International Journal of Forecasting 32 (2016) 875–887

Contents lists available at ScienceDirect

International Journal of Forecasting

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

In-sample confidence bands and out-of-sample forecast bands for time-varying parameters in observation-driven models

Francisco Blasques [a], Siem Jan Koopman [a,b,*], Katarzyna Łasak [a], André Lucas [a]

[a] Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands

[b] CREATES, Aarhus University, Denmark

ARTICLE INFO

Keywords:
Autoregressive conditional duration
Generalized autoregressive conditional heteroskedasticity
Score driven models
Time-varying mean
Delta-method

ABSTRACT

We study the performances of alternative methods for calculating in-sample confidence and out-of-sample forecast bands for time-varying parameters. The in-sample bands reflect parameter uncertainty, while the out-of-sample bands reflect not only parameter uncertainty, but also innovation uncertainty. The bands are applicable to a wide range of estimation procedures and a large class of observation driven models with differentiable transition functions. A Monte Carlo study is conducted to investigate time-varying parameter models such as generalized autoregressive conditional heteroskedasticity and autoregressive conditional duration models. Our results show convincing differences between the actual coverages provided by the different methods. We illustrate our findings in a volatility analysis for monthly Standard & Poor's 500 index returns.

1. Introduction

Over recent decades, time-varying parameter models have become increasingly popular in empirical economics and finance. The rapid development of new methods for filtering time-varying parameters in dynamic models with nonlinear and non-Gaussian features has made these models more accessible, flexible and attractive. Initially, starting in the 1960s, time-varying parameters for the mean equation in linear Gaussian models were typically handled by the Kalman filter and related methods. For linear Gaussian state space models, the Kalman filter can be used to calculate the conditional means and variances of unobserved time-varying parameters (or linear functions thereof) in a computationally efficient way; for a detailed treatment, see for example Durbin and Koopman (2012). In this modeling framework, the construction of in-sample confidence bands and out-of-sample forecast bands is straightforward, and is performed on a routine basis, as expressions for the conditional variances of the time-varying parameters are available explicitly. In the case of nonlinear and/or non-Gaussian extensions of state space models, the computation of confidence bands can be somewhat more involved. One example is the stochastic volatility model, for which the analysis is typically based on simulation-based methods; see the discussion by Shephard (2005).

Since the 1980s, new classes of model for time-varying parameters have been developed. Specifically, models for the time-varying conditional variance have received considerable attention in the empirical economics and finance literature. For example, the generalized autoregressive conditional heteroskedasticity (GARCH) model of Bollerslev (1986) and Engle (1982) has led to a range of model

* Corresponding author at: Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands.

E-mail addresses: f.blasques@vu.nl (F. Blasques), s.j.koopman@vu.nl (S.J. Koopman), k.a.lasak@vu.nl (K. Łasak), a.lucas@vu.nl (A. Lucas).

http://dx.doi.org/10.1016/j.ijforecast.2015.11.018
0169-2070/

formulations for time-varying parameters. In the standard ARCH and GARCH models, the conditional variance is obtained by filtering past observations through a volatility updating equation. The relative simplicity of GARCH models has spurred their widespread adoption by both academics and professionals.

In most empirical studies, the estimated volatility from the GARCH model is presented without in-sample bands reflecting the parameter uncertainty in the volatility updating equation. Similarly, volatility forecasts may feature bands that reflect innovation uncertainty, but they typically ignore the parameter uncertainty. Exact analytical results are not available because the filters are highly nonlinear functions of past observations, and as a result, statistical software rarely provides either in-sample confidence bands or out-of-sample forecast bands to represent the estimates of such time-varying parameters. This argument also applies to other models that are related to GARCH, including the autoregressive conditional duration (ACD) model of Engle and Russell (1998), the multiplicative error model of Engle (2002), the observation-driven Poisson count model of Davis, Dunsmuir, and Streett (2003), and the score driven models of Creal, Koopman, and Lucas (2013). All of these models belong to the class of observation-driven models, as opposed to parameter-driven models; see Cox (1981) for a detailed description of these two classes of time series models.

