International Journal of Forecasting 30 (2014) 550–553

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Discussion

Discussion of the paper: The way out of recessions: A forecasting analysis for some Euro area countries

Gianni Amisano

European Central Bank, Germany

1. Introduction

The paper starts from the premise that standard SETAR models might not be appropriate for describing cyclical output dynamics, due to the fact that, in observed cyclical fluctuations, the speed of the recovery is affected by the intensity of the recession. This feature cannot be captured by a standard two-state SETAR model, but it can be given an economic interpretation as arising from the process of creative destruction taking place during recessions (the cleansing effect) and the consequent interaction between new entrepreneurial energies coming into play and the sluggishness of demand adjustments.

As a solution to this problem, the authors use the bounce-back function (BBF) mechanism proposed by Kim, Morley, and Piger (2005). The resulting model allows the local mean in a SETAR model to be affected by the duration and the intensity of the recession phase, as described by Eq. (4) in the paper, reproduced below for convenience:

$$\mu_t = \gamma_0(1 - s_t) + \gamma_1 s_t + \lambda_1 s_t \sum_{j=\ell+1}^{\ell+m} s_{t-j} + \lambda_2(1 - s_t) \sum_{j=\ell+1}^{\ell+m} s_{t-j} + \lambda_3 \sum_{j=\ell+1}^{\ell+m} \Delta y_{t-j-1} s_{t-j}. \quad (4)$$

This is a flexible specification that nests several special cases: (1) the simple SETAR specification, when $\lambda_1 = \lambda_2 = \lambda_3 = 0$; (2) the U-shaped pattern, when $\lambda_1 = \lambda_2$; and (3) a time homogeneous pattern depending on the regime duration alone, irrespective of episode-specific developments, when $\lambda_3 = 0$. The last term of Eq. (4) is the most interesting one, because it allows the intensity of the recession to affect the strength of the recovery. This requires λ_3 to be significant and negative.

My overall impression is that the paper is very interesting, has a clear and relevant motivation, and is well connected with the recent literature on nonlinear models for the business cycle, and, in addition, the empirical analysis contains interesting results. However, the paper unfortunately has certain shortcomings that, in my view, limit the relevance and usefulness of the results. I explain these features in detail in the remainder of my discussion.

2. Some comments

2.1. Univariate approach

The analysis is univariate and is based on quarterly output growth series for individual euro area countries and the euro area as a whole. This is a particular way to look at business cycle features. By definition, the business cycle is a concept that encompasses the behaviours of several economic indicators, and the main way of studying its properties has traditionally involved looking at the joint behaviours of multiple economic time series. On the other hand, Stock and Watson (forthcoming) study business cycle dating and show how to aggregate signals from univariate models. The application documented in the paper, which is purely univariate, must therefore be considered as providing only partial and preliminary indications, and it would be interesting to see how a VAR extension of the model, aimed at dealing with a set of key macro indicators for each country and/or different countries, would work in practical applications.

E-mail address: gianni.amisano@ecb.europa.eu.

http://dx.doi.org/10.1016/j.ijforecast.2014.03.008

2.2. Are recessions all alike?

When λ_3 in Eq. (1) is equal to zero, the BBF-SETAR specification imposes a pattern that is constant across recession episodes, but different mechanically across countries, since it is based on univariate, country-specific equations. An empirically relevant fact is that recessions tend to vary over time (reflecting the fact that recessions are generated by different kinds of shocks), but to be affected by some kind of homogeneity across different countries. Unfortunately, the λ_3 parameter is insignificant in the empirical analysis documented in the paper, and therefore the resulting preferred model for all countries means that the BBF pattern is only determined by the length of the recession. In my view, this feature is difficult to interpret.

2.3. SETAR versus Markov switching

In their Section 2, the authors state that the SETAR model has two advantages with respect to the Markov Switching (MS) model:

1. it is easier to estimate the SETAR model (via nonlinear least squares), and its outcome does not depend on initialisation, as in the MS case;
2. the driver of regime changes is perfectly observable (the lagged dependent variable), whereas there is an additional latent variable in the MS model.

I do not think that either of these points is relevant in guiding the applied researcher in a choice between a SETAR and an MS model. In particular, as far as point 1 is concerned, for the MS equivalent of the SETAR model used in the paper, namely a model where only the conditional mean changes across regimes, there are very few computational complications, the resulting ML estimation procedure is very tractable, and convergence is very fast. It is only when shock variances differ across regimes that complications arise, and SETAR estimation also becomes more difficult in these cases. In addition, the dependence on initialisation is an issue that has to be investigated in any kind of model, a problem which can usually be solved by running different estimation runs, starting from different initialisations. Finally, on this point, MS models also lend themselves to an easy implementation of the EM algorithm, in much the same way that SETAR estimation breaks down to OLS estimation conditional on the threshold variable.

As regards the second point, the fact that the driver of the change depends on a latent variable (that is, that regime transitions are not deterministic functions of the data and the parameters) is an asset, not a limitation of the model, and allows for additional, useful flexibility. Ultimately, such a lack of flexibility is the reason why the SETAR model has to be augmented with awkward bounce-back factors.

In my view, the serious drawback of the MS model is that, in its simplest version, transitions are completely exogenous, which might not be appealing, whereas in the SETAR model transitions are completely deterministic and entirely endogenous. However, these specifications are both extreme. There is an interesting middle ground, namely that provided by MS models with time-varying transition probabilities as functions of observable indicators. Here, transitions are still partially stochastic but not exogenous, when the variable(s) affecting transition probabilities are (lagged) endogenous variable(s). See Amisano and Fagan (2013) as an example of the way of specifying and estimating these models.

2.4. RMSE comparisons

In their Section 4, the authors describe a thorough forecasting comparison exercise in which the selected SETAR-BBF model is compared to alternative models, in terms of pseudo-out-of-sample forecasting over the 2000Q1–2010Q4 sample (and subperiods within this evaluation period). While it emerges that the favourite BBF model has decent forecasting properties relative to linear and nonlinear alternatives, it is not entirely clear whether the RMSE differences are significant at all. It would have been useful to have tested for RMSE differences using formal tests following Diebold and Mariano (1995) or Giacomini and White (2006). Without this, the reader is left with the suspicion that none of these differences are relevant. In addition, another important element which is missing from the comparison is the relative performance over a longer horizon.

2.5. Density forecast evaluation

Another very good feature of the paper is that the authors extend their assessment to the density forecast properties of the model. This is a good way to go beyond point forecasting in a direction that is useful for actual decisions and realistic uses of forecasting models. The authors construct density forecasts which are reported in the form of 5% and 95% quantiles of the one-step-ahead forecasting distribution at each point in time in the evaluation sub-sample.

This is done in two different ways, namely by fixing parameters at their estimated values and by formally incorporating parameter uncertainty. Then, the authors provide graphs for the resulting quantiles (Figs. 2 and 3) and discuss their results by comparing these quantiles with actual realisations.

Unfortunately, the authors make no attempt to connect to the very active and interesting literature on the evaluation of density forecasts. Corradi and Swanson (2006) provide an interesting review of these techniques. In particular, the authors limit themselves to just checking when the actual values are outside forecasting density quantiles. This should be happening $100 \cdot (1-\alpha)\%$ of the time in any case, even when the model is correctly specified, where α is the size of the predictive interval, and this is the principle behind the Christoffersen (1998) coverage test.

The authors should instead focus on the requirement that the model delivers correct coverage, i.e., the correct calibration of the predictive density. One way to check this would be to use Probability Integral Transforms (PITs). Suppose that model A assigns one-step-ahead predictive