International Journal of Forecasting 34 (2018) 598–611

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Markov-switching dynamic factor models in real time

Maximo Camacho [a,*], Gabriel Perez-Quiros [b,c], Pilar Poncela [d]

[a] *Universidad de Murcia, Spain*
[b] *Banco de España, Spain*
[c] *CEPR, United Kingdom*
[d] *European Commission, Joint Research Centre and Universidad Autónoma de Madrid, Spain*

A R T I C L E   I N F O

Keywords:
Business cycles
Output growth
Time series

A B S T R A C T

We extend the Markov-switching dynamic factor model to account for some of the specificities of the day-to-day monitoring of economic developments from macroeconomic indicators, such as mixed sampling frequencies and ragged-edge data. First, we evaluate the theoretical gains of using data that are available promptly for computing probabilities of recession in real time. Second, we show how to estimate the model that deals with unbalanced panels of data and mixed frequencies, and examine the benefits of this extension through several Monte Carlo simulations. Finally, we assess its empirical reliability for the computation of real-time inferences of the US business cycle, and compare it with the alternative method of forecasting the probabilities of recession from balanced panels.

1. Introduction

The slump of 2008–2009 was the most sustained economic slump that the United States has weathered since World War II. One of the lessons that the Great Recession left for economists was that policymakers and business people, who had become accustomed to the serene conditions of the Great Moderation, have dramatically increased their interest in determining as quickly as possible when the economy has suffered from a business cycle phase shift. In this context, time series models, which can automate the increasing complexity of the signal extraction problem in economics, help economic agents to produce and update their real-time views of any developments in economic activity. These models deal with economic indicators that share the two properties of the business cycle that were documented early on by Burns and Mitchell (1946): their signals about economic developments are spread over the different aggregates, and they exhibit business cycle asymmetries.

Diebold and Rudebusch (1996) were the first to suggest a unified model that captures these two business cycle features from a set of economic indicators. They argued that comovements among the individual economic indicators can be modelled by using the linear coincident indicator approach described by Stock and Watson (1991), while the existence of two separate business cycle regimes can be modelled using the Markov-switching specification advocated by Hamilton (1989). Integrating these approaches, Chauvet (1998), Kim and Nelson (1998) and Kim and Yoo (1995) combined the dynamic-factor and Markov-switching frameworks to propose various different versions of statistical models that capture comovements and regime shifts simultaneously. Camacho, Perez Quiros, and Poncela (2015) find that the fully non-linear multivariate specification outperforms the "shortcut" of using a linear factor model to obtain a coincident indicator that is then used to compute the Markov-switching probabilities. More recently, Chauvet and Hamilton (2006), Chauvet and Piger (2008), and Hamilton (2011) examined the empirical reliability of these models for computing real-time inferences of the US business cycle states.

* Corresponding author.
E-mail addresses: mcamacho@um.es (M. Camacho), gabriel.perez@bde.es (G. Perez-Quiros), pilar.poncela@ec.europa.eu (P. Poncela).

https://doi.org/10.1016/j.ijforecast.2018.05.002
0169-2070/

One important limitation of these Markov-switching dynamic factor models (MS-DFM) is that they were designed originally for dealing with balanced panels of business cycle indicators. This crucial assumption means that MS-DFM exhibit two drawbacks when applied to the (timely) day-to-day monitoring of economic activity in real time. The first drawback is that the typical lack of synchronicity in the flow of macroeconomic information implies that some indicators are published with a time delay, which requires one to deal with unbalanced panels of data. Failing to account for this publication pattern would imply that the users of traditional MS-DFM who develop early assessments of economic developments from balanced panels of data will unavoidably incur one of the following two substantial costs. The first cost appears when the forecasts are made from the latest available balanced panel. In this case, the forecasts lose the latest and most valuable information that is contained in the promptly-issued indicators at the time of the assessments. The second cost is that of being late, when the analysts decide to wait until all of the business cycle indicators become available, and the inferences then actually refer to the past. A significant example of this limitation is the two-month lag in the reporting of the real-time recession probability chart released by the St Louis Fed, which uses the MS-DFM originally developed by Chauvet (1998).

