (MCMC) [12].

The Bayesian approach offers several advantages compared to classical methods. For instance, it is natural to incorporate external information through the prior distributions. In particular, this will be materialized in Section 2.4.1 where expert information can be incorporated through the spike and slab prior. Another useful advantage is that, due to the Bayesian nature of the model, it is straightforward to obtain predictive intervals through the predictive distribution (see Section 2.5).

## 2.3 Model specification

### 2.3.1 State-space formulation

The continuous-time Nerlove-Arrow model must be first cast in discrete time so as to formulate our model in state-space. From equation (1), we get:

$$A_t = (1 - \delta)A_{t-1} + qu_{t-1} + \epsilon_t$$

where  $A_t$  is the *goodwill* stock,  $u_t$  is the advertising spending rate,  $q$  is the effectiveness coefficient and the random disturbance  $\epsilon_t$  captures the net effects of the variables that affect the goodwill but cannot be modelled explicitly. This discrete counterpart of Nerlove-Arrow is a distributed-lag structure with geometrically declining weights, i.e., a *Koyck model* [26, 27]. Since in our setting the model includes the effect of  $k$  different channels in the goodwill, we shall modify the previous equations according to:

$$A_t = (1 - \delta)A_{t-1} + \sum_{i=1}^k q_i u_{i(t-1)} + \epsilon_t$$

Now, following the notation from (2) and (3), the discrete-time Nerlove-Arrow model in state-space form will read as follows:

**Evolution equation:**

$$\theta_t = G_t \theta_{t-1} + H_t \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, W_t), \quad (4)$$

where

$$\theta_t = \begin{bmatrix} A_t \\ q_1 \\ \vdots \\ q_k \end{bmatrix}, \quad G_t = \begin{bmatrix} (1 - \delta) & u_{1(t-1)} & \dots & u_{k(t-1)} \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}, \quad H_t = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Note that the  $q_i$  are constant over time and that the matrix  $G_t$  depends on the known inversion levels at time  $t - 1$  and on an unknown parameter ( $\delta$ ) to be estimated from the data.

**Observation equation:**

$$Y_t = F_t \theta_t + \epsilon'_t \quad \epsilon'_t \sim \mathcal{N}(0, V_t) \quad (5)$$

where  $Y_t$  are the observed sales at time  $t$  and  $F_t = [1, 0, \dots, 0]$ .

## 2.4 Modularity and additional structure

The above model is very flexible in the sense that it can be defined *modularly*, in as much as the different hidden states evolve independently of the others (*i.e.* the evolution matrix can be cast in block-diagonal form). This greatly simplifies their implementation and allows for simple building-blocks with characteristic behavior. Typical blocks specify trend and seasonal components — which can be helpful to discover additional patterns in the time series— or explanatory variables that can be added to further reduce the uncertainty in the

model and bridge the gap between time series and regression models. Via the *superposition principle* [24, Chapter 3] we could include additional blocks in our model:

$$Y_t = Y_{NA,t} + Y_{R,t} + Y_{T,t} + Y_{S,t}$$

where  $Y_{NA,t}$  corresponds to the discretized Nerlove-Arrow equation, defined in (4) and (5);  $Y_{R,t}$  is a regression component; containing the effects of external explanatory variables  $X_t$ ;  $Y_{T,t}$  is a trend component or a simpler local level component; and  $Y_{S,t}$  is a seasonal component.

#### 2.4.1 Regression components. Spike and slab variable selection

To take into account the effects of external explanatory variables, a static regression component can be easily incorporated into the model through

$$Y_{R,t} = X_t\beta + \epsilon_t,$$

where the state  $\beta$  is constant over time to favor parsimony.

A spike and slab prior [28] is used for the static regression component since it can incorporate *prior information* and also facilitate variable selection. This is specially useful for models with large number of regressors, a typical setting encountered in business scenarios. Let  $\gamma$  denote a binary vector that indicates whether the regressors are included in the regression. Specifically,  $\gamma_i = 1$  if and only if  $\beta_i \neq 0$ . The subset of  $\beta$  for which  $\gamma_i = 1$  will be denoted  $\beta_\gamma$ . Let  $\sigma_\epsilon^2$  be the residual variance from the regression part. The spike and slab prior [29] can be expressed as

$$p(\beta, \gamma, \sigma_\epsilon^2) = p(\beta_\gamma | \gamma, \sigma_\epsilon^2) p(\sigma_\epsilon^2 | \gamma) p(\gamma).$$

A usual choice for the  $\gamma$  prior is a product of Bernoulli distributions:

$$\gamma \sim \prod_i \pi_i^{\gamma_i} (1 - \pi_i)^{1 - \gamma_i}.$$

The manager of the firm may elicit these  $\pi_i$  in various ways. A reasonable choice when detailed prior information is unavailable is to set all  $\pi_i = \pi$ . Then, we may specify an expected number of non-zero coefficients by setting  $\pi = k/p$ , where  $p$  is the total number of regressors. Another possibility is to set  $\pi_i = 1$  if the manager believes that the  $i$ -th regressor is crucial for the model.

## 2.5 Model estimation and forecasting

Model parameters can be estimated using Markov Chain Monte Carlo simulation, as described in Chapter 4 of [24] or [12]. We follow the same scheme.

