Energy Economics 33 (2011) 321–337

Contents lists available at ScienceDirect

Energy Economics

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

Forecasting petroleum futures markets volatility: The role of regimes and market conditions

Nikos K. Nomikos *, Panos K. Pouliasis

Faculty of Finance, Cass Business School, London EC1Y 8TZ, United Kingdom

A R T I C L E I N F O

Article history:
Received 27 July 2009
Received in revised form 15 November 2010
Accepted 21 November 2010
Available online 27 November 2010

JEL classification:
C32
C53
G10
G14
Q40

Keywords:
Petroleum markets
Regime-dependent volatility
Forecasting
Reality check
Value-at-risk

A B S T R A C T

In this paper we employ regime volatility models to describe time dependency in petroleum markets. Using a sample of NYMEX and ICE futures contracts, we establish the existence of a regime process and link this process to market fundamentals. This formulation results in two distinct states: a highly persistent conditional volatility process, characterised by long memory and low sensitivity to market shocks, and a relatively short-lived nonstationary process with less memory but higher sensitivity to shocks. Moreover, to investigate the relationship between disequilibrium and volatility of oil futures across high and low volatility regimes we use augmented regime GARCH models to address in a realistic way the potential diverse response of volatility to forward curve shocks. The performance of these models is compared to benchmarks, using both statistical tests and risk management loss functions. To test the robustness of the forecasting strategies, we also perform a reality check employing the stationary bootstrap approach. The findings of this paper have important implications for decision making concerning trading and risk management, as well as energy market operations, such as refining and budget planning, by providing valuable information on the oil price volatility dynamics and the ability to predict risk.

1. Introduction

In the volatile world of energy markets, quantifying and mitigating price risk presents a number of challenges due to the time-dependence in volatility, non-linear dynamics and heavy tails in the distribution of oil returns. Petroleum price volatility has always been at the core of economic research agenda not only because of its effect on the cash flows of oil-related businesses, but also due to the far-reaching implications of oil price uncertainty on the macroeconomy (Hamilton, 2003 and Chen and Chen, 2007) and the financial markets (Driesprong et al., 2008 and Aloui and Jammazi, 2009). It is not surprising therefore that in the energy economics literature there is a plethora of empirical studies examining the issue of modelling volatility and risk management.

Traditionally, the family of Autoregressive Conditional Heteroscedasticity (ARCH) models – introduced by Engle (1982) – have been widely used to describe the conditional volatility of oil prices, due to their flexibility. However, empirical research suggests that in the presence of asymmetries, fat tails and time-dependent higher order moments, the standard Generalised ARCH model of Bollerslev (1986) is not appropriate and thus, numerous extensions have been developed in the literature either by assuming different distributions of the error structure or by adding asymmetric terms, such as leverage effects, in the variance process. Kang et al. (2009) for instance, compare the forecasting ability of different GARCH models in the WTI, Brent and Dubai crude oil futures markets and find that Fractionally Integrated GARCH processes provide more accurate volatility forecasts, concluding that persistence and long memory are essential elements of energy markets volatility. Agnolucci (2009) investigates the market volatility of WTI futures and finds that extensions of GARCH models with asymmetric effects and different error distributions out-perform implied volatility models' predictive accuracy. Fan et al. (2008) show that the assumption of normality leads to underestimation of risk and GARCH models based on the Generalised Error Distribution (GED) produce more reliable forecasts compared to ordinary GARCH models. Hung et al. (2008) also highlight the importance of selecting the appropriate distribution in a GARCH context and find that crude oil and oil products' Value-at-Risk (VaR) is better captured by fat-tail distributions. Overall, the findings of this study imply that the assumption of fat tails plays an important role in VaR estimates since it directly affects the required quantiles. Costello et al. (2008) on the other hand, employ a GARCH filter and rely on historical simulations (semi-parametric GARCH) to forecast VaR

* Corresponding author. Faculty of Finance, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK. Tel.: +44 20 7040 0104.
E-mail addresses: n.nomikos@city.ac.uk (N.K. Nomikos), panos.pouliassis.1@city.ac.uk (P.K. Pouliasis).

0140-9883/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.eneco.2010.11.013

whereas Huang et al. (2009) employ an alternative CAViaR (Conditonal Autoregressive VaR) technique based on regression quantiles. Other studies testing different variants of GARCH models include Duffie and Gray (1995), Sadorsky (2006), Cheong (2009) and Wei et al. (2010).

A major shortcoming of GARCH models is that they induce a high degree of persistence in shocks, that falsely implies high predictability but, in essence reflects regime shifts or structural breaks in the volatility process (Lamoureux and Lastrapes, 1990). This means that a regime-switching GARCH model may be more suitable for modelling volatility particularly in the energy markets where structural breaks are quite common.[1] Another advantage of a regime GARCH process is its ability to deal with fat-tails (see Haas et al., 2004a, 2004b for more details and derivation of higher moments of mixed normal distributions); this is very important for modelling volatility in the oil futures markets where demand shocks result in an asymmetrically higher volatility when the market is at the steep part of the supply stack.

