Computational Statistics and Data Analysis ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

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

Testing for jumps in conditionally Gaussian ARMA–GARCH models, a robust approach

Sébastien Laurent[a,b,c,d,*], Christelle Lecourt[a,b,e], Franz C. Palm[d]

[a] Aix-Marseille University (Aix-Marseille School of Economics), France
[b] CNRS & EHESS, France
[c] IAE-Aix, France
[d] Maastricht University, School of Business and Economics, The Netherlands
[e] CeReFim, Université de Namur, Belgium

ARTICLE INFO

Article history:
Received 3 May 2013
Received in revised form 29 April 2014
Accepted 22 May 2014
Available online xxxx

Keywords:
Jumps
GARCH
Test
Forecasting

ABSTRACT

Financial asset prices occasionally exhibit large changes. To deal with their occurrence, observed return series are assumed to consist of a conditionally Gaussian ARMA–GARCH type model contaminated by an additive jump component. In this framework, a new test for additive jumps is proposed. The test is based on standardized returns, where the first two conditional moments of the non-contaminated observations are estimated in a robust way. Simulation results indicate that the test has very good finite sample properties, i.e. correct size and high proportion of correct jump detection. The test is applied to daily returns and detects less than 1% of jumps for three exchange rates and between 1% and 3% of jumps for about 50 large capitalization stock returns from the NYSE. Once jumps have been filtered out, all series are found to be conditionally Gaussian. It is also found that simple GARCH-type models estimated using filtered returns deliver more accurate out-of sample forecasts of the conditional variance than GARCH and Generalized Autoregressive Score (GAS) models estimated from raw data.

1. Introduction

The distributional properties of speculative prices have been extensively studied in the finance literature. High frequency returns of most financial assets exhibit volatility clustering and large jumps often caused by the arrival of important news surprises (e.g. news announcements).

In this paper, we propose a new semi-parametric statistical procedure to detect additive jumps in financial series. It is similar to the non-parametric tests for jumps proposed by Lee and Mykland (2008) and Andersen et al. (2007b). Our test is expected to be useful when these two tests are not applicable, i.e. when intraday data are not available or, if they are, the asset is not liquid enough to be frequently traded. A test for jumps based on daily returns is therefore needed because, as shown by Andersen et al. (2007a), the largest shocks have a relatively smaller effect on future volatility than smaller shocks. As a consequence, GARCH models neglecting jumps usually overestimate the volatility during several days, if not weeks, after the occurrence of these jumps. See for instance the example of the stock price of Apple in Boudt et al. (2013).

* Correspondence to: GREQAM, Château La Farge, Route des Milles, 13290 Les Milles, France.
E-mail addresses: sebastien.laurent@iae-aix.com, sebastien.laurent@univ-amu.fr (S. Laurent), christelle.lecourt@univ-amu.fr (C. Lecourt), f.palm@maastrichtuniversity.nl (F.C. Palm).

http://dx.doi.org/10.1016/j.csda.2014.05.015
0167-9473/

Similar to Lee and Mykland (2008) and Andersen et al. (2007b), who standardize their non-parametric test statistics by a robust to jumps estimate of instantaneous volatility (based on realized bipower variation), we standardize our test statistic using the conditional volatility based on a robustified GARCH volatility estimate and a robust conditional mean estimate. Our test therefore incorporates the idea that when spot or instantaneous volatility is high (also in the absence of jumps), returns may also be high, even as high as those due to jumps. It is built using the same framework as that of Franses and Ghijsels (1999), Lee and Mykland (2008) and Andersen et al. (2007b), that returns are conditionally Gaussian on days without jumps.

