Energy Economics 43 (2014) 82–88

Contents lists available at ScienceDirect

Energy Economics

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

Daily seasonality in crude oil returns and volatilities

Benjamin R. Auer *

University of Leipzig, Department of Finance, Grimmaische Str. 12, 04109 Leipzig, Germany

ARTICLE INFO

Article history:
Received 28 October 2013
Received in revised form 4 February 2014
Accepted 16 February 2014
Available online 22 February 2014

JEL classification:
C10
G10

Keywords:
Crude oil
GARCH
Seasonality

ABSTRACT

In this article, we test for the existence of daily seasonality in returns and volatilities of crude oil. Using a dummy-augmented GARCH specification for the period from May 1987 to October 2013, our key findings are as follows: (i) Volatilities on Mondays are significantly higher than on all other weekdays, providing the important insight that seasonal effects should be considered when forecasting crude oil volatility. (ii) Returns on the other hand tend to be lower on Mondays than on other weekdays, suggesting profitable investment strategies based on this seasonal pattern. In fact, the analysis of a simple long-short trading rule based on the Monday effect provides some evidence that it can outperform a passive buy-and-hold approach. However, it cannot do so to an extent that is statistically significant. (iii) Our seasonality results are fairly robust to the choice of other frequently used GARCH model variants, like GARCH-M, TGARCH and CGARCH.

1. Introduction

In recent years, a considerable number of studies have contributed to a better understanding of crude oil price and return dynamics. Among the various research topics, modelling volatility and testing the efficiency of the crude oil market are of special interest for energy researchers, market participants and policy makers. This is because volatility is an important input for the valuation of oil-based derivatives, hedging activities, and decisions to invest in oil inventories and facilities for the production, transportation or consumption of oil (see Narayan and Narayan, 2007). Furthermore, higher volatility in oil prices creates uncertainty, leading to economic instability and negative impacts on economic growth for both oil-exporting and oil-importing countries (see Federer, 1996; Jimenez-Rodriguez and Sanchez, 2005; Regnier, 2007). The efficiency issue is of particular relevance because inefficient anomalous market behaviour may allow return predictability and the implementation of profitable investment strategies (see Lucey and Pardo, 2005; Lucey and Tully, 2006).

Literature on volatility has so far concentrated on modelling and forecasting volatility using generalised autoregressive conditional heteroscedasticity (GARCH) models and their variants (see Adrangi and Chatrath, 2002; Cabedo and Moya, 2003; Fong and See, 2002; Giot and Laurent, 2003; Kang et al., 2009; Mohammadi and Su, 2010; Morana, 2001; Narayan and Narayan, 2007; Sadeghi and Shavvalpour, 2006; Sadorsky, 2006),[1] the analysis of the relationship between oil price volatility and the macroeconomy (see Chen and Chen, 2007; Federer, 1996; Huang et al., 2005; Lee et al., 1995; Yang et al., 2002), its connection to stock price movements (see Huang et al., 1996; Sadorsky, 1999, 2003) and a comparison of crude oil volatility to the volatilities of other commodities (see Pindyck, 1999; Plourde and Watkins, 1998; Regnier, 2007). The comparably smaller efficiency literature focuses on testing for efficiency by means of time-varying long-range dependence, Hurst exponent dynamics, unit root tests and variance ratio tests (see Alvarez-Ramirez et al., 2008; Charles and Darné, 2009; Maslyuk and Smyth, 2008; Tabak and Cajueiro, 2007).

The main goal of this article is to contribute to both strands of the literature by investigating the existence of daily seasonality in crude oil returns and volatilities. The phenomenon of daily seasonality has been extensively analysed in equity markets, with one of the most influential results being that Mondays tend to show the lowest, often negative, returns despite having the highest, or higher than average, risk as proxied by standard deviation (see Maberly, 1995).[2] Identifying similar effects in crude oil returns would have two major implications: First,

[1] Mohammadi and Su (2010) summarise studies using other methodologies like vector autoregressive models, cointegration and error correction modelling, neural networks and Jump-diffusion processes.

