Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Forecasting the real price of oil under alternative specifications of constant and time-varying volatility

CAMA Working Paper 71/2017
November 2017

Beili Zhu
Centre for Applied Macroeconomic Analysis, ANU

Abstract

This paper constructs a monthly real-time oil price dataset using backcasting and compares the forecast performance of alternative models of constant and timevarying volatility based on the accuracy of point and density forecasts of real oil prices of both real-time and ex-post revised data. The paper considers Bayesian autoregressive and autoregressive moving average models with respectively, constant volatility and two forms of time-varying volatility: GARCH and stochastic volatility. In addition to the standard time-varying models, more flexible models with volatility in mean and moving average innovations are used to forecast the real price of oil. The results show that time-varying volatility models dominate their counterparts with constant volatility in terms of point forecasting at longer horizons and density forecasting at all horizons. The inclusion of a moving average component provides a substantial improvement in the point and density forecasting performance for both types of time-varying models while stochastic volatility in mean is superfluous for forecasting oil prices.

Keywords

Forecasting, oil price, real-time data, time-varying volatility, moving average, stochastic volatility in mean

JEL Classification

C11, C53, C82, Q43

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

Forecasting the real price of oil under alternative specifications of constant and time-varying volatility

Beili Zhu

Centre for Applied Macroeconomic Analysis (CAMA)

Crawford School of Public Policy, The Australian National University

Abstract

This paper constructs a monthly real-time oil price dataset using backcasting and compares the forecast performance of alternative models of constant and time-varying volatility based on the accuracy of point and density forecasts of real oil prices of both real-time and ex-post revised data. The paper considers Bayesian autoregressive and autoregressive moving average models with respectively, constant volatility and two forms of time-varying volatility: GARCH and stochastic volatility. In addition to the standard time-varying models, more flexible models with volatility in mean and moving average innovations are used to forecast the real price of oil. The results show that time-varying volatility models dominate their counterparts with constant volatility in terms of point forecasting at longer horizons and density forecasting at all horizons. The inclusion of a moving average component provides a substantial improvement in the point and density forecasting performance for both types of time-varying models while stochastic volatility in mean is superfluous for forecasting oil prices.

Keywords: forecasting; oil price; real-time data; time-varying volatility; moving average; stochastic volatility in mean

JEL classification: C11, C53, C82, Q43

* The author thanks Professor Renée Fry-McKibbin and Professor Joshua Chan for supervision and valuable suggestions. The author gratefully acknowledges ARC Discovery Project DP120103443 for financial support. The author also thanks CAMA seminar participants for the useful comments.

1 Introduction

Accurately forecasting the real price of oil is important due to the fact that the oil price affects various economic activities and thus the economic outlook worldwide. The oil price influences natural resource development, the manufacturing industry as well as oil importing and exporting industries. Therefore, it is not surprising that central banks and governments heavily rely on oil price forecasts to assess the general economic outlook. Figure 1 depicts the movements of the two types of real oil prices from January 1974 to December 2014. The real oil price has fluctuated dramatically over the last forty years, especially after the global financial crisis in 2007. Because of the volatile oil price, forecasting the real price of oil is not easy for forecasters. Given the importance and difficulty of oil price forecasting, this paper compares the oil price forecasting performance of alternative models with constant and time-varying volatility. In addition, new forecasting methods are introduced to forecast the oil price. Specifically, the paper extends the standard time-varying volatility model by allowing for MA components and volatility in the conditional mean. Another contribution of this paper is that in addition to the point forecast, density forecasts are used to measure the accuracy of the forecast performance amongst the different models. The paper also constructs a real-time dataset of variables required in forecasting the real price of oil using backcasting methods and extends the real-time oil price series to 2015. Both the real-time and ex-post revised data are used for the forecasting comparisons.

There is a voluminous literature on forecasting both the real and nominal oil price. Two strands of forecasting methods of the oil price co-exist: one is regression based (Baumeister and Kilian, 2011; Baumeister and Kilian, 2012) while the other is survey based (Chernenko, Schwarz and Wright, 2004; Sanders, Manfredo and Boris, 2008). When it comes to regression based methods, existing studies on oil price forecasting mainly use econometric models with constant variance and do not allow for volatility clustering, which is a prominent feature of oil data (Yea, Zyrenb and Shore, 2004; Baumeister and Kilian, 2012; Baumeister and Kilian, 2014a). However, the oil price is highly volatile and subject to dramatic shocks. Figure 2 shows the percentage change in monthly real U.S. refiners acquisition cost for imports and monthly real WTI price from January 1974 to December 2014. Both types of oil price exhibit larger swings in both directions during the half century. The assumption of constant variance cannot capture these features. Most papers in the literature forecast the volatility of the oil price using two classes of time-varying volatility models. One is the general autoregressive conditional heteroscedastic (GARCH) model and its variants (Sadorsky