In addition, oil market volatility is characterised by different dynamics under different market conditions. For instance, Fong and See (2002, 2003) document strong evidence of regime switching in the temporal volatility dynamics of oil futures, consistent with the theory of storage; an increase in backwardation is more likely to increase regime persistence in the high volatility state, due to low inventories. Alizadeh et al. (2008) employ a Markov Regime Switching (MRS) approach for determining optimum hedge ratios in NYMEX energy futures markets. In a low variance regime, error correction coefficients are in accordance with convergence towards a long-run equilibrium relationship, while the high variance state is characterised by insignificant speed of adjustment coefficients, which effectively results in a widening of the basis thus explaining the high variance regime; hence, the adjustment process undergoes regime shifts and does not behave uniformly to shocks to equilibrium across different states. Another study by Vo (2009) combined the concept of regime switching with that of stochastic volatility to forecast the dynamics of WTI crude oil. The author finds that the simple MRS model captures better the in-sample dynamics in terms of mean absolute errors whereas out-of-sample, stochastic volatility with regime shifts is favoured.

Building on these previous studies, this paper investigates the volatility dynamics for the NYMEX WTI crude and heating oil as well as the ICE Brent crude and gas oil futures contracts. In doing so, it contributes to the existing literature in a number of ways. First, we employ various volatility regime models, to accommodate some of the stylised features of the oil markets such as volatility clustering, non-normality, time-varying skewness and excess kurtosis. In particular, we consider the Mix (distribution) GARCH and the MRS GARCH models based on the mixed conditional heteroscedasticity models of Haas et al. (2004a) and Alexander and Lazar (2006) and the Markov model of Haas et al. (2004b), respectively. Our study is different from the above mentioned research in the sense that we provide a thorough empirical application of the provided framework in the energy markets. Although volatility modelling and forecasting in a regime framework has been widely documented in equity and foreign exchange markets (see Marcucci, 2005; Li and Lin, 2004; Giannikis et al., 2008), few studies have analysed in depth the nature of the volatility regimes of oil futures prices and the forecasting ability of those models in the specific market.

Second, we extend previous research by including the squared lagged basis of futures prices in the specification of the conditional variance in what is termed the GARCH-X model (Lee, 1994; Ng and Pirrong, 1996). A principal feature of the basis is that the time paths of spot and futures prices are influenced by the extent of deviations from their long-run equilibrium (Engle and Granger, 1987). As prices respond to the magnitude of disequilibrium then, in the process of adjusting, they may become more volatile. If this is the case then the inclusion of the basis term in the conditional variance specification may lead to the estimation of more accurate volatility forecasts. Examining different volatility components will enable us to investigate whether the dependence of volatility to the basis changes across different regimes and uncover how these asymmetries are transmitted. To the authors knowledge this is the first time that the GARCH-X methodology is tested in a regime volatility setting. Implementing such models allows us to draw some new interesting insights regarding the effect of disequilibrium and the persistence of volatility under different market conditions.

Third, we extend the above framework to a conditional extreme value theory (EVT) setting and use the estimated volatility models as filters, in order to combine the forecasts with EVT-based methods for quantile estimation and link the regime volatility background with tail estimation. From a risk management perspective, the tails of the conditional distributions of the models may contain important information that needs to be considered. Existing literature that addresses the issue is limited to the EVT-Switching ARCH model of Samuel (2008), applied in estimating VaR in the stock index market. In the oil market there is limited evidence on conditional EVT based VaR provided by Krehbiel and Adkins (2005) for the NYMEX complex and Marimoutou et al. (2009) for WTI and Brent crude oil.

Fourth, the forecasting performance of the proposed models is assessed and contrasted using a battery of forecast statistics which measure both the tracking errors from actual volatility measures, as well as the degree of volatility under or over-prediction. In addition, we evaluate the effectiveness of the proposed models in VaR applications for both long and short positions and this way, we provide robust evidence on the performance of the proposed volatility models. VaR forecasts are assessed by means of risk management loss functions and their relative performance is ranked using White's (2000) Reality Check.

Finally, volatility and VaR forecasts are tested across periods of backwardation and contango. Many authors (see Fama and French, 1987 and Geman and Ohana, 2009) have shown that price volatility has a negative correlation with inventory levels, in line with the theory of storage. Consequently, it is worth examining the performance of different models under conditions of backwardation and contango, since the risk-return profile of energy prices is known to change fundamentally, between the two different states.

The remainder of this paper is organised as follows. Section 2 demonstrates the Regime GARCH models estimation procedure. In Section 3, the data and their properties are discussed. This is followed by an evaluation of the proposed strategies in Section 4. Finally, conclusions are given in the last section.

2. Methodology

To estimate the volatility models, the methodology used in this study follows the Mix-GARCH model of Haas et al. (2004a) and Alexander and Lazar (2006) and the MRS-GARCH model of Haas et al. (2004b). Both assume more than one individual component variances and differ in the way that they treat regime probabilities. For the former, what is important is the overall regime probability; for the latter, the probability of each observation belonging to any given regime is more important. However, both models assume that asset returns are generated from different information distributions and in this regard, they can accommodate parameter shifts or switches among a finite number of regimes; this is expected to improve the performance of these models in financial applications, such as VaR.

[1] See for instance Wilson et al. (1996). Employing an iterative cumulative sums-of-squares (ICSS) approach, they show evidence of sudden changes in the unconditional volatility of oil futures contracts. In particular, 15 significant volatility changes were detected from 1984 to 1992, whereas 5 of these exceeded 100% in absolute terms e.g. the eight day period following the invasion of Kuwait in 1991 was associated with a 213 percent upward change in the unconditional volatility.