International Journal of Forecasting 33 (2017) 482–501

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Adaptive models and heavy tails with an application to inflation forecasting

Davide Delle Monache[a], Ivan Petrella[b,c,*]

[a] Banca d'Italia, Via Nazionale 91, 00184, Rome, Italy
[b] Warwick Business School, University of Warwick, UK
[c] CEPR, UK

ARTICLE INFO

Keywords:
Adaptive algorithms
Inflation
Score-driven models
Student-*t*
Time-varying parameters

ABSTRACT

This paper introduces an adaptive algorithm for time-varying autoregressive models in the presence of heavy tails. The evolution of the parameters is determined by the score of the conditional distribution, with the resulting model being observation-driven and being estimated using classical methods. In particular, we consider time variation in both the coefficients and the volatility, emphasizing how the two interact with each other. Meaningful restrictions are imposed on the model parameters in order to achieve local stationarity and bounded mean values. The model is applied to the analysis of inflation dynamics with the following results: allowing for heavy tails leads to significant improvements in terms of both the fit and forecasts, and the adoption of the Student-*t* distribution proves to be crucial for obtaining well-calibrated density forecasts. These results are obtained using the US CPI inflation rate and are confirmed by other inflation indicators, as well as for the CPI inflation of the other G7 countries.

1. Introduction

The last two decades have seen an increasing level of interest in the use of models with time-varying parameters (TVP) for the analysis of macroeconomic variables. Stock and Watson (1996) renewed the interest in this area by documenting the widespread forecasting gains of TVP models.[1] Recently, Cogley and Sargent (2005), Primiceri (2005), and Stock and Watson (2007) highlighted the importance of allowing for time variation in the volatility as well as in the coefficients.[2] However, most studies to date have considered TVP models under the assumption that the errors are normally distributed. Although this assumption is convenient, it limits the ability of the model to capture the tail behaviors that characterize a number of macroeconomic variables. As the recent recession has shown, departures from Gaussianity are important for properly accounting for the risks associated with black swans (see e.g. Cúrdia, Del Negro, & Greenwald, 2014).

This paper considers an adaptive autoregressive model where the errors are Student-*t* distributed. Following Creal, Koopman, and Lucas (2013) and Harvey (2013), the variation in the parameters is driven by the score of the conditional distribution. In this framework, the

* Corresponding author at: Warwick Business School, University of Warwick, UK.
E-mail addresses: davide.dellemonache@bancaditalia.it (D. Delle Monache), Ivan.Petrella@wbs.ac.uk (I. Petrella).

[1] Attempts to take into account the well-known instabilities in macroeconomic time series can be traced back to the work of Cooley and Prescott (1973, 1976), Rosenberg (1972), and Sarris (1973).

[2] D'Agostino, Gambetti, and Giannone (2013) highlight the relative forecast accuracy gains of TVP models over traditional constant parameter models in a multivariate setting.

http://dx.doi.org/10.1016/j.ijforecast.2016.11.007
0169-2070/

distribution of the innovations not only modifies the likelihood function (as in the t-GARCH of Bollerslev, 1987, for example), but also implies a different updating mechanism for the TVP. In this regard, Harvey and Chakravarty (2008) highlight the fact that the score-driven model for a time-varying scale with Student-t innovations leads to a filter that is robust to outliers, while Harvey and Luati (2014) show that the same intuition holds true in models with time-varying locations. The resulting model is observation-driven. Unlike parameter-driven models, the parameters' values are obtained as functions of the observations only and the likelihood function is available in closed form, meaning that the model is estimated using classical methods.[3]

As was stressed by Stock (2002) in his discussion of Cogley and Sargent's (2002) study, estimating TVP models without controlling for the possible heteroscedasticity is likely to overstate the time variation in the coefficients (see also Benati, 2007). This paper considers time variation in both the coefficients and the volatility, and emphasizes the interactions between the two in a score-driven model. Moreover, we show how restrictions can be imposed on the model's parameters so as to achieve local stationarity and a bounded long-run mean. These two restrictions, which are used commonly in applied macroeconomics, have not yet been considered in the context of score-driven models.