Let  $\Theta$  be the set of model parameters other than  $\beta$  and  $\sigma_\epsilon^2$ . The posterior distribution can be simulated with the following Gibbs sampler:

1. Simulate  $\theta \sim p(\theta|y, \Theta, \beta, \sigma_\epsilon^2)$ .
2. Simulate  $\Theta \sim p(\Theta|y, \theta, \beta, \sigma_\epsilon^2)$ .
3. Simulate  $\beta, \sigma_\epsilon^2 \sim p(\beta, \sigma_\epsilon^2|y, \theta, \Theta)$ .

Repeatedly iterating the above steps gives a sequence of draws  $\rho^{(1)}, \rho^{(2)}, \dots, \rho^{(K)} \sim p(\Theta, \beta, \sigma_\epsilon^2, \theta)$ . In our experiments, we set  $K = 4000$  and discard the first 2000 draws to avoid burn-in issues.

In order to sample from the predictive distribution, we follow the usual Bayesian approach summarized by the following predictive equation, in which  $y_{1:t}$  denotes the sequence of observed values, and  $\bar{y}$  denotes the set of values to the forecast

$$p(\bar{y}|y_{1:t}) = \int p(\bar{y}|\rho)p(\rho|y_{1:t})d\rho.$$

Thus, it is sufficient to sample from  $p(\bar{y}|\rho^{(i)})$ , which can be achieved by iterating equations (3) and (2). With these predictive samples  $\bar{y}^{(i)}$  we can compute statistics of interest regarding the predictive distribution  $p(\bar{y}|y_{1:t})$  such as the mean or variance (MC estimates of  $E[\bar{y}|y_{1:t}]$  and  $Var[\bar{y}|y_{1:t}]$ , respectively) or quantiles of interest.

### 2.5.1 Robustness

We can replace the assumption of Gaussian errors with student- $t$  errors in the observation equation, thus leading to the model

$$Y_t = F_t\theta_t + \epsilon'_t \quad \epsilon'_t \sim \mathcal{T}_\nu(0, \tau^2).$$

Typically, in these settings we set  $\nu > 1$  to make the variance finite, and this variance parameter can be estimated from data using Empirical Bayes methods, for instance. In this manner, we allow the model to predict occasional larger deviations, which is reasonable in the context of forecasting sales. For instance, a special event not taken into account through the predictor variables may lead to an increase in the sales for that week.

## 3 Case Study. Data and parameter estimation

### 3.1 Data analysis

The time series analyzed in this case study contains the total weekly sales of a country-wide franchise of fast food restaurants, Figure 1, covering the period January 2011 - June 2015,

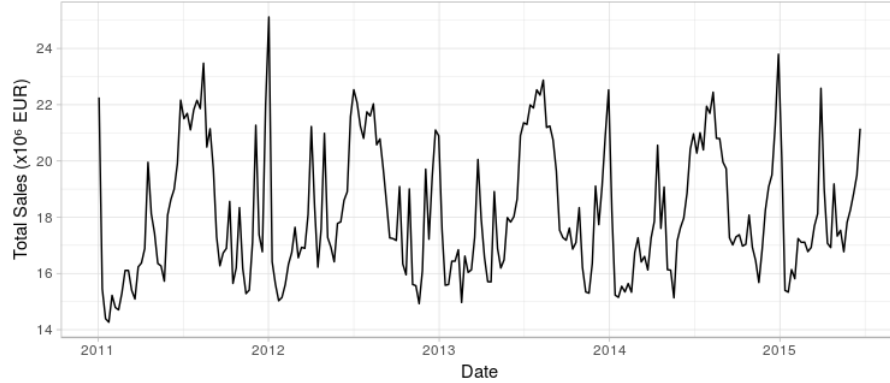


Figure 1: Total weekly sales. Jan-2011 to Jun-2015

thus comprising 234 observations. The total weekly sales is in fact the aggregated sum from the individual sales of the whole country network of 426 franchises, also allowing a fine-grained study down to the store level, although we shall not carry it out in this paper. Along with the sales figures, the series includes the investment levels  $\{u_{it}\}$  in advertising during this period for seven different channels *viz.* OOH (Out-of-home, *i.e.* billboards), Radio, TV, Online, Search, Press and Cinema, Figure 2,  $i = 1, \dots, 7$ .

From such graphs we observe that:

- In Figure 1, the series has a characteristic seasonal pattern that shows peaks in sales coinciding with Christmas, Easter and summer vacations.
- In Figure 2 we observe that the investment levels at each channel vary largely in scale, with investments in TV, and OOH dominating the other channels.
- The investment strategies adopted by the firm at each channel are also qualitatively different. Some of them show spikes while others depict a relatively even spread investment across time.

A handful of other predictors  $X_t$  which are also known to affect sales will be used in the model, all of them weekly sampled:

- Global economic indicators: unemployment rate (Unemp\_IX), price index (Price\_IX) and consumer confidence index (CC\_IX).
- Climate data: average weekly temperature (AVG\_Temp) and weekly rainfall (AVG\_Rain) along the country.
- Special events: holidays (Hols) and important sporting events (Sport\_EV).