[2] Recent studies have also started to analyse other markets. For example, Lucey and Tully (2006) focus on cash and future data for gold and silver. Blose and Gondhalekar (2013) list articles observing the Monday effect for T-Bills, exchange rates, real estate investment trusts and futures markets.

* Tel.: +49 341 97 33 672; fax: +49 341 97 33 679.
E-mail address: auer@wifa.uni-leipzig.de.

http://dx.doi.org/10.1016/j.eneco.2014.02.005
0140-9883/

higher returns on specific weekdays imply the potential of generating abnormal profits based on the identified pattern. At least timing or market entry/exit strategies if not actual trading strategies could be successful (see Lucey and Pardo, 2005). Second, differences in volatilities during a week would provide new explanatory variables relevant for improving currently used volatility forecasting models (see Kang et al., 2009; Mohammadi and Su, 2010).

In order to provide such insights, we estimate a dummy-augmented GARCH model for crude oil returns. Unlike conventional dummy variable regressions traditionally used in seasonality analysis, our specification allows in avoiding misleading inference caused by departures from normality, serial correlation and heteroscedasticity (see Chien et al., 2002; Connolly, 1989). It is also a natural choice for testing any seasonal effect on volatility. Furthermore, to check the robustness of our results, we extend our analysis by considering alternative popular GARCH models, namely, the GARCH-M, TGARCH and CGARCH models. Finally, in a last part of our analysis, we outline a simple long-short investment strategy based on the identified seasonal pattern and compare its risk-adjusted performance (as measured by the classic reward-to-risk ratio proposed by Sharpe, 1966) to a passive buy-and-hold strategy.[3]

The remainder of the article is organised as follows: Section 2 discusses our econometric methodology to test for daily seasonality. Section 3 describes key properties of our dataset. Section 4 contains the empirical analysis, including reported results, while Section 5 concludes.

2. Methodology

Our econometric framework is based on the well-known GARCH model of Bollerslev (1986) because empirical results show that this simple model design is usually sufficient to capture the volatility dynamics of most time series and even proves to be superior to more advanced model variants (see, for example, Hansen and Lunde, 2005). We extend its basic specification in two ways. First, we add lagged oil returns to the mean equation in order to capture potential serial correlation in returns. This way we obtain a specification similar to ARMA-GARCH models (see Mohammadi and Su, 2010; Narayan and Narayan, 2007). Second, we add dummy-variables to both the mean and the variance equation because this allows in testing for seasonal effects in daily returns and volatilities (see Bhattacharya et al., 2003; Lucey and Tully, 2006). These modifications lead to an AR(m)-GARCH(p,q) model of the following form[4]:

$$R_t = \alpha + \sum_{i=1}^{m} \beta_i R_{t-i} + \sum_{j=1}^{4} \gamma_j D_{j,t} + \epsilon_t \tag{1}$$

$$\epsilon_t | \psi_{t-1} \sim N(0, h_t) \tag{2}$$

$$h_t = \delta + \sum_{k=1}^{p} \zeta_k \epsilon_{t-k}^2 + \sum_{l=1}^{q} \eta_l h_{t-l} + \sum_{j=1}^{4} \theta_j D_{j,t}. \tag{3}$$

Here, the log return R_t of crude oil at day t is considered to be linearly related to its lagged values R_{t-i}, $i = 1, \ldots, m$, four dummies $D_{j,t}$, $j = 1, \ldots, 4$, and an error term ϵ_t. The dummy $D_{1,t}$ ($D_{2,t}$, $D_{3,t}$, $D_{4,t}$) takes the value 1 if the day t is a Tuesday (Wednesday, Thursday, Friday), and is 0 otherwise. Monday is resembled by $D_{1,t} = D_{2,t} = D_{3,t} = D_{4,t} = 0$. ϵ_t depends on past information ψ_{t-1} and is assumed to follow a conditional normal distribution. Thus, our model allows a conditionally heteroscedastic error distribution and, as a direct consequence, even captures fat-tailed behaviour (see Tsay, 2005, chpt. 3.5). As in the classic GARCH model, the conditional variance h_t depends upon the squared residuals ϵ_{t-k}^2, $k = 1, \ldots, p$, of the process and lagged values h_{t-l}, $l = 1, \ldots, q$, of the conditional variance. In our model, it is additionally related to four seasonal dummies.

