ARTICLE IN PRESS

Journal of Econometrics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Econometrics

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

Bad environments, good environments: A non-Gaussian asymmetric volatility model[☆]

Geert Bekaert[a,b], Eric Engstrom[c,*], Andrey Ermolov[a]

[a] Columbia University, United States

[b] National Bureau of Economic Research, United States

[c] Board of Governors of the Federal Reserve System[1], United States

ARTICLE INFO

Article history:
Received 25 July 2013
Received in revised form
15 April 2014
Accepted 9 June 2014
Available online xxxx

JEL classification:
G11
G17

Keywords:
Non-Gaussianities
GARCH
Asymmetric volatility
Conditional skewness
Risk management

ABSTRACT

We propose an extension of standard asymmetric volatility models in the generalized autoregressive conditional heteroskedasticity (GARCH) class that admits conditional non-Gaussianities in a tractable fashion. Our "bad environment–good environment" (BEGE) model utilizes two gamma-distributed shocks and generates a conditional shock distribution with time-varying heteroskedasticity, skewness, and kurtosis. The BEGE model features nontrivial news impact curves and closed-form solutions for higher-order moments. In an empirical application to stock returns, the BEGE model outperforms asymmetric GARCH and regime-switching models along several dimensions.

Published by Elsevier B.V.

1. Introduction

Since the seminal work of Engle (1982) and Bollerslev (1986) on volatility clustering, thousands of articles have applied models in the generalized autoregressive conditional heteroskedasticity (GARCH) class to capture volatility clustering in economic and financial time series data. In the basic GARCH (1, 1) model, the conditional variance is a deterministic function of the past conditional variance and contemporaneous squared shocks to the process describing the data. Nelson (1991) and Glosten et al. (1993, GJR henceforth), motivated by empirical work on stock return data, provide important extensions, accommodating asymmetric responses of conditional volatility to negative versus positive shocks. Engle and Ng (1993) compare the response of conditional variance to shocks ("news impact curves") implied by various econometric models and find evidence that the GJR model fits stock return data the best.

The original models in the GARCH class assumed Gaussian innovations, but nonetheless imply non-Gaussian unconditional distributions. However, time-varying volatility models with Gaussian innovations generally do not generate sufficient unconditional non-Gaussianity to match certain financial asset return data (see, e.g. Poon and Granger, 2003). Additional evidence of conditional non-Gaussianity has come from two corners. First, empirical work by Evans and Wachtel (1993), Hamilton and Susmel (1994), Kim and White (2004), and many others has documented conditional non-Gaussianities in economic data. Second, in finance, a voluminous literature on the joint properties of option prices and stock returns (see, e.g. Broadie et al., 2009) has also suggested the need for models with time-varying non-Gaussianities. In principle, one can estimate GARCH models consistently using quasi maximum likelihood (see Lumsdaine, 1996; Hansen and Lee, 1994), not worrying about modeling the non-Gaussianity in the shocks. However,

☆ The authors thank Nancy R. Xu, seminar participants at CK GSB (Beijing), the University of Sydney, UNSW (Sydney), SAIF (Shanghai), SMU (Singapore), Temple University, HEC/EPFL in Lausanne, and the 2014 NBER Summer Institute for useful comments. The authors are especially grateful for the suggestions of two anonymous referees, and the editor, Yacine Ait-Sahalia, which greatly improved the paper. Geert Bekaert acknowledges financial support from Netspar. All errors are the sole responsibility of the authors.

* Correspondence to: Federal Reserve, Washington, DC 20551, United States. Tel.: +1 202 452 3044; fax: +1 202 728 5887.

E-mail address: eric.c.engstrom@frb.gov (E. Engstrom).

1 The views expressed in this document do not necessarily reflect those of the Board of Governors of the Federal Reserve System, or its staff.

http://dx.doi.org/10.1016/j.jeconom.2014.06.021
0304-4076/Published by Elsevier B.V.

fitting the actual non-Gaussianities in the data can lead to more efficient estimates and may be important if the model is to be used in applications (for example, option pricing or risk management) that require an estimate of the conditional distribution. Several authors have introduced non-Gaussian shocks in GARCH frameworks (see, e.g., Bollerslev (1987) and Hsieh (1989), who used the t-distribution, and Mittnik et al. (2002), who used shocks with a distribution in the stable Paretian class). However, as we will show, extant models in this vein generally cannot fit time-varying non-Gaussianities that are evident in the data.

We present an extension of models in the GARCH class that accommodates conditional non-Gaussianity in a tractable fashion, offering simple closed-form expressions for conditional moments. Our "bad environment–good environment" (BEGE) model utilizes two gamma-distributed shocks that together imply a conditional shock distribution with time-varying heteroskedasticity, skewness, and kurtosis. This is accomplished by allowing the shape parameters of the two distributions to vary through time. Hence, our model features rich variation in higher-order moments. We apply the model to stock returns, showing that the model outperforms extant alternatives using a variety of specification tests. In the stock market context, one shape parameter determines the conditional distribution of the "good environment", with positive skewness and "good volatility"; the other shape parameter drives the "bad environment", with negative skewness and "bad volatility". Of course, conditional non-Gaussian models exist outside the GARCH class that may also fit the data quite well. Regime-switching models, in particular, have shown promise in many applications. We therefore also estimate several types of regime-switching models on our stock returns data sample and show that the BEGE model significantly outperforms various models in this class.

The remainder of the article is organized as follows. In Section 2, we present the BEGE model, describe how it nests the standard GJR–GARCH model as a special case, and present various models in the regime-switching class. In Section 3, we describe the estimation methodology and the specification tests that we conduct. In Section 4, we confront several models from the above classes, including the BEGE model, with monthly US stock return data from 1929 through 2010.