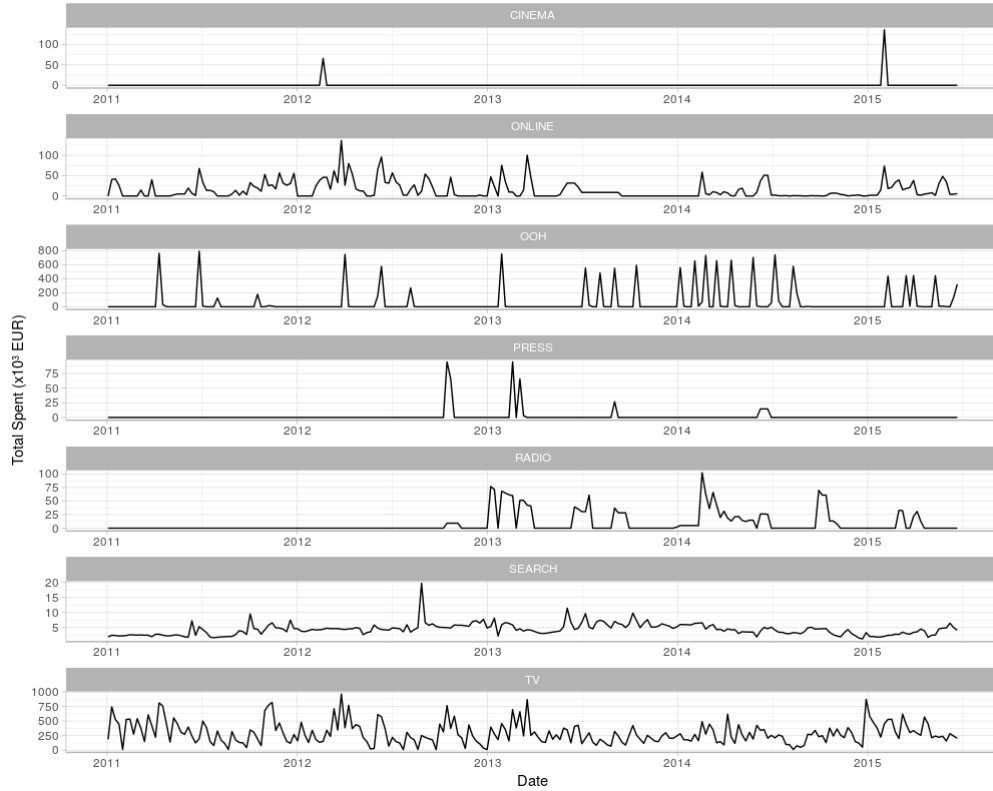


Figure 2: Advertising expenditure. Jan-2011 through Jul-2015. Note the different scales on the y-axis.

### 3.2 Experimental setup

Following the notation in Section 2.3, we consider three model variants for the particular dataset in increasing order of complexity:

- **Baseline model**, which makes use of no external variables

$$Y_t^B = Y_{NA,t} + Y_{T,t}.$$

- **Auto regression**: this model incorporates the external ambient and investment variables, so the equation of the model becomes

$$Y_t^{RA} = Y_{NA,t} + Y_{T,t} + Y_{R,t}.$$

We select an expected model size of 5 in the spike and slab prior, letting all variables to be treated equally.

- **Regression (forcing)**: the model has same equation as before (we will refer to it as RF). However, in the prior we force investment variables to be used by setting their corresponding  $\pi_i = 1$ , and imposing an expected model size of 5 for the rest of the variables.

In all cases, only the five principal advertising channels (TV, OOH, ONLINE, SEARCH and RADIO) will be used; the remaining two (CINEMA and PRESS) are sensibly lower both in magnitude and frequency than the others so we can safely disregard them in a first approximation.

As customary in a supervised learning setting with time series data, we perform the following split of our dataset: since it comprises four years of sales, we take the first two years of observations as training set, and the rest as holdout, in which we measure several predictive performance criteria. Before fitting the data, we scale the series to have zero mean and unit variance as this increases MCMC stability. Reported sales forecasts are transformed back to the original scale for easy interpretation. The models were implemented in R using the `bsts` package [17].

## 4 Discussion of results

It is customary to aim at models achieving good predictive performance. For this reason, we test the predictive performance of our three models using two metrics:

- Mean Absolute Percentage Error:

$$\text{MAPE} = \frac{100\%}{T} \sum_{t=1}^T \frac{|y_t - \hat{y}_t|}{y_t},$$

where  $y_t$  denotes the actual value;  $\hat{y}_t$ , the mean one-step-ahead prediction; and,  $T$  is the length of the hold-out period.

- Cumulative Predicted Sales over a year  $Y$ :

$$\text{CPS}_Y = \sum_{t \in \mathcal{T}(Y)} \hat{y}_t$$

where  $\mathcal{T}(Y)$  denotes the set of time-steps  $t$  contained in year  $Y$ .

These scores are reported in Table 1 with sales in million EUR. Note that, unsurprisingly, the models which include external information (RA and RF) achieve better accuracy than the baseline. In addition, we note that predictions are unbiased, since cumulative predictions are extremely close to their observed counterparts. Overall, we found the predictive performance of our models to be successful for a business scenario, as we achieve under 5% relative

absolute error using the variants augmented with external information. This is clearly useful for the decision maker who may forecast their weekly sales one week ahead to within a 5% error in the estimation.

| Model                  | MAPE  | CPS <sub>2013</sub> | CPS <sub>2014</sub> |
|------------------------|-------|---------------------|---------------------|
| B                      | 5.85% | 9660                | 9627                |
| RA                     | 4.62% | 9680                | 9613                |
| RF                     | 4.59% | 9665                | 9582                |
| Cumulative True Sales: |       | 9666                | 9610                |

Table 1: Accuracy measures for each model variant.