In order to find an adequate model for capturing the time-series dynamics of crude oil returns, we first determine suitable values for m, p and q in a setup that does not yet contain seasonal dummies. This is achieved by employing a specification procedure similar to Bollerslev (1988) and Choudhry (2000).[5] Its basic idea is to estimate models for different combinations of m, p and q and to evaluate their general descriptive validity. We start with a basic specification that assumes no serial correlation in returns ($m = 0$) but a typical GARCH process in the conditional variance ($p = q = 1$). Model parameters are estimated by the maximum likelihood procedure (see Davidson and MacKinnon, 2004; Greene, 2008, chpts. 13.6 and 19.3). To assess the adequacy of the model, we test the standardised residuals ($\epsilon_t / h_t^{0.5}$) and the squared standardised residuals (ϵ_t^2 / h_t) for serial correlation by means of the Ljung and Box (1978) test. Absence of serial correlation in the standardised (squared standardised) residuals implies the lack of need to encompass a higher order AR (GARCH) process in the mean (variance) equation. In case of serial correlation, we add AR and/or GARCH terms. We continue this procedure until we come up with a specification of (at conventional levels of 1%, 5% or 10%) significant AR (β_i) and GARCH (ζ_k, η_l) parameters that is free of serial correlation.

In a second step, we then extend the identified specification by our four dummy-variables and use their coefficients to test for daily seasonality. As Monday is the reference category, the dummies in the mean and variance equations describe differential effects. This means that, for example, a significantly positive (negative) estimate of γ_4 (θ_4) tells us that returns (volatilities) on Fridays are significantly higher (lower) than those on Mondays.

To evaluate the robustness of our seasonality results to alternative GARCH specifications, we follow Capie et al. (2005), Kang et al. (2009), Mohammadi and Su (2010) and Wei et al. (2010) by also estimating alternative commonly used GARCH models. These models have a number of theoretical advantages and are constructed as follows:

First, the GARCH-M model of Engle et al. (1987) considers the possibility of a tradeoff between risk and returns by adding the conditional standard deviation h_t to the mean equation. Thus, in our application, the mean Eq. (1) takes the form

$$R_t = \alpha + \sum_{i=1}^{m} \beta_i R_{t-i} + \sum_{j=1}^{4} \gamma_j D_{j,t} + \kappa h_t + \epsilon_t, \tag{4}$$

while the variance Eq. (3) remains unchanged. If $\kappa > 0$, then there is a positive tradeoff between risk and return as suggested by portfolio theory.

Second, the TGARCH model proposed by Glosten et al. (1993) and Zakoian (1994) incorporates the feature that negative shocks may give rise to higher volatility than positive shocks of equal magnitude.[6] To consider this in our model, we specify the variance as

$$h_t = \delta + \sum_{k=1}^{p} \zeta_k \epsilon_{t-k}^2 + \sum_{l=1}^{q} \eta_l h_{t-l} + \sum_{v=1}^{r} \lambda_v T_{t-v} \epsilon_{t-v}^2 + \sum_{j=1}^{4} \theta_j D_{j,t}, \tag{5}$$

[3] Schuhmacher and Eling (2011, 2012) highlight that the Sharpe ratio has a decision-theoretic foundation for a wide range of (skewed and fat-tailed) non-normal distributions.

[4] For simplicity, our theoretical model description does not consider potential MA terms in the mean equation because they turn out to have no significant impact on our empirical results.

[5] Note that most recent studies do not apply specification searches to find optimal values for m, p and q (see Kang et al., 2009; Mohammadi and Su, 2010; Narayan and Narayan, 2007). They set $m = 0$ and/or $p = q = 1$. This way, the mean and/or variance equations may be misspecified and the possibility that higher order variants may have better explanatory power is not considered.

[6] The EGARCH model of Nelson (1991) would be an alternative to model asymmetric effects of shocks on the conditional variance. For a recent application, see Narayan and Narayan (2007).