International Journal of Forecasting 31 (2015) 620–634

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Forecasting realized volatility with changing average levels

Giampiero M. Gallo [a,*], Edoardo Otranto [b]

[a] Dipartimento di Statistica, Informatica, Applicazioni "G. Parenti", Università di Firenze, Viale G. B. Morgagni, 59 - 50134 Firenze, Italy

[b] Dipartimento di Scienze Cognitive, della Formazione e degli Studi Culturali and CRENoS, Università di Messina, Via Concezione, 6 - 98121 Messina, Italy

ARTICLE INFO

Keywords:
Evaluating forecasts
MEM
Regime switching
Robustness
Volatility forecasting

ABSTRACT

We explore the abilities of regime switching with Markovian dynamics (MS) and of a smooth transition (ST) nonlinearity within the class of Multiplicative Error Models (MEMs) to capture the slow-moving long-run average in the realized volatility. We compare these models to some alternatives, including considering (quasi) long memory features (HAR class), the benefits of log transformations, and the presence of jumps. The analysis is applied to the realized kernel volatility series of the S&P500 index, adopting residual diagnostics as a guidance for model selection. The forecast performance is evaluated and tested via squared and absolute losses both in- and out-of-sample, as well as based on a robustness check on different sample choices. The results show very satisfactory performances of both MS and ST models, with the former also allowing for the dating and interpretation of regimes in terms of market events.

1. Introduction

A consolidated body of literature in financial econometrics is devoted to the measurement of asset volatility, exploiting the information contained in asset price data collected at a very high frequency. The volatility estimators, known as realized volatility (RV) measures, have allowed for deeper insights into the dynamics of volatility, which are traditionally analyzed in a modeling and forecasting framework within the GARCH paradigm as conditional variances of returns (see Bollerslev, 1986, Engle, 1982, and further extensions; for a review, see Teräsvirta, 2009; for nonlinear models, see Teräsvirta, 2011). Starting from the plain realized volatility, studied in detail by Andersen, Bollerslev, Diebold, and Labys (2000, 2003), various other measures have been introduced to take into consideration the presence of jumps and other market microstructure issues (for a review, see Andersen, Bollerslev, & Diebold, 2010). The most recent addition to the family of volatility estimators is the realized kernel volatility (developed by Barndorff-Nielsen, Hansen, Lunde, & Shephard, 2008), which is designed to possess robustness to market microstructure noise.

While volatility measurement from an end-of-day perspective has reached a mature stage, the modeling and forecasting of conditional volatility is open to refinements, building on existing dynamic models (cf. among others, Brownlees & Gallo, 2010; Cipollini, Engle, & Gallo, 2013; Hansen, Huang, & Shek, 2012; Shephard & Sheppard, 2010). In Multiplicative Error Models (MEMs, developed by Engle, 2002, and Engle & Gallo, 2006), the volatility series is specified as the product of a time-varying scale factor that evolves autoregressively and a random disturbance with a suitable distribution. As it is applied to non-negative values, a MEM captures dynamics without resorting to logs, thus producing forecasts of volatility (not log-volatility); adopting multiplicative (rather than additive) disturbances accommodates heteroskedasticity naturally ("vol-of-vol").

* Corresponding author.
E-mail addresses: gallog@disia.unifi.it (G.M. Gallo), eotranto@unime.it (E. Otranto).

http://dx.doi.org/10.1016/j.ijforecast.2014.09.005

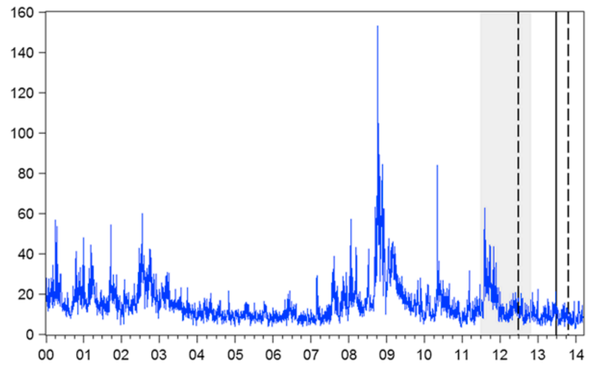

Fig. 1. Realized kernel volatility of the S&P500 index (Jan. 3, 2000 to Mar. 10, 2014); the gray shadowed area represents the out-of-sample period of the main analysis (between Jul. 5, 2011, and Oct. 26, 2012). For the robustness check, the vertical dashed lines identify the first out-of-sample period (Jul. 5, 2012 to Oct. 25, 2013) and the solid line is the beginning of the second out-of-sample period (Jul. 5, 2013).