Figure 3 displays the predictive ability of model RF over the hold-out period. Note that the model is sufficiently flexible to adapt to fluctuations such as the peaks at Christmas. Predictive intervals also adjust their width with respect to the time to reflect varying uncertainty, yet in the worst cases they are sufficiently narrow. Further information can be tracked in Figure 4, where mean standardized residuals are plotted for each model variant. Notice that the residuals for models RA and RF are roughly comparable, being both sensibly smaller than those of the baseline. This means that the simpler Nerlove-Arrow model benefits from the addition of the ambient variables  $X_i$ , as suggested in our findings from Table 1.

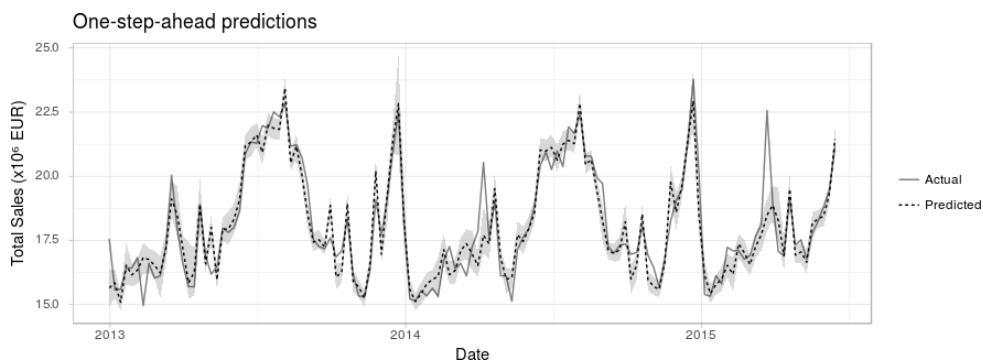


Figure 3: One-step-ahead forecasts of the RF model versus actual data during hold-out period. 95% predictive intervals are depicted in light gray.

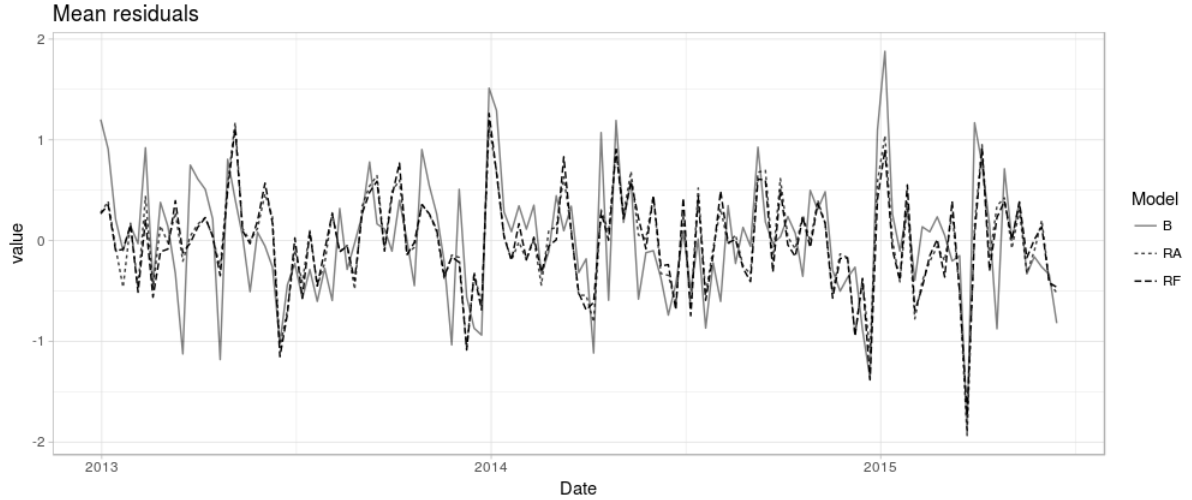


Figure 4: Mean residuals for each model

Having built good predictive models, we inspect them more closely with the aim of performing valuable inferences for our business setting. The average estimated parameter and expected standard deviations of the  $q_i$  coefficients for the different advertising channels are displayed in Table 2. We also show the weights of the ambient variables  $X_i$  for the augmented models RA and RF in Table 3, as well as the probability of a variable being selected in the MCMC simulation for a given model in Figure 5. Convergence diagnostics of the MCMC scheme are reported in Appendix A.

