Computational Statistics and Data Analysis 70 (2014) 212–226

Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

Dynamic seasonality in time series

Mike K.P. So *, Ray S.W. Chung

Department of Information Systems, Business Statistics and Operations Management, The Hong Kong University of Science and Technology, Clear Water Bay Road, Kowloon, Hong Kong

ARTICLE INFO

Article history:
Received 13 November 2012
Received in revised form 21 June 2013
Accepted 13 September 2013
Available online 21 September 2013

Keywords:
Dynamic seasonality
Financial time series
GARCH models
Conditional heteroskedasticity
Model selection

ABSTRACT

This study introduces a new class of time series models capturing dynamic seasonality. Unlike traditional seasonal models that mainly focus on the mean process, our approach accommodates dynamic seasonality in the mean and variance processes. This feature allows us to statistically infer dynamic seasonality in heteroskedastic time series models. Quasi-maximum likelihood estimation and a model selection procedure are adopted. A simulation study is carried out to evaluate the efficiency of the estimation method. In the empirical examples, our model outperforms a deterministic seasonality model and Holt–Winters method in forecasting monthly Nino Region 3 Sea Surface Temperature Index and intraday stock return variations in an out-of-sample analysis.

1. Introduction

Modeling seasonality in time series is an important topic in statistics. Omitting seasonality from a model may induce substantial bias. For example, if we want to model the average monthly temperature, ignoring the seasonal component can lead to huge bias in predicted values. In practice, many time series exhibit seasonal patterns in various forms. Monthly tourist arrivals and monthly electricity demand are examples of time series subject to seasonality (Soares and Medeiros, 2008; Chang et al., 2009; Taylor, 2010b). Practitioners commonly employ either dummy variables or the Fourier transform to model seasonality. The former method relies on introducing a set of indicator variables to capture the seasonality in each season, and the latter method involves fitting a Fourier series to capture the curve of seasonality. These two methods were explained by Soares and Medeiros (2008).

The methods described above assume that the seasonal component within the same season is time invariant. Some practitioners are also interested in methods that allow the seasonal component within each season to change over time. This paper refers to this phenomenon as *dynamic seasonality*. The Holt–Winters (HW) method, first proposed by Holt (1957) and made popular by Winters (1960), is commonly adopted for dynamic seasonality modeling. It is based on the application of an exponential smoothing mechanism to capture seasonality dynamically. Williams (1987) modified the HW method to allow the smoothing parameters to evolve over time. Ord et al. (1997) incorporated a state space model framework into the HW method. With this framework, the HW method can be taken as parametric, and hence more sophistical statistical inference like interval estimation is feasible (Koehler et al., 2001). More details of the above discussion can be found in Hyndman et al. (2008b). Taylor (2010a,b) extended the HW method to seasonality over multiple cycles. The HW method is appealing because exponential smoothing is a simple, robust, and effective method (Chatfield and Yar, 1988).

Practitioners are also interested in using parametric models to capture dynamic seasonality. One such model is the seasonal autoregressive integrated moving average (ARIMA) class of model. This type of model captures the dynamics of

* Corresponding author. Tel.: +852 23587726.
E-mail addresses: immkpso@ust.hk (M.K.P. So), swchungaa@ust.hk (R.S.W. Chung).

0167-9473/$ – see front matter
http://dx.doi.org/10.1016/j.csda.2013.09.010

seasonality through the Box–Jenkins mechanism. The seasonal ARIMA class of model utilizes linear structure to capture the autocorrelation of the time series. Box et al. (1994) set out the details of seasonal ARIMA models. Another common type of parametric model for dynamic seasonality is the periodic autoregressive (PAR) model class. This type of model can be regarded as a generalization of seasonal autoregressive models. PAR models assume that, because observations in different seasons have different autoregressive structures, the time series in each season has its own set of parameters. Studies such as those of Franses (1995) have suggested that, because it may not be possible to separate seasonality from the stochastic trend, PAR models, which do not require this assumption, can help us model seasonal time series with a stochastic trend. Franses and Paap (2004) describe PAR models in more detail.

