



ELSEVIER

Journal of Econometrics 81 (1997) 273–280

**JOURNAL OF  
Econometrics**

## Recognizing changing seasonal patterns using artificial neural networks

Philip Hans Franses<sup>a,\*</sup>, Gerrit Draisma<sup>a</sup>

<sup>a</sup> *Econometric Institute, Erasmus University Rotterdam, NL-3000 DR Rotterdam, Netherlands*

---

### Abstract

In this paper we propose a graphical method based on an artificial neural network model to investigate how and when seasonal patterns in macroeconomic time series change over time. Neural networks are useful since the hidden layer units may become activated only in certain seasons or periods, and since this activity can be stepwise or smooth. The graphical method is based on the partial contribution of the hidden layer units to the overall fit. We apply our method to quarterly Industrial Production in France and Netherlands. © 1997 Elsevier Science S.A.

*Key words:* Pattern recognition; Neural networks; Seasonality

*JEL classification:* C22

---

### 1. Introduction

Several recent studies have focused on changing seasonality in macroeconomic time series. Hylleberg et al. (1990) (HEGY) test for the presence of seasonal unit roots, which correspond to stochastic trends at the seasonal frequencies, and hence reflect that seasonal patterns change permanently because of certain shocks. Canova and Hansen (1995) propose a test for constant seasonality versus the alternative hypothesis of seasonal stochastic trends. A common property of these two approaches is that changing seasonality is reflected by test statistics, and that

---

\*Corresponding author.

This paper was presented at an Institute of Economics seminar in Aarhus (November 1994) and at the (EC)<sup>2</sup> conference in Berlin (December 1994). We thank the participants for their helpful comments. Special thanks go to Herman Bierens, Svend Hylleberg, Johan Kaashoek, Halbert White, and to Helmut Lütkepohl (the Associate Editor) and two anonymous referees for their helpful suggestions. The first author thanks the Royal Netherlands Academy of Arts and Sciences for its financial support. The Gauss program for parameter estimation can be obtained via e-mail (franses@few.eur.nl).

no statements are made about how and when seasonal fluctuations change. For example, it is not investigated if seasonal patterns change only slowly, if the observations in only a few or in all seasons change, or if these changes correspond to an influential event like an oil crisis. There are two recent exceptions to the above approaches. Canova and Ghysels (1994) investigate whether parameter estimates for seasonal dummy variables obtain different values in different (a priori defined) business cycle periods. Secondly, Lin and Teräsvirta (1994) use a smooth transition regime-switching model where part of the regression output gives an indication of possible regimes, without fixing these regimes a priori. Lin and Teräsvirta (1994) use seasonal dummy variables to describe seasonality, and their regimes are characterized by functions of deterministic trend variables. In the present paper, we build upon the Lin and Teräsvirta approach by using single hidden layer artificial neural networks in order to additionally allow for stochastic changing seasonality. These models can capture a large part of the variation of any time series and can therefore be useful for recognizing, e.g., regime-switches and structural breaks. We focus on the partial contribution of the hidden layer cells to the overall fit. This contribution is presented in graphs which can yield visual evidence of how and when seasonality has changed.

The outline of this paper is as follows. In Section 2, we discuss augmented single hidden layer neural networks for the purpose of recognizing changing seasonality. In Section 3, we apply our approach to two Industrial Production series. We find that changing seasonality tends to occur only in a few and not in all seasons, and that sometimes these changes can occur quite deterministically. In Section 4, we conclude this paper with a discussion.