| Channel | Model B         |                 | Model RA        |                 | Model RF        |                 |
|---------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
|         | mean            | sd              | mean            | sd              | mean            | sd              |
| y_AR    | <b>8.00e-01</b> | <b>5.10e-02</b> | <b>5.17e-01</b> | <b>4.50e-02</b> | <b>5.07e-01</b> | <b>4.55e-02</b> |
| OOH     | 9.30e-03        | 3.11e-02        | 1.15e-03        | 9.54e-03        | <b>6.10e-02</b> | <b>3.54e-02</b> |
| ONLINE  | 1.10e-03        | 9.25e-03        | 1.95e-03        | 1.23e-02        | <b>9.04e-02</b> | <b>4.01e-02</b> |
| RADIO   | -5.34e-04       | 6.80e-03        | -2.61e-05       | 2.30e-03        | -2.57e-02       | 3.42e-02        |
| TV      | -2.85e-04       | 5.12e-03        | -5.53e-05       | 3.71e-03        | -6.14e-02       | 4.21e-02        |
| SEARCH  | 1.51e-05        | 2.74e-03        | 8.58e-05        | 2.49e-03        | 1.38e-02        | 3.50e-02        |

Table 2: Expected value and standard error of coefficients  $q_i$ . Statistically significant coefficients appear in bold.

| Channel  | Model RA         |                 | Model RF         |                 |
|----------|------------------|-----------------|------------------|-----------------|
|          | mean             | sd              | mean             | sd              |
| Sport_EV | <b>-2.06e-01</b> | <b>3.80e-02</b> | <b>-2.00e-01</b> | <b>3.71e-02</b> |
| AVG_Temp | <b>2.57e-01</b>  | <b>4.43e-02</b> | <b>2.43e-01</b>  | <b>4.65e-02</b> |
| Hols     | <b>3.10e-01</b>  | <b>3.70e-02</b> | <b>3.15e-01</b>  | <b>3.72e-02</b> |
| AVG_Rain | -2.86e-02        | 5.01e-02        | -1.96e-02        | 4.14e-02        |
| Price_IX | 1.38e-03         | 1.15e-02        | 3.60e-03         | 1.92e-02        |
| Unemp_IX | 6.34e-05         | 3.32e-03        | 2.36e-04         | 7.79e-03        |
| CC_IX    | 6.27e-07         | 2.94e-03        | 2.85e-04         | 5.50e-03        |

Table 3: Expected value and standard error of coefficients  $\beta_i$ . Statistically significant coefficients are in bold.

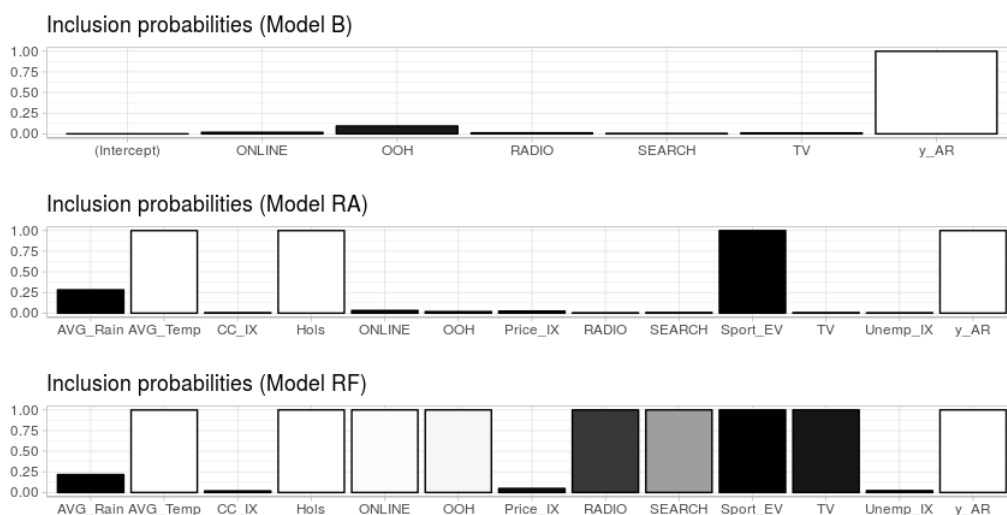


Figure 5: Selection probabilities in the MCMC simulation for each predictor variable in the three models. The color code shows variables positively (white) or negatively (black) correlated with sales.

Looking at the ambient variables, the following comments seem in order:

- From Table 3 we see that the socio-economic indicators (unemployment rate, inflation and consumer confidence) do not seem statistically relevant for the problem at hand.
- Looking at the sign of the coefficients of the most significant regressors  $X_i$  (`Hols`, `Sport_EV`, `AVG_Temp` and `AVG_Rain`) we see that they are as we would naturally expect.

Moreover, their absolute value is well above the error in both models RA and RF, a strong indicator of their influence in the expected weekly sales (*cf.* Table 3, Figure 5).

- We see, for instance, that sporting events are negatively correlated with sales. This can be interpreted as follows: major sporting events in the country where the data have been recorded receive a large media coverage and are followed by a significant fraction of the population. The restaurant chain in this study has no TVs broadcasting in their restaurants, so customers probably choose alternative places to spend their time on a day when `Sport_EV` = 1.
- The sign of `Hols` and `AVG.Temp` is positive, showing strong evidence for the fact that sales increase in periods of the year where potential customers have more leisure time, like national holidays or the summer vacation.
- One would expect `AVG.Rain` to be negatively correlated with restaurant sales, but in our study (despite having negative sign) it is not statistically significant. A possible explanation is that `AVG.Rain` records average rainfall over a large country. A more fine-grained analysis on a restaurant by restaurant basis incorporating local weather conditions will be performed in a future study, probably revealing a larger influence of this variable on the sales prediction.