The adaptive model in this paper is related to an extensive body of literature that has investigated ways of improving forecasting performances in the presence of instability. Pesaran and Pick (2011), Pesaran, Pick, and Pranovich (2013) and Pesaran and Timmermann (2007) focus on optimal weighting schemes in the presence of structural breaks. Giraitis, Kapetanios, and Yates (2014) propose a non-parametric estimation approach for time-varying coefficient models. The weighting functions implied by these models typically decrease monotonically with time, a feature which they share with traditional exponential weighted moving average forecasts (see e.g. Cogley, 2002). Our model features time variation in both location and scale, and Student-t errors. This implies a non-linear filtering process with a weighting pattern that cannot be replicated by the procedures proposed in the literature. The benefit of this approach is that observations that are perceived as outliers, based on the estimated conditional location and scale of the process, have effectively no weight in updating the TVP. The resulting pattern of weights is both non-monotonic and time-varying, since it is a function of the estimated TVP. Therefore, the model implies more rapid updates of the coefficients in periods of high volatility. Furthermore, in periods of low volatility, even deviations from the mean that are not extremely large in absolute terms are more likely to be 'classified' as outliers. As such, they are disregarded by the filter, which is robust to extreme events. These model features are important in the analysis of macroeconomic time series that display instability and changes in volatility, as is demonstrated empirically through an application to inflation dynamics.

Understanding inflation dynamics is key for policy makers. In particular, modern macroeconomic models highlight the importance of forecasting inflation for the conduct of monetary policy (see e.g. Svensson, 2005). There are at least three reasons why our model is particularly suitable for inflation forecasting. First, simple univariate autoregressive models have been shown to work well in the context of inflation forecasting (see Faust & Wright, 2013). Second, Pettenuzzo and Timmermann (2016) show that TVP models outperform constant-parameters models and that models with small/frequent changes, like that proposed in this paper, produce more accurate forecasts than models whose parameters exhibit large/rare changes. Third, while important changes in the dynamic properties of inflation are well documented (see e.g. Stock & Watson, 2007), most of the empirical studies use a Bayesian framework and have a number of shortcomings: (i) they are computationally demanding; (ii) large numbers of draws need to be discarded when restrictions are imposed to achieve stationarity, thus leading to a potentially large inefficiency[4]; and (iii) normally distributed errors are usually assumed. The last point is particularly relevant, as it is well known, at least since the seminal work of Engle (1982), that the distribution of inflation displays non-Gaussian features. The adaptive model presented in this paper tackles all of these shortcomings.

Our model produces reasonable patterns for the long-run trend and the underlying volatility when it is used to analyze inflation. By introducing the Student-t distribution, we make the model more robust to short-lived spikes in inflation (for instance, those in the last part of the sample). At the same time, the specifications with Student-t innovations display substantially more variation in the volatility. In practice, the variance is affected by the outliers less when using Student-t innovations, and can adjust conveniently to accommodate changes in the dispersion of the central part of the distribution. The introduction of heavy tails improves both the fit and the out-of-sample forecasting performance of the model. The density forecasts produced under a Student-t distribution are substantially better than those produced by its Gaussian counterpart, the model of Stock and Watson (2007), or a TVP-VAR with stochastic volatility (see e.g. D'Agostino et al., 2013). In fact, well-calibrated density forecasts are obtained only when we allow for heavy tails. While the baseline analysis focuses on US CPI inflation, which is noticeably noisier and harder to forecast than other measures of inflation, we show that the improvement in density forecasting performance is also obtained for other inflation measures, such as those derived from the PCE and GDP deflators. Given the differences in inflation dynamics between countries (Cecchetti, Hooper, Kasman, Schoenholtz,

3 In the parameter-driven models, the dynamics of the parameters are driven by additional idiosyncratic innovations. Therefore, analytical expressions for the likelihood function are hardly available in closed form, and the use of computationally-intensive simulation methods is usually required (see e.g. Koopman, Lucas, & Scharth, 2016).

4 Chan, Koop, and Potter (2013) and Koop and Potter (2011) deal with local stationarity and bounded trends in the context of TVP models, and discuss the computational costs associated with those restrictions in a Bayesian setting.