## 2. Artificial neural networks

A useful modelling framework that gives insight into how and when possible changes in time series patterns occur should be quite flexible. It may have a nonlinear component since we wish to allow for slowly changing (seasonal) fluctuations, possibly confined to only a few seasons. The model should also give a close fit in order to be able to detect changes with a high degree of precision. We propose to consider the so-called augmented single hidden layer neural network, see, e.g., Kuan and White (1994) and Ripley (1994). In the case of a univariate quarterly time series  $y_t$ , such a model can be written as

$$y_t = \mu_t + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + \sum_{j=1}^q \beta_j G(\phi_{0j}^* \mu_t + \phi_{1j}^* y_{t-1} + \cdots + \phi_{pj}^* y_{t-p}) + \varepsilon_t, \quad (1)$$

where  $\mu_t$  contains deterministic terms. It is common practice to define the activation function as  $G(a) = (1/(1 + \exp(-a)))$ . The number of hidden layer cells is

$q$ , and the number of parameters to be estimated in (1) is  $p + m + (p + m + 1)q$ , where  $m$  is the number of parameters corresponding to  $\mu_t$ . Notice that the  $G(a)$  function can approximate a linear function for certain parameter combinations. Hence, one may want to abstain from the linear part in (1). However, we are interested in the partial contribution of the nonlinear component, and therefore we decide to use (1). Finally, it is clear from the nonlinear component that it is difficult to interpret the values of the individual parameters  $\beta_j$  and  $\phi_{ij}^*$ ,  $i = 0, 1, \dots, p$ . It seems however possible to interpret the output  $\beta_j G(\cdot)$  of the  $q$  hidden layer units, and it is exactly this opportunity that we exploit in this paper for the purpose of recognizing changing seasonal patterns.

Neural network models can be shown to approximate an observed time series arbitrarily well; see Hornik et al. (1989). This implies that also time-series processes with time-varying parameters can be described using models like (1). Hence, there seems no need to modify (1) to allow for, e.g., periodic variation in the  $\phi_i$  parameters,  $i = 1, 2, \dots, p$ . Moreover, time-varying parameter values across the sample are also allowed since the  $G(\cdot)$  function can easily cope with such parameter variation. Of course, model (1) is fairly general and it can be simplified, e.g., by including only  $\mu_t$  in the  $G(\cdot)$  function. The latter is not pursued here since we wish to allow for possible time-varying dynamics.

Model (1) has a linear part and a nonlinear part, where the latter part can be made more important by increasing  $q$ . The penalty for this flexibility is that the total number of parameters can become quite large. Typically, to choose the number of hidden layer units  $q$ , one uses the Schwarz (SC) model selection criterion, which among all such statistical model selection criteria is the criterion which penalizes too many parameters most heavily.

A model closely related to (1) is the smooth transition regression model in Lin and Teräsvirta (1994). In our notation it reads as

$$y_t = \mu_t + \phi_1 y_{t-1} + \dots + \phi_p y_{t-p} + (\phi_0^* \mu_t + \phi_1^* y_{t-1} + \dots + \phi_p^* y_{t-p}) F(t) + \varepsilon_t, \quad (2)$$

where  $F(t)$  is a logistic function of (powers of) time trend  $t$ . The main difference between (2) and (1) is that the argument of  $F$  contains functions of  $t$ , while that of  $G$  in (1) contains the complete set of explanatory variables. If (2) is used to describe changing seasonality, one assumes that these changes are deterministic.

The neural network part of (1) takes care of the nonlinearity in the relation between the dependent variable and the explanatory variables. The hidden layer cells may change activity gradually over time or instantaneously, when their input passes a certain threshold. This threshold activity feature allows us to investigate whether the linear part, consisting of deterministic terms and lagged  $y_t$  variables, is not sufficient to describe one or more observations. This insufficient ability to describe  $y_t$  may only be temporarily because of an outlier, i.e., a hidden layer cell

is then only active for one observation, but such a cell may also be active in a certain period, suggesting the possibility of structural breaks. Of course, one may calculate the output of hidden layer cells for each of the seasons to investigate whether breaks occur in a few or in all seasons. When a hidden layer cell becomes active only gradually, one may conclude that model or parameter changes occur only gradually. When the output of a hidden layer cell is constant throughout the sample, it is clear that it can not be distinguished from deterministic terms.