Next, we turn our attention to the investment variables across different advertising channels.

- The advertising channels  $u_i$  are almost never selected in model RA, and their  $q_i$  coefficients are not significantly higher than their errors to be considered influential in the model. In model RF, however, there is strong evidence that their effect is more than a random fluctuation (*cf.* Table 2, Figure 5).
- The negative sign in both `RADIO` and `TV` in all three models suggests that (at least locally) part of the expenditures in these two channels should be diverted towards other channels with positive sign on their  $q_i$  coefficients, specially to the channel with the strongest positive coefficient (`ONLINE` and `OOH`).
- It is interesting that the trend that shows the year-to-year advertising budget of this firm (*cf.* figure 6) has a significant reduction in `TV` expenditures and a big increase in `OOH`. `RADIO` however is not reduced accordingly but increased, and `ONLINE` — which our model considers the best local inversion alternative — follows the inverse path. It has to be noted, however, that `ONLINE` expenditures are typically correlated with discounts, coupons and offers, and this information was not available to us.
- The autoregressive term is close to 0.5, which means that the immediate effect of advertising is roughly half to the *long run accumulated effects*.

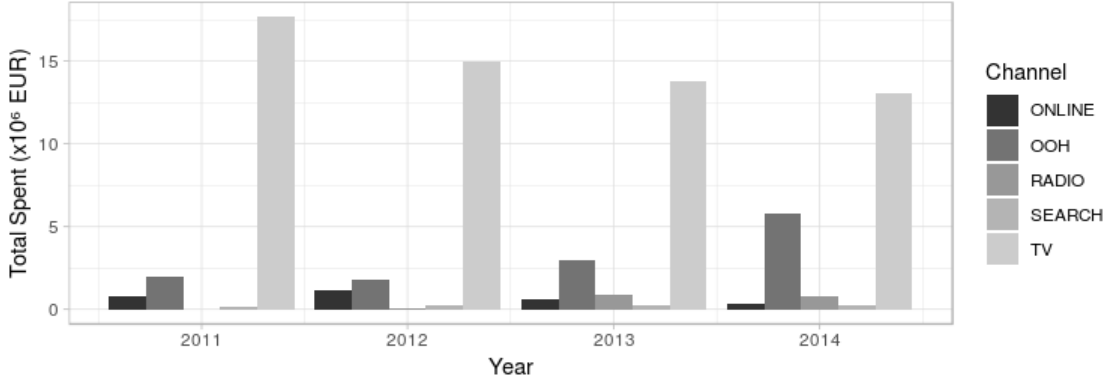


Figure 6: Yearly total expenditures on advertising campaigns for the restaurant chain, per advertising channel.

#### 4.1 Budget allocation. Model-based solutions

We propose a model which can be used as a decision support system for the manager, helping her in adopting the investment strategy in advertising. The company is interested in maximizing the expected sales for the next period, subject to a budget constraint for the advertising channels and also a risk constraint, i.e., the variance of the predicted sales must be under a certain threshold. This optimization problem based on one-step-ahead forecasts can be formulated as a non linear, but convex, problem that depends on the parameter  $\sigma^2$ :

$$\begin{aligned}
 & \underset{u_{(t+1),1} \dots u_{(t+1),k}}{\text{maximize}} && E[\bar{y}_{t+1} | y_{1:t}, x_{t+1}, u_{t+1}] \\
 & \text{subject to} && \sum_{i=1}^k u_{(t+1),i} \leq b_{t+1} \\
 & && \text{Var}[\bar{y}_{t+1} | y_{1:t}, x_{t+1}, u_{t+1}] \leq \sigma^2,
 \end{aligned}$$

where  $b_t$  is the total advertising budget for week  $t$  and  $\sigma$  is a parameter that controls the risk of the sales. We have made explicit the dependence on the regressor variables  $x_t$  and advertisement investments  $u_t$  in the mean and variance expressions. Solving for different values of  $\sigma$ , we obtain a continuum of Pareto optimal investment strategies that we can present to the manager, each one representing a different trade-off between risk and expected sales that we can plot in a risk-return diagram. This approach greatly reduces the decision space for the manager.

A possible alternative would be to rewrite the objective function as

$$E[\bar{y}_{t+1} | y_{1:t}, x_{t+1}, u_{t+1}] - \lambda \sqrt{\text{Var}[\bar{y}_{t+1} | y_{1:t}, x_{t+1}, u_{t+1}]}$$

which may be regarded as a lower quantile if the predictive samples  $\bar{y}_{t+1}^{(i)}$  are normally distributed. As in the previous approach, different values of  $\lambda$  represent different risk-return trade-offs.

If the errors in the observation equation (3) are Normal, the computations for expected predicted sales and the above variance can be done exactly and quickly using conjugacy as in [30, 3.7.1]. Otherwise, the desired quantities can be computed through Monte Carlo simulation, as in Section 2.5.

Note that due to the nature of the state-space model, it is straightforward to extend the previous optimization problem over  $k$  timesteps, for  $k > 1$ . The two objectives would be the expected total sales and the total risk over the period  $t + 1, \dots, t + k$ . For normally distributed errors, the predicted sales on the period also follow a Normal distribution, which makes computations specially simple.