We analyze various different methods of constructing in-sample and out-of-sample bands. For our in-sample bands, we compare two analytical methods and one simulation-based method. All of these bands reflect the parameter uncertainty only. The approximate analytical bands require only simple computations, and are not subject to random fluctuations due to simulation error. These analytical bands can be used when the updating equation is differentiable and the (asymptotic) distribution of the estimator for the static parameters is known. For the computation of forecast bands, we compare three simulation-based procedures. The first method takes only innovation uncertainty into account. The second and third methods incorporate both parameter and innovation uncertainty. Although these methods require simulations, the forecast bands are relatively quick to compute. In particular, we argue that the necessary computations are more efficient than the bootstrapped forecast bands proposed by Pascual, Romo, and Ruiz (2006) for GARCH models, for example.

All of the methods that we consider can be implemented readily in software packages. We investigate the coverage probabilities of each of these different approaches in detail over a range of different time-varying parameter models, and find that simulation-based methods are the most reliable, but that the approximate analytical methods also perform well in many settings.

To provide evidence of how effective the different methods are, we present the results of a Monte Carlo study in which we compute in-sample confidence bands and out-of-sample forecast bands for time series generated using GARCH, score-driven, ACD and time-varying mean (local level) models. The results reveal that the actual coverage of our (preferred) analytical bands is close to the nominal coverage level obtained by simulation. The simulation-based in-sample confidence bands and out-of-sample forecast bands all obtain accurate coverage levels. An empirical illustration for the GARCH model applied to a time series of monthly log-returns from the Standard & Poor's 500 index reveals the practical importance of these bands. We also show that the choice of the method for computing in-sample bands is empirically relevant, and that our analytical bands provide a good approximation to the more computationally-intensive simulation-based bands.

The remainder of this paper is organized as follows. Section 2 introduces the class of observation-driven models. Section 3 introduces different methods of computing in-sample bands for the time-varying parameter. Section 4 presents different simulation-based methods for the computation of the out-of-sample forecast bands. Section 5 analyzes the relative performances of the bands in a Monte Carlo study. Section 6 presents our empirical findings for the Standard & Poor's 500 monthly returns. Section 7 concludes.

2. Observation-driven models

In observation-driven models, the time-varying parameter is filtered using an updating equation that depends on past observations. In these models, the focus is on the specification of the mechanism through which past realizations of the variable of interest affect the current value of the time-varying parameter.

Consider a model for an observed time series $y_1, \dots, y_T$ given by

$$y_t \sim p_y(y_t | f_t; \boldsymbol{\theta}), \quad t = 1, \dots, T, \tag{1}$$

where density $p_y(\cdot)$ is implied by an "observation equation" for y_t and depends on the time-varying parameter f_t and the static parameter $\boldsymbol{\theta}$. For example, $y_t = f_t + \varepsilon_t$ for a time-varying mean, or $y_t = \mu + f_t^{1/2}\varepsilon_t$ for a fixed mean and time-varying variance, with, for instance, $\varepsilon_t \sim \text{NID}(0, 1)$, where NID is normally independently distributed with zero mean and unity variance. The time-varying parameter is defined formally as a function $f_t := f_t(y^{1:t-1}, f_1; \boldsymbol{\theta})$ that depends on the past observations $y^{1:t-1} := \{y_1, y_2, \dots, y_{t-1}\}$, an initial value f_1, and a static parameter vector $\boldsymbol{\theta}$. The updating function for the time-varying parameter can be expressed in different ways. For example, if we consider a linear updating equation consisting of lagged values of y_t and f_t, we obtain

$$f_{t+1} = \omega + \beta f_t + \alpha s(y_t, f_t; \boldsymbol{\theta}), \tag{2}$$

with initialization f_1, and where $s(y_t, f_t; \boldsymbol{\theta})$ is a (possibly nonlinear) function of y_t, f_t, and $\boldsymbol{\theta}$. The function $s(\cdot)$ can be chosen in a flexible way, and is often just a transformation of y_t, as we will show in the examples below. The coefficients ω, α and β are part of the parameter vector $\boldsymbol{\theta}$. The recursive nature of the formulation implies that f_{t+1} is a (nonlinear) function of $y_t, \dots, y_1, f_1$ and $\boldsymbol{\theta}$. Hence, the updating equation (Eq. (2)) is consistent with the definition of f_t, that is, $f_t := f_t(y^{1:t-1}, f_1; \boldsymbol{\theta})$. We can also