Expert Systems With Applications 55 (2016) 222–242

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Combining high frequency data with non-linear models for forecasting energy market volatility[☆]

Jozef Baruník[a,b,*], Tomáš Křehlík[a,b]

[a] Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague, Opletalova 26, 110 00, Czech Republic
[b] Institute of Information Theory and Automation, The Czech Academy of Sciences, Pod Vodarenskou Vezi 4, 182 00, Prague, Czech Republic

ARTICLE INFO

JEL Classification:
C14
C53
G17

Keywords:
Artificial neural networks
Realized volatility
Multiple-step-ahead forecasts
Energy markets

ABSTRACT

The popularity of realized measures and various linear models for volatility forecasting has been the focus of attention in the literature addressing energy markets' price variability over the past decade. However, there are no studies to help practitioners achieve optimal forecasting accuracy by guiding them to a specific estimator and model. This paper contributes to this literature in two ways. First, to capture the complex patterns hidden in linear models commonly used to forecast realized volatility, we propose a novel framework that couples realized measures with generalized regression based on artificial neural networks. Our second contribution is to comprehensively evaluate multiple-step-ahead volatility forecasts of energy markets using several popular high frequency measures and forecasting models. We compare forecasting performance across models and across realized measures of crude oil, heating oil, and natural gas volatility during three qualitatively distinct periods: the pre-crisis period, the 2008 global financial crisis, and the post-crisis period. We conclude that the newly proposed approach yields both statistical and economic gains, while reducing the tendency to over-predict volatility uniformly during all the tested periods. In addition, the proposed methodology is robust to a substantial structural break induced by the recent financial crisis. Our analysis favors median realized volatility because it delivers the best performance and is a computationally simple alternative for practitioners.

1. Introduction

Predicting energy price variability has become one of the most significant issues faced by the natural gas industry and energy companies in recent decades. With their considerable volatility, the leading products of energy markets, i.e., crude oil, natural gas, and heating oil,[1] contributed to a climate of uncertainty and distrust of energy companies and investors, on one hand, and of consumers, regulators, and legislators, on the other. The high level of volatility in energy markets is likely due to supply uncertainty—such as from a variety of macroeconomic and political factors in the case of crude oil or simply storage constraints in the case of natural gas—and short-term inelasticity of demand, i.e., the difficulty of reducing consumption within a short period of time. The combination of these two factors makes it extremely difficult for both consumers and producers to forecast their costs and profits. The desire to protect market participants against the losses resulting from this unpredictability has led to immense interest in empirical research aiming to predict the variability in energy prices. In this paper, we contribute to this literature by proposing a novel framework to forecast energy commodity volatility that couples realized measures with generalized regression based on artificial neural networks. We demonstrate that our approach delivers precise forecasts even in the regime-switching moment of financial crisis.

Volatility research from previous decades is influenced mainly by the work of Engle (1982) and Bollerslev (1986, 1987) and has shown that price variability is much easier to understand than it is to forecast the direction of future price changes. However, the lion's share of previous research has focused on the financial markets, and the focus has only recently turned to the energy

☆ Support from the Czech Science Foundation under project no. P402/12/G097 DYME – "Dynamic Models in Economics" is gratefully acknowledged. Křehlík gratefully acknowledges financial support from the Grant Agency of Charles University under projects 588314 and 837413.

* Corresponding author at: Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague, Opletalova 26, 110 00, Czech Republic. Tel.: +420 776 259273.

E-mail addresses: barunik@fsv.cuni.cz (J. Baruník), tomas.krehlik@gmail.com (T. Křehlík).

[1] According to the CME Group Leading Products Resource, crude oil, natural gas, and heating oil futures are traded with the highest average volume among energy commodities (http://www.cmegroup.com/education/featured-reports/cme-group-leading-products.html).

http://dx.doi.org/10.1016/j.eswa.2016.02.008
0957-4174/

markets.[2] (Kang & Yoon, 2013; Kuper & van Soest, 2006; Linn & Zhu, 2004; Mohammadi & Su, 2010; Pindyck, 2004; Sévi, 2014; Wei, Wang, & Huang, 2010; Wilson, Aggarwal, & Inclan, 1996; Yang, Hwang, & Huang, 2002).

More recent advances in financial econometrics have led to the development of new estimators of volatility using high frequency data that make volatility observable. Although the pioneering studies in the realized volatility literature recognize the benefits of using high frequency data in terms of increased accuracy (Merton, 1980; Zhou, 1996), subsequent research[3] proposes several estimators that improve model efficiency, robustness to market microstructure effects, and the ability to separately estimate the variation due to the continuous part of the price process, on one hand, and the variation due to the jump part of the price process, on the other. For excellent reviews of the realized volatility literature, see Andersen, Bollerslev, Christoffersen, and Diebold (2006); McAleer and Medeiros (2008), or Barndorff-Nielsen and Shephard (2007). Moreover, recent studies utilize high frequency data in energy markets (Baum & Zerilli, 2016; Prokopczuk, Symeonidis, & Wese Simen, 2015).