We estimate the parameters in (1) by minimizing the residual sum of squares (RSS) using a combination of two minimization algorithms: the simplex method of Nelder and Mead, and the Broyden–Fletcher–Goldfarb–Shanno algorithm as implemented in the Gauss optimization toolbox (version 3.0). In the first rounds of optimization we impose a penalty on the size of the parameters  $\beta_j$  and  $\phi_{ij}^*$ , for  $i=0, \dots, p$ , as suggested by Ripley (1994). In the last round, the penalty on the  $\phi_{ij}^*$  parameter is removed. We experiment with five different starting-values to start up the estimation routine. The estimation procedure takes about 30 min on a 486 computer for each model.

Finally, in this paper we depict and interpret the estimated output of the hidden layer cells, i.e.

$$\hat{\beta}_j G(\hat{\phi}_{0j}^* + \hat{\phi}_{1j}^* y_{t-1} + \dots + \hat{\phi}_{pj}^* y_{t-p}). \quad (3)$$

One can also consider the neural network when it is applied to the residuals of the regression of  $y_t$  on explanatory variables. Hence, one may consider the model in a Lagrange Multiplier testing context; see Bierens (1990) and Lee et al. (1993) for such testing approaches. Some of our (unreported) experimentation with simulated and real-life data reveals that this approach does often yield similar conclusions. An additional argument for using (1) for  $y_t$  instead of for residuals is that the initial regression may smooth away small changes in parameters and model structure.

### 3. Empirical results

In this section we investigate the usefulness of model (1) to indicate how and when seasonal patterns change in the (logs of the) quarterly Industrial Production index in France and The Netherlands, observed for the period 1969.1–1992.3. The data are extracted from the OECD Main Economic Indicators (1985 = 100). To investigate the presence of changing seasonal patterns in our data, we apply the tests proposed in HEGY for seasonal unit roots, and we find evidence for such roots in both time series.

We transform our time series  $x_t$  to the interval [0,1] as follows:  $y_t = (x_t - \min(x_t)) / (\max(x_t) - \min(x_t))$ . We do not difference the  $y_t$  using a first- or fourth-order differencing filter since we do not know a priori which filter is required to obtain stationarity. The possible presence of seasonal and nonseasonal unit roots




Fig. 1. Output of hidden layer units, per quarter: France (dotted line is time series, straight line is output hidden layer).

will appear in the lag polynomials in (1). It may however be that the empirically selected value of  $q$  is affected by the presence of unit roots. An investigation of such effects is considered outside the scope of this paper.

We estimate the parameters of neural network models as in (1) where we set  $q$  equal to 0, 1, 2 or 3. Hence, we also consider the linear model for which  $q$  equals 0. We select the model with that value of  $q$  for which the SC criterion is minimized. Furthermore, we set  $p$  in (1) equal to 4, and we assume that  $\mu_t$  contains four seasonal dummy variables and a deterministic trend. As could be expected, the fit of the empirical models is exceptionally good with  $R^2$  values