2. The BEGE–GARCH model

Before introducing the BEGE model, we begin with a review of the seminal GJR asymmetric GARCH model.

2.1. Traditional GJR–GARCH

Consider a time series r_{t+1} with conditional mean μ_t. The GJR model assumes that the series follows

$$r_{t+1} = \mu_t + u_{t+1},$$
$$u_{t+1} \sim N(0, h_t),$$
$$\text{and} \quad h_t = h_0 + \rho_h h_{t-1} + \phi^+ u_t^2 I_{u_t \geq 0} + \phi^- u_t^2 \left(1 - I_{u_t \geq 0}\right). \tag{1}$$

That is, the innovation to returns, u_{t+1}, has time-varying conditional variance, $var_t(r_{t+1}) = h_t$, which is assumed to be a linear function of its own lagged value and squared innovations to returns. One key feature of this model that enables it to better fit many economic time series is the differential response of the conditional variance of shocks following positive versus negative innovations. In stock return and economic activity data, it is typically found that $\phi^- > \phi^+$, indicating that negative shocks result in more of an increase in variance than do positive shocks.

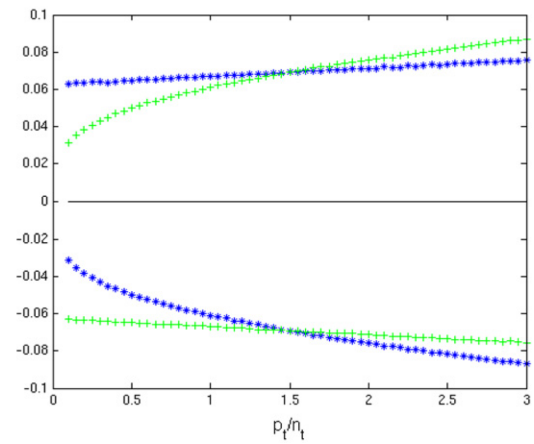

Fig. 1. BEGE distribution tail properties.
This plot shows the 99th percentiles and 1st percentiles for two sequences of BEGE distributions, which take the form.

$$u_{t+1} = \omega_{p,t+1} - \omega_{n,t+1}; \omega_{p,t+1} \sim \widetilde{\Gamma}(p_t, \sigma_p); \omega_{n,t+1} \sim \widetilde{\Gamma}(n_t, \sigma_n),$$

where $\widetilde{\Gamma}$ denotes the centered gamma distribution. Throughout, we maintain that $\sigma_n = \sigma_p = 0.015$. The lines of blue asterisks show the quantiles for distributions in which p_t is fixed at 1.5, but n_t varies from 0.1 through 3.0. Conversely, the lines of green plus symbols show the quantiles for distributions in which p_t varies from 0.1 through 3.0 while n_t is held fixed at 1.5.

2.2. BEGE–GJR–GARCH

The BEGE model that we propose relaxes the assumption of Gaussianity by assuming that the u_{t+1} innovation consists of two components. We assume that $\omega_{p,t+1}$, a good environment shock, and $\omega_{n,t+1}$, a bad environment shock, are drawn from "demeaned" (or "centered") gamma distributions that have a mean equal to zero.[2] The overall innovation is a linear combination of the two component shocks, which are assumed to be conditionally independent. The gamma distributions are assumed to have constant scale parameters, but we let their shape parameters vary through time. More precisely, the BEGE framework assumes:

$$u_{t+1} = \sigma_p \omega_{p,t+1} - \sigma_n \omega_{n,t+1}, \quad \text{where}$$
$$\omega_{p,t+1} \sim \widetilde{\Gamma}(p_t, 1), \quad \text{and}$$
$$\omega_{n,t+1} \sim \widetilde{\Gamma}(n_t, 1), \quad \text{and} \tag{2}$$

where $\widetilde{\Gamma}(k, \theta)$ denotes a centered gamma distribution with shape and scale parameters, k and θ, respectively. Thus, p_t (n_t) is the shape parameter for the good (bad) environment shock. Fig. 1 provides a visual representation of the flexibility of the BEGE distribution. Plotted are the 1st and 99th percentiles of two sequences of hypothetical distributions. The blue stars illustrate a series of BEGE distributions for which p_t is fixed at 1.5, but n_t varies from 0.1 to 3.0, which are the values across the horizontal axis. The lower line of blue asterisks shows the 1st percentiles of these distributions, while the upper line of blue stars shows the 99th percentiles. Clearly, increases in n_t have an outsized effect on the lower tail, particularly at low values of n_t. The upper tail is relatively insensitive to changes in n_t. The green plus symbols show results from the complementary exercise: holding n_t fixed at 1.5 and varying p_t from 0.1 through 3.0. Clearly p_t impacts the upper tail of the distribution much more than it impacts the lower tail. These results

[2] The centered gamma distribution with shape parameter k and scale parameter θ, which we denote $\widetilde{\Gamma}(k, \theta)$, has probability density function, $\phi(x) = \frac{1}{\Gamma(k)\theta^k}(x + k\theta)^{k-1} \exp\left(-\frac{1}{\theta}(x + k\theta)\right)$ for $x > -k\theta$, and with $\Gamma(\cdot)$ representing the gamma function.

Please cite this article in press as: Bekaert, G., et al., Bad environments, good environments: A non-Gaussian asymmetric volatility model. Journal of Econometrics (2014), http://dx.doi.org/10.1016/j.jeconom.2014.06.021