The volatility profile over long periods typically exhibits fairly persistent local average volatilities (cf., as a leading example, the realized kernel volatility of the S&P500 in Fig. 1). Competing models must capture this empirical regularity with possibly alternative or complementary features (discrete changes in volatility, slow-moving long-run dynamics, other nonlinearities, or explicit long memory), and will be evaluated for both their in-sample fitting capabilities and their out-of-sample forecasting performances. The specific suggestions that we make involve extending MEMs to have regime-switching (Hamilton, 1989, 1990) or smooth-transition (Teräsvirta, 1994) representations, keeping variants of the heterogeneous autoregressive model (HAR, also known as a quasi-long memory model) of Corsi (2009) for RV and log-RV, for comparison purposes.

The consideration of time-varying local averages is not new to the volatility literature (see below for a review); our extended class of Markov Switching (MS) MEMs highlights regime-specific dynamics. Relative to other MEMs, our approach is distinct from both the mixture MEM (either that of De Luca & Gallo, 2004, and Lanne, 2006, which suffers from the shortcoming of time-independent regime indicator variables, cf. Bauwens, Hafner, & Laurent, 2012; or that of De Luca & Gallo, 2009, who make the mixing weights dependent on a lagged observable variable) and the P-Spline MEM used by Brownlees and Gallo (2010) for realized volatility with a time varying average level (which is a statistical fit to a supposed smooth underlying trend). A standard HAR may face serious challenges in fitting the observed pattern given its linear nature; the HAR-J of Andersen, Bollerslev, and Diebold (2007) reacts to a separate measure of jumps as a possible accommodation of abrupt changes, but its out-of-sample performance depends crucially on jointly forecasting jumps as well. The issue of whether a log-transformation mitigates the importance of peaks in the series, flattening out the extent to which local averages may be present, seems to be mostly an empirical question. With many models at hand, we need a suitable model selection strategy; we adopt a nonparametric Bayesian procedure (Otranto & Gallo, 2002) for selecting the number of regimes, and take the view that the presence of autocorrelation in residuals and/or squared residuals may be seen as evidence of the presence of some form of model misspecification (in the dynamics and/or nonlinearity). In-sample fitting capabilities will not necessarily correspond to good forecasting performances, especially at different forecasting horizons; and finally, but possibly as a by-product, it is useful to evaluate the compatibility of the results with the interpretation of major events in financial markets.

The paper is organized as follows: in Section 2, we review the main literature on the nonlinear modeling of the realized volatility. In Section 3, we discuss the issues behind discrete and slow-moving local average volatilities, with a description of the Markov Switching and Smooth Transition extensions within the MEM class; we analyze the estimators' properties and provide some Monte Carlo evidence as to the small-sample properties of the estimators. In the same section, we also detail some of the alternative parameterizations that may capture different features of the data. An empirical application to the realized kernel volatility S&P500 series is the object of Section 4. The estimation of 13 models is a preliminary step to determining certain features of the data and allowing for a reduction in the number of models based on their in-sample properties and forecasting performances (one- and ten-step-ahead forecasts with the Diebold & Mariano, 1995 (DM), test and the Model Confidence Set (MCS) sequential test of Hansen, Lunde, & Nason, 2011, using absolute and squared losses). The inference on the regimes and the comparison with the smooth transition function provide some interesting insights on the classification of the periods and the calculation of the corresponding average levels. The robustness of the model performances is also checked on different time spans. Concluding remarks follow. Finally, some extra details on model estimation and comparison are provided in a web appendix.

2. High persistence in volatility

There is a major debate in the literature about the nature of the high level of persistence in realized volatility, and whether it may be the result of some nonlinearity in the process. The HAR model (Corsi, 2009), although not a formal long-memory model, can reproduce the observed hyperbolic-type decay of the autocorrelation function of (log-)volatility by specifying a sum of volatility components over different horizons. Andersen et al. (2007) insert a volatility jump component for capturing the abrupt changes that characterize the realized volatility, resulting in significant improvements in the forecasting performance. Baillie and Kapetanios' (2007) reasoning about the existence of both non-linear and long memory components in many economic and financial time series is developed by McAleer and Medeiros (2008), who introduce a multiple regime smooth transition extension of the HAR; their model is also able to capture the presence of sign and size asymmetries. Bordignon and Raggi (2012) propose an elegant solution for combining, in a single model, both non-linearity effects, through a Markov switching process, and high persistence, through fractionally integrated dynamics, that are capable of improving the accuracy of both in- and out-of-sample forecasts. Alternatively, concentrating on long memory explanations, Andersen et al. (2003)