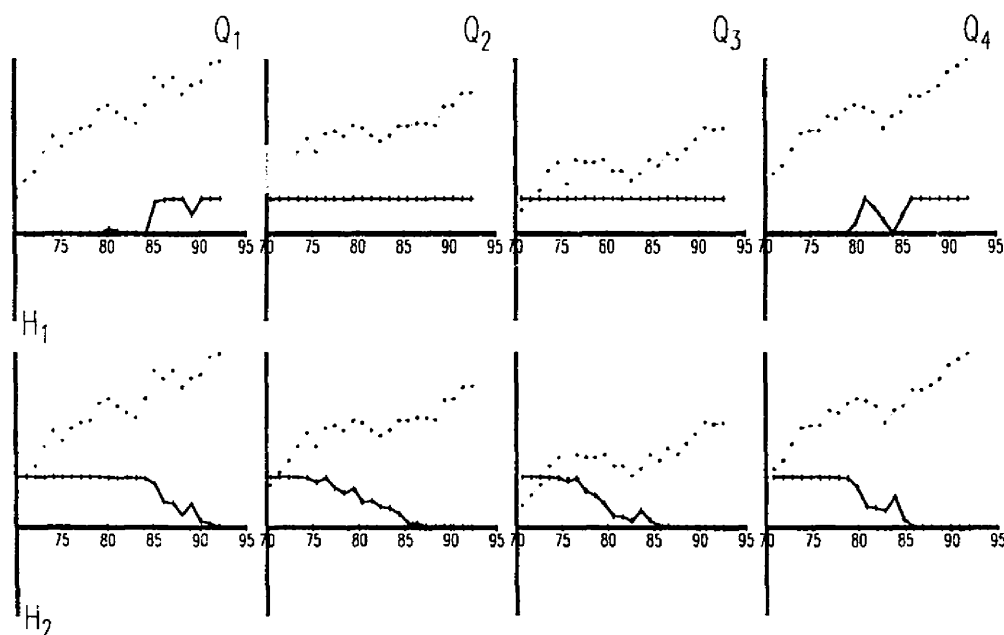


Fig. 2. Output of hidden layer units, per quarter: Netherlands (dotted line is time series, straight line is output hidden layer).

close to 1. For  $q = 0, 1, 2, 3$ , the SC values for France are  $-609.1, -621.9, -631.9$ , and  $-647.6$ , and for Netherlands these are  $-531.0, -536.4, -550.1, -515.3$ . These results indicate that  $q = 3$  for France and  $q = 2$  for Netherlands.

In Figs. 1 and 2 we display the output of the hidden layer cells (denoted as  $H_1, H_2$  and  $H_3$ ) in each of the four quarters ( $Q_1$  to  $Q_4$ ) for the two Industrial Production variables. For France, in Fig. 1, we observe that the output of the first hidden layer cell ( $H_1$ ) seems roughly constant over time, except for quarter 3 for which we observe a gradual increase.  $H_2$  only seems to be activated by single observations, and hence may reflect outlying data points around 1975.1 and 1974.4.  $H_3$  is only active in quarters 3 and 4, where it seems that it takes a constant value until 1979 and a different constant value after 1982. Hence, the output of the last hidden layer cell may be interpreted as indicating a structural shift in the beginning of the eighties in quarters 3 and 4.

The graphs in Fig. 2 for Netherlands indicate that the seasonal patterns seem to change only quite slowly over time. Notice that these graphs are similar to Fig. 4 in Lin and Teräsvirta (1994) where the transition function  $F(t)$  is depicted for a model for  $y_t - y_{t-1}$ . Our  $H_1$  is only active in quarter 1 and 4, where especially after 1985 this activity remains constant.  $H_2$  yields somewhat more complicated output patterns. For the first quarter the output gradually decreases after 1985,

while it does so after 1975 for the second and third quarter. The  $Q_4$  output only starts decreasing after 1980. In sum, the seasonal patterns for Industrial Production in Netherlands change gradually between 1975 and 1985. Lin and Teräsvirta (1994) report results in their Fig. 3 which are also similar to ours, except that their results are obtained using model (2) and for  $y_t - y_{t-4}$  as the dependent variable. The additional information we obtain from our graphs is that the changes do not appear to occur simultaneously in all seasons.

In Franses and Draisma (1995), we present additional results for Industrial Production indices of Germany, Italy, Japan, UK and the USA. The SC values suggest the usefulness of (1), except for Japan where  $q$  can be set equal to 0. Graphs as in Figs. 1 and 2 yield results which can be summarized as follows. For Germany we find that seasonal patterns change around 1980 quite rapidly for the second and third quarter. For Italy, the changes in seasonal fluctuations may have occurred around 1978 and 1985 in the second and third quarter. The seasonal pattern in the UK data appears to change between 1975 and 1981, and this change may have occurred in all seasons. Finally, for the USA we find the single hidden layer cell only to be active in the third quarter after around 1983, and hence we find that changing seasonality is not confined to certain business cycle stages.