We apply our test to daily returns and detected as jumps less than 1% of observations for the three exchange rates and between 1% and 3% of observations for all the stocks considered in the application (about 50 large capitalization stocks from the NYSE). Interestingly, once jumps have been filtered out, all series are found to be conditionally Gaussian. Importantly, a rejection of the conditional normality of the jump-corrected data would be taken as an indication that our theoretical framework is inappropriate. To investigate this issue, we simulated data using a model with ARCH effects and Student-t innovations. This model is able to generate volatility clustering and large once-off events that might be indistinguishable from additive jumps. However, this model does not contain additive jumps and therefore our test should not be applied in this case. Results suggest that while a small proportion of jumps is detected by our test, a Jarque–Bera test has very good power (greater than 95% for a degree of freedom of the Student distribution < 7) to reject the assumption of conditional normality of the jump-corrected data.

The rest of this article is organized as follows. In Section 2, the theoretical framework and the model setting are described and the proposed semi-parametric test and its asymptotic distribution are presented. In Section 3, simulation results comparing our test with alternative tests in the literature are presented. Section 4 contains the findings of an empirical study of jump detection in daily data. Finally, Section 5 concludes the paper.

2. Model and test

Andersen et al. (2007a), Harvey and Chakravarty (2008) and Muler and Yohai (2008) among others found that the jumps affect future asset return volatility less than what standard return volatility models predict. Andersen et al. (2007a) show that conditioning also on the past jumps in an autoregressive (AR) model for the realized volatility tempers the persistence in the volatility forecasts, indicating that the jumps in asset prices tend to lead the short-lived increases in volatility. In a univariate GARCH setting, Sakata and White (1998), Franses and Ghijsels (1999), Carnero et al. (2007, 2008), Charles and Darné (2005) and Muler and Yohai (2008) show that, in the presence of additive jumps, the Gaussian Quasi-Maximum Likelihood (QML) estimator of GARCH models tends to overestimate the volatility for the days following a jump and to produce upward biased estimates of the long-run volatility.

2.1. Data generating process

One of the most popular model for the financial return data sampled, at say the daily frequency, is certainly the GARCH(1, 1) of Bollerslev (1986). To account for leverage effect, Glosten et al. (1993) proposed a generalization of this model, called GJR. A random variable r_t follows a normal-ARMA(p, q)–GJR$(1, 1)$ model if it can be described by the system (2.1)–(2.4):

$$r_t = \mu_t + \varepsilon_t \tag{2.1}$$

$$\mu_t = c + \sum_{i=1}^{\infty} \zeta_i \varepsilon_{t-i} \tag{2.2}$$

$$\varepsilon_t = \sigma_t z_t \quad \text{and} \quad z_t \overset{\text{i.i.d.}}{\sim} N(0, 1) \tag{2.3}$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \gamma_1 D_{t-1} \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \tag{2.4}$$

where $D_{t-1} = 1$ if $\varepsilon_{t-1} < 0$ and 0 otherwise, ζ_i's are the coefficients of $\zeta(L) = \phi^{-1}(L)\theta(L) = 1 + \sum_{i=1}^{\infty} \zeta_i L^i$, L is the lag operator, $\phi(L) = 1 - \sum_{i=1}^{p} \phi_i L^i$ and $\theta(L) = 1 - \sum_{i=1}^{q} \theta_i L^i$ are the AR and MA polynomials of orders p and q respectively (with roots outside the unit circle). Therefore, μ_t is the conditional mean of r_t while σ_t^2 is its conditional variance of r_t. Note that Model (2.1)–(2.4) reduces to a normal-ARMA(p, q)–GARCH$(1, 1)$ when $\gamma_1 = 0$. Eq. (2.4) can be further extended to account for long-memory like in the FIGARCH model of Baillie et al. (1996) or the FIAPARCH of Tse (1998). This extension is beyond the scope of the paper.

If we add an independent jump component $a_t I_t$ to r_t, we obtain

$$r_t^* = r_t + a_t I_t, \tag{2.5}$$

where r_t^* denotes observed financial returns, I_t is a binary variable taking value 1 in case of a jump on day t and 0 otherwise and a_t is the size of the jumps (either positive or negative). We assume that a_t and I_t are independent of each other, that I_t is independently distributed across time and independent of the past of r_t. Jump size a_t is either a sequence of numbers or a function of the past squared values of r_t, e.g. a_t is proportional to the conditional standard-deviation of r_t. In the latter case