The above optimization problem may not need to be solved by exhaustively searching over the space of possible channel investments. In a typical business setting, the manager would consider a discrete set of  $S$  investment strategies that are easy to interpret, so she may perform  $S$  simulations of the predictive distribution and use the above strategy to discard the Pareto suboptimal strategies.

## 5 Summary and future work

We have developed a data-driven approach for the management of advertising investments of a firm. First, using the firm's investment levels in advertising, we propose a formulation of the Nerlove-Arrow model via a Bayesian structural time series to predict an economic variable (global sales) which also incorporates information from the external environment (climate, economical situation and special events). The model thus defined offers low predictive errors while maintaining interpretability and can be built in a modular fashion, which offers great flexibility to adapt it to other business scenarios. The model performs variable selection and allows to incorporate *prior* information via the *spike-and-slab* prior. It can handle non-gaussian deviations and also provides hints to which of the advertising channels are having positive effects upon sales. The model can be used as a basis for a decision support system by the manager of the firm, helping with the task of allocating ad investments.

Possible extensions of this model include:

1. Use a different model to explain the influence of advertising upon sales. This model could take into account interactions between the channels or allow for different long-term effects for each of them.
2. Develop a model-based strategy for long-run temporal and cross-sectional budget allocation.



3. Model each restaurant individually instead of using total aggregated sales. For this approach to be sound, it would entail to use both local values of the ambient variables  $X_i$  and some of the investment channels  $u_i$  (e.g. `OOH`).
4. Include in the model the effect of special discounts, promotions and coupons, since these are probably highly correlated with some of the channels, specially `ONLINE` and `SEARCH`.

## Acknowledgements

The authors acknowledge financial support from the Spanish Ministry of Economy and Competitiveness, through the “Severo Ochoa Programme for Centres of Excellence in R&D” (SEV-2015-0554). V.G. acknowledges support for grant FPU16/05034. The research of D.G-U is supported in part by Spanish MINECO-FEDER Grant MTM2015-65888-C4-3 and MTM2015-72907-EXP. The authors also thank the members of the SPOR-Datalab group at ICMAT for their suggestions and support.

## References

- [1] Gert Assmus, John U Farley, and Donald R Lehmann. How advertising affects sales: Meta-analysis of econometric results. *Journal of Marketing Research*, pages 65–74, 1984.
- [2] Gerard J Tellis, GJ Tellis, and T Ambler. Advertising effectiveness in contemporary markets. *The SAGE Handbook of Advertising*, page 264, 2007.
- [3] Xueming Luo and Pieter J de Jong. Does advertising spending really work? the intermediate role of analysts in the impact of advertising on firm value. *Journal of the Academy of Marketing Science*, 40(4):605–624, 2012.
- [4] Thorsten Wiesel, Koen Pauwels, and Joep Arts. Practice prize paper—marketing’s profit impact: quantifying online and off-line funnel progression. *Marketing Science*, 30(4):604–611, 2011.
- [5] Gerard J Tellis. Generalizations about advertising effectiveness in markets. *Journal of Advertising Research*, 49(2):240–245, 2009.
- [6] Robert Dorfman and Peter O Steiner. Optimal advertising and optimal quality. *The American Economic Review*, 44(5):826–836, 1954.
- [7] John DC Little. Aggregate advertising models: The state of the art. *Operations research*, 27(4):629–667, 1979.

- [8] Marc Nerlove and Kenneth J Arrow. Optimal advertising policy under dynamic conditions. *Economica*, pages 129–142, 1962.
- [9] ML Vidale and HB Wolfe. An operations-research study of sales response to advertising. *Operations research*, 5(3):370–381, 1957.
- [10] John DC Little. Brandaid: A marketing-mix model, part 1: Structure. *Operations Research*, 23(4):628–655, 1975.
- [11] Kyle Bagwell. The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844, 2007.
- [12] Steven L Scott and Hal R Varian. Predicting the present with bayesian structural time series. *International Journal of Mathematical Modelling and Numerical Optimisation*, 5(1-2):4–23, 2014.
- [13] Norbert Wiener. Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications. 1949.
- [14] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.
- [15] Simo Särkkä. *Bayesian Filtering and Smoothing*, volume 3. Cambridge University Press, 2013.
- [16] Giovanni Petris. An r package for dynamic linear models. *Journal of Statistical Software*, 36(12):1–16, 2010.
- [17] Steven L Scott. *bsts: Bayesian structural time series*, 2016.
- [18] Patrizia Campagnoli, Pietro Muliere, and Sonia Petrone. Generalized dynamic linear models for financial time series. *Applied Stochastic Models in Business and Industry*, 17(1):27–39, 2001.
- [19] Kay H Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, Steven L Scott, et al. Inferring causal impact using bayesian structural time-series models. *The Annals of Applied Statistics*, 9(1):247–274, 2015.
- [20] Steven L Scott and Hal R Varian. Bayesian variable selection for nowcasting economic time series. In *Economic Analysis of the Digital Economy*, pages 119–135. University of Chicago Press, 2015.
- [21] Prasad A Naik, Murali K Mantrala, and Alan G Sawyer. Planning media schedules in the presence of dynamic advertising quality. *Marketing Science*, 17(3):214–235, 1998.