#### 4. Discussion

In this paper we have illustrated that the partial contribution of hidden layer units in neural network models can convey visual information on how and when seasonal fluctuations change. We have observed that the observations often change in only one or a few seasons. Furthermore, we have concluded that such changes can be gradually, and only sometimes quite rapidly.

Neural network models can be useful in their own right for out-of-sample forecasting and pattern recognition. However, one may also want to consider more parsimonious models that can generate similar features as put forward by the neural networks and that can be given economic interpretation. For univariate time series analysis, our graphs of the output of hidden layer units may suggest how and when to allow for structural shifts in parameters. A gradual change in one or more seasons may lead to the construction of innovative outlier models. In turn, such models may indicate a trade-off between seasonal unit roots and seasonal mean shifts.

Changing seasonal patterns may also occur when time-invariant models are used while periodic models would have been more useful. When a model as  $y_t = \alpha_s y_{t-1} + \varepsilon_t$  is the data generating process, where the  $\alpha_s$  parameter can take different values in different seasons, one may expect that a nonperiodic model for  $y_t$  can result in systematic forecasting errors in various seasons. Within the neural network model, these errors will be taken care of by the hidden layer units.

The finding in Franses (1993) that Industrial Production in Netherlands can be described by a periodic autoregression seems to correspond with the reported empirical findings.

When hidden layer units become only active in some seasons and not in others, this may have an implication for seasonally adjusting a time series. A dominant component in most seasonal adjustment procedures is a linear two-sided moving average filter with weights that sum to one. Such filters are applied to all seasons, and hence, when observations in only one season change, one may expect inappropriate adjustment in other seasons.

There are several suggestions for further research. An extensive Monte Carlo experiment of the properties of neural network models in case of (seasonal or nonseasonal) mean shifts seems useful. Our method may be turned into a diagnostic testing format similar as in Bierens (1990) and Lee et al. (1993) or in order to perform simulation experiments. A further research topic is to study the impact of seasonal and nonseasonal unit roots on the number of hidden layer cells. This topic is related to analyzing stationarity and unit roots in nonlinear models. Finally, an application of our graphical approach to large sets of seasonal and nonseasonal time series may reveal its practical usefulness for detecting changing time series patterns.

## References

- Bierens, H.J., 1990. A consistent conditional moment test of functional form. *Econometrica* 58, 1443–1458.
- Canova, F., Ghysels, E., 1994. Changes in seasonal patterns: are they cyclical? *Journal of Economic Dynamics and Control* 18, 1143–1171.
- Canova, F., Hansen, B.E., 1995. Are seasonal patterns constant over time? A test for seasonal stability. *Journal of Business and Economic Statistics* 13, 237–252.
- Franses, P.H., 1993. Periodically integrated subset autoregressions for Dutch industrial production and money stock. *Journal of Forecasting* 12, 601–613.
- Franses, P.H., Draisma, G., 1995. Recognizing changing seasonal patterns using artificial neural networks. *Econometric Institute Report* 9514, Erasmus University Rotterdam.
- Hornik, K., Stinchcombe, M., White, H., 1989. Multi-layer feed-forward networks are universal approximators. *Neural Networks* 2, 359–366.
- Hylleberg, S., Engle, R.F., Granger, C.W.J., Yoo, B.S., 1990. Seasonal integration and cointegration. *Journal of Econometrics* 44, 215–238.
- Kuan, C.M., White, H., 1994. Artificial neural networks: an econometric perspective. *Econometric Reviews* 13, 1–91.
- Lee, T.-H., White, H., Granger, C.W.J., 1993. Testing for neglected nonlinearity in time series models: a comparison of neural network methods and alternative tests. *Journal of Econometrics* 56, 269–290.
- Lin, C.F., Teräsvirta, T., 1994. Testing the constancy of regression parameters against continuous structural change. *Journal of Econometrics* 62, 211–228.
- Ripley, B.D., 1994. Neural networks and related methods for classification (with discussion). *Journal of the Royal Statistical Society B* 56, 409–456.