The second drawback of standard MS-DFM is that they consider variables sampled at only a single frequency. In practice, although some of the macroeconomic indicators that are observed for inferring business cycle states are sampled quarterly, others, which may potentially be useful in real-time inferences, are sampled at a higher frequency. For example, the National Bureau of Economic Research (NBER) Dating Committee acknowledges that recessions are defined as significant declines in economic activity that are normally visible in the real Gross Domestic Product (GDP), which is available quarterly, and the real income, employment, industrial production, and wholesale-retail sales, which are available monthly.

This paper examines the extent to which the incoming information provided by new releases of promptly published economic indicators, potentially sampled at different frequencies, could help to improve real-time inference about the business cycle. Using a theoretical MS-DFM, we show the extent to which inferences about the state of the economy can be improved upon by including indicators that are available early. Interestingly, we find that the improvements in performance depend on the factor loadings, the idiosyncratic variances, the dynamics of the common factor, and the differences between the means in the business cycle states.

Next, we extend the MS-DFM to allow economic agents who track business cycle developments in real time to use whatever business cycle economic indicators they wish, regardless of their publication delays, and of whether they are sampled monthly or quarterly. Based on a Markov-switching extension of the linear dynamic factor model proposed by Mariano and Murasawa (2003), our procedure deals with missing observations by using a time-varying nonlinear Kalman filter. Whenever the data are not observed, the missing observations are replaced by random draws from a variable whose distribution cannot depend on the parameter space that characterizes the Kalman filter. The corresponding row in the Kalman recursion is then skipped, and the measurement equation for the missing observation is set to the random choice.

By means of several Monte Carlo experiments, we measure the magnitude of the gains from using our extension to compute business cycle inferences. We show that our proposal outperforms the MS-DFM that requires balanced panels, especially when the forecasting horizon increases, when the two states are separated well, and when the variance and inertia of the idiosyncratic components are low. Finally, we use a real-time data set to show that our extension of the MS-DFM leads to significant improvements in computing real-time business cycle inferences relative to forecasting from balanced and/or lagged panels of indicators using the four constituent monthly series of the Stock-Watson coincident index. Notably, we show that adding the GDP does not seem to produce significant improvements in this setting.

The structure of this paper is organized as follows. Section 2 assesses the real-time features of the dataflow within a factor model framework. Section 3 examines the relative performance gains of dealing with ragged-edge data through a Monte Carlo experiment. Section 4 illustrates these results for US real-time data. Section 5 concludes.

2. The model

2.1. Model features

Our framework is the single-index Markov-switching dynamic factor model that was proposed in the mid-nineties by Chauvet (1998), Kim and Nelson (1998), and Kim and Yoo (1995), which incorporates both comovements and business-cycle shifts into a statistical model. The model postulates that a vector of N economic indicators, $\boldsymbol{y}_t = (y_{1,t}, \ldots, y_{N,t})'$, which are hypothesized to move contemporaneously with the overall economic conditions, can be decomposed as the sum of two components. The first component is a linear combination of r unobserved factors, $\boldsymbol{f}_t = (f_{1,t}, \ldots, f_{r,t})'$, which accounts for the common comovements. The second component is the $N \times 1$ time series vector $\boldsymbol{u}_t$, which represents the idiosyncratic movements in the series. This suggests the formulation:

$$\boldsymbol{y}_t = \boldsymbol{\Lambda} \boldsymbol{f}_t + \boldsymbol{u}_t, \tag{1}$$

where $\boldsymbol{\Lambda}$ is the $N \times r$ factor loading matrix and $\boldsymbol{u}_t$ is the vector of idiosyncratic components.

We account for the business cycle asymmetries by assuming that the dynamic behavior of the common factors is governed by an unobserved regime-switching state variable, s_t. Within this framework, one can label $s_t = 0$ and $s_t = 1$ as the expansion and recession states at time t. In addition, it is standard to assume that the state variable evolves according to an irreducible two-state Markov chain whose transition probabilities are defined by

$$\begin{aligned} &p(s_t = j \mid s_{t-1} = i, s_{t-2} = h, \ldots, I_{t-1}) \\ &\quad = p(s_t = j \mid s_{t-1} = i) = p_{ij}, \end{aligned} \tag{2}$$