- [22] Frank M Bass, Norris Bruce, Sumit Majumdar, and BPS Murthi. Wearout effects of different advertising themes: A dynamic bayesian model of the advertising-sales relationship. *Marketing Science*, 26(2):179–195, 2007.
- [23] James Durbin and Siem Jan Koopman. *Time Series Analysis by State Space Methods*, volume 38. OUP Oxford, 2012.
- [24] Giovanni Petris, Sonia Petrone, and Patrizia Campagnoli. Dynamic linear models. In *Dynamic Linear Models with R*, pages 31–84. Springer, 2009.
- [25] Mike West and Jeff Harrison. *Bayesian forecasting and dynamic models*. Springer Science & Business Media, 2006.
- [26] Darral G Clarke. Econometric measurement of the duration of advertising effect on sales. *Journal of Marketing Research*, pages 345–357, 1976.
- [27] Leendert Marinus Koyck. *Distributed Lags and Investment Analysis*, volume 4. North-Holland Publishing Company, 1954.
- [28] Toby J Mitchell and John J Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.
- [29] Edward I George and Robert E McCulloch. Approaches for bayesian variable selection. *Statistica sinica*, pages 339–373, 1997.
- [30] Jon Wakefield. *Bayesian and frequentist regression methods*. New York, NY: Springer, 2013.

## A MCMC Convergence Diagnostics

In order to assess the convergence of the MCMC scheme described in Section 2.5, we used the Gelman-Rubin convergence statistic  $\hat{R}$ . We report its value for each latent dimension of our best performing model, the RF variant, in Table 4. All values are under 1.1, confirming correct convergence. In addition, we display the trace plots for each variable at Figure 7.

|             |           |           |          |           |          |          |           |
|-------------|-----------|-----------|----------|-----------|----------|----------|-----------|
| coefficient | y_AR      | OOH       | TV       | ONLINE    | SEARCH   | RADIO    | Hols      |
| $\hat{R}$   | 0.9997814 | 0.9997621 | 1.00001  | 0.9998523 | 1.000033 | 1.000522 | 0.9999129 |
| coefficient | AVG_Temp  | AVG_Rain  | Unemp_Ix | CC_IX     | Price_IX | Sport_EV |           |
| $\hat{R}$   | 0.999768  | 0.9997925 | 1.030499 | 1.039316  | 1.005669 | 1.000532 |           |

Table 4: Gelman-Rubin statistic results

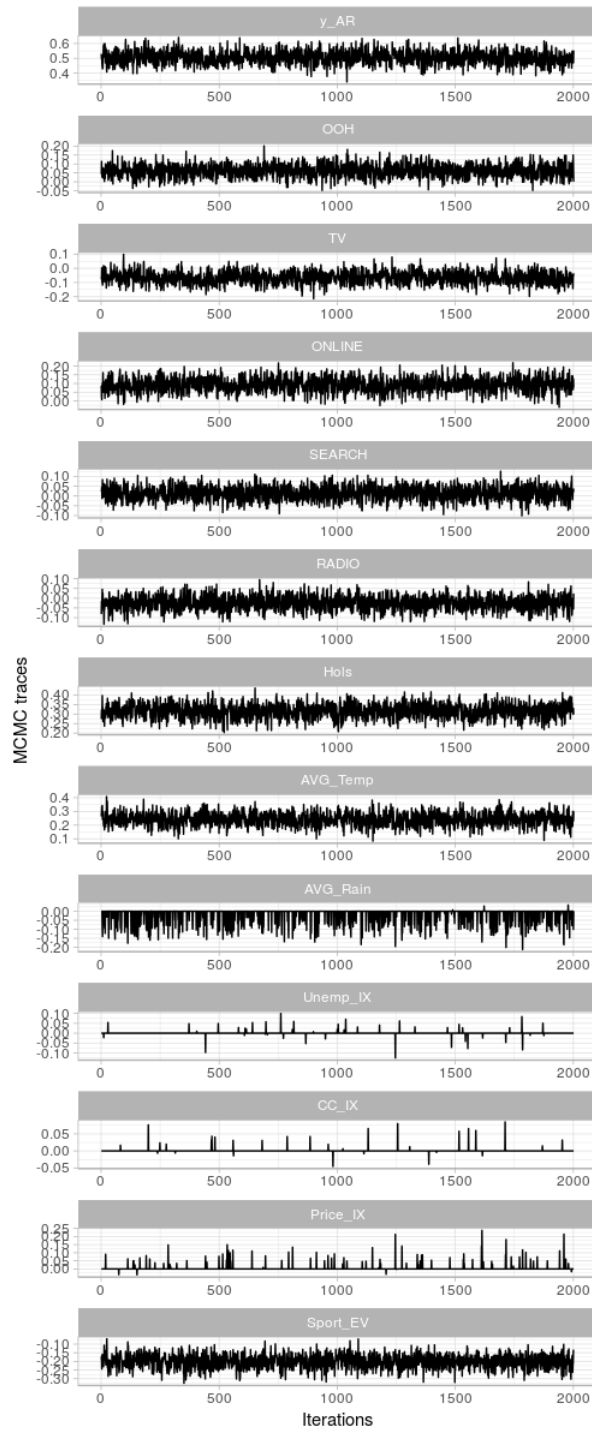


Figure 7: MCMC traces after a burn-in of 2000 iterations.