However, estimating realized volatility is only the first step to more accurate predictions and using the appropriate model is the second step. Heterogeneous autoregressive (HAR) and autoregressive fractionally integrated (ARFIMA) models became widely used to forecast realized volatility because these models effectively capture the long memory of volatility (Andersen, Bollerslev, Diebold, & Labys, 2003; Corsi, 2009). In contrast to FIGARCH models that capture the long memory of volatility using daily returns data,[4] these approaches are more flexible and easier to estimate when high frequency data are available. Although both the HAR and the ARFIMA have been developed to capture the specific long memory feature of volatility, more complex patterns may be revealed and explored. Changes in market conditions and many types of noises induced by measurement error lead to non-linear patterns that cannot be captured by linear models that are based on restrictive distributional assumptions. Microstructure noise that can arise through the bid-ask bounce, asynchronous trading, infrequent trading or price discreteness are important examples of measurement error.

Artificial neural networks (ANN) may be understood as a generalization of these classical approaches that may help to uncover more complex volatility patterns. Concisely, neural networks are semi-parametric non-linear models that can approximate any reasonable function (Haykin, 2007; Hornik, Stinchcombe, & White, 1989). The number of models using machine learning is rapidly growing in the academic literature but applications that apply these models in energy market in energy markets are limited. Among the few that do, Fan, Liang, and Wei (2008) proposes a generalized pattern matching based on a genetic algorithm to predict crude oil prices on a multi-step-ahead basis. Xiong, Bao, and Hu (2013); Yu, Wang, and Lai (2008) proposes an empirical model that decomposes neural networks to forecast crude oil prices. Jammazi and Aloui (2012) uses a hybrid model for crude oil forecasting, Panella, Barcellona, and D'Ecclesia (2012) use a mixture of Gaussian neural network to forecast energy commodity prices, and Papadimitriou, Gogas, and Stathakis (2014) investigates the efficiency of support vector machines in forecasting next-day electricity prices. Although the focus has remained solely on forecasting prices, research using neural networks to forecast volatility continues to be developed.

This paper's primary contribution is that it proposes a model that couples measures of volatility from high frequency data with artificial neural networks to reliably forecast energy price volatility. Whereas researchers in financial econometrics have performed pioneering work using stock market index data (McAleer & Medeiros, 2011) or exchange rate data (Sermpinis, Theofilatos, Karathanasopoulos, Georgopoulos, & Dunis, 2013), we are the first to comprehensively test this strategy against competing models in the energy literature. Rather than choosing from among the plethora of advanced machine learning algorithms, we use the simplest and most popular feed-forward neural network as the first step in this field. Our main motivation is to show whether there are statistical and economic gains that can be realized by coupling high frequency data with easy-to-implement artificial neural networks.

This paper also contributes to the literature by comprehensively evaluating the most popular models and realized measures. These realized volatility measures rely on different assumptions, and there are no studies guiding practitioners to use a specific measure when working with the volatility forecasting of energy markets. To bridge this gap, we focus on the three most liquid energy commodities—crude oil, heating oil, and natural gas—during the period from January 5, 2004 to December 31, 2012 and put the models into a horse race through several discrete sub-periods to determine which model produces uniformly lower errors in multiple-step-ahead volatility forecasts. The period under study is especially interesting because it includes a sub-period of high and rapidly rising prices, a sub-period encapsulating the interruption of price increases in 2008 due to global turmoil in the financial markets, and the last sub-period that witnessed profound regime change over the most recent few years in which price variability became much calmer. In particular, the last period is particularly interesting from the forecaster's perspective, as it appears that demand for liquid transport fuels has peaked in the developed economies with car engines becoming more efficient and amid partial substitution by biofuels. On the supply side, high prices reversed the previous trend toward growing dependence on the conventional oil fields of the OPEC member states. Sophisticated modeling strategies should reflect these changes.

We test the ANN against widely used the HAR and ARFIMA long-memory models and a benchmark low frequency-based GARCH model. The tests are performed within the recently proposed frameworks of the Model Confidence Set (MCS) developed by Hansen, Lunde, and Nason (2011) and Superior Predictive Ability (SPA) developed by Hansen (2005) with several popular loss functions used in the literature. Moreover, we use realized variance (RV), realized kernel (RK), two-scale realized variance (TSRV), bipower variation (BV), median realized volatility (MedRV), and the recently proposed jump-adjusted wavelet two-scale realized variance (JWTSRV) as measures of volatility. Motivated by the possibility of reducing model uncertainty, we also experiment with the linear combination of forecasts from the popular HAR model and artificial neural network. This experiment yields the lowest error uniformly through all tested periods regardless of which realized measure is used. These low error levels also translate to economic benefits in terms of Value-at-Risk. One of the loss functions we use in the exercise allows us to assess whether the models tend to over-predict volatility as commonly found using GARCH-type models[5]. A uniform finding is that coupling neural networks with high frequency data results in substantial reductions in the

[2] For a complete review of GARCH-type models used in the energy literature, see Wang and Wu (2012).

[3] Andersen and Bollerslev (1998); Andersen, Bollerslev, Diebold, and Labys (2001, 2003); Bandi and Russell (2006); Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008); Hansen and Lunde (2006); Zhang, Mykland, and Aït-Sahalia (2005).

[4] Kang and Yoon (2013) recently investigate the ability of FIGARCH models to capture energy market volatility.

[5] For example, see Nomikos and Pouliasis (2011), who confirm the strong tendency of GARCH-type models to over-predict the volatility of crude oil, heating oil, and gasoline, which is further confirmed by Wang and Wu (2012), who find that multivariate GARCH-type models also suffer from over-predictions.