Although it is easy to implement the HW method, it assumes that all seasons have the same set of smoothing parameters. Hyndman et al. (2008a) noted that the HW method can be represented in a state space model in which the seasonal component has a seasonal unit root. Therefore, the HW method may not be suitable for periodically stationary time series. Because of the restriction whereby the seasonal component and the stochastic trend are inseparable, neither seasonal ARIMA models nor PAR models can return an estimate of the seasonal component. If we want to undertake further analysis based on the seasonal component (e.g., a policy-maker may want to examine seasonal variation of the unemployment rate through the seasonal component), seasonal ARIMA models and PAR models may not be the best choices. This restriction also makes it difficult to interpret the dynamics of these models. Moreover, the seasonality considered in the literature is usually considered as seasonality in mean. However, the seasonality of financial time series is often reflected in the fluctuation of the process. For example, if we plot the autocorrelation function (ACF) of any equity return, we will not find a seasonal pattern. Nevertheless, if we instead plot the ACF of the squared return as a proxy of the variance of the return random variable, a clear seasonal pattern will be revealed. This type of seasonality, which is associated with the variance process, is regarded as seasonality in variance. Bildik (2001) illustrated in detail the possible causes of seasonality in the variance observed in financial time series. Giot (2000, 2005) introduced a decomposition to extract seasonality in variance and used the average squared return of each season to model the corresponding seasonal component. Taylor and Xu (1997) and Chang and Taylor (2003) incorporated dummy variables into a GARCH model, one of the most commonly used classes of volatility models for financial time series introduced by Engle (1982) and Bollerslev (1986), to capture seasonality in variance. Although several methods deal with seasonality in variance deterministically, only a few of them can handle dynamic seasonality in variance. One example is the periodic GARCH (PGARCH) model proposed by Bollerslev and Ghysels (1996). However, the PGARCH model has some limitations. For instance, it can be used for forecasting but cannot return an estimate of the seasonal component. We have two main objectives in this paper. First, we introduce a new class of periodically stationary dynamic seasonality in the mean models to alleviate shortcomings found in existing methods. Second, we develop a new model to account for dynamic seasonality in variance.

The remainder of this paper is organized as follows. Section 2 introduces the dynamic seasonality models developed in this paper. Section 3 outlines the quasi-maximum likelihood estimation process and a model selection procedure that assist us in finding a suitable dynamic seasonality model to fit the data. Section 4 outlines a simulation study, and Section 5 gives two empirical examples to demonstrate the advantages of considering dynamic seasonality. Section 6 concludes the paper.

2. Dynamic seasonality models

2.1. Seasonality in the mean

This section introduces the structure of our dynamic seasonality in the mean (DSM) model, a model that takes account of the dynamic nature of the seasonality in mean. Let y_t, $t = 1 \ldots, n$, be the time series of interest. Assume that y_t is periodically stationary and that d seasons are identified in the time series. Let s_t be the seasonal factor at time t. Following Hylleberg (1986) and Bell and Hillmer (1992), we consider the additive decomposition of y_t as follows:

$$y_t = \mu_t + s_t + \sigma_t \varepsilon_t, \quad \varepsilon_t \sim D(0, 1), \tag{1}$$

where μ_t is the trend of y_t, σ_t^2 is the conditional variance given $\Im_{t-1}$ (the information on y up to time $t-1$), ε_t is independent innovation,and $D(0, 1)$ is a distribution with mean 0 and variance 1. The components s_t, μ_t, and σ_t^2 are assumed to be $\Im_{t-1}$-measurable. We argue that any possible autocorrelation in ε_t can be explained by the ARMA structure embedded in μ_t, and thus assume ε_t to be independent. This additive decomposition is common in modeling seasonality in the mean. It is flexible enough to accommodate models with seasonality in the multiplicative form instead of the additive form by taking a logarithmic transformation (Bell and Hillmer, 1992), because we allow the distribution of the error ε_t to be unspecified. As a result, the representation given by (1) can describe many kinds of seasonality in the mean.

The decomposition of y_t in (1) gives $E[y_t|\Im_{t-1}] = \mu_t + s_t$. In other words, $y_t - \mu_t$ provides information relevant to the seasonal factor at time t. Therefore, we propose to model the seasonal factor s_t by

$$\begin{aligned} s_t &= \sum_{k=1}^{d} I_{t,k}[\delta_{k,0} + \delta_{k,1}(y_{t-d} - \mu_{t-d}) + \delta_{k,2}s_{t-d}], \\ \delta_{k,0} &> 0, \qquad \delta_{k,1}, \delta_{k,2} \geq 0 \quad \text{and} \quad \delta_{k,1} + \delta_{k,2} < 1, \end{aligned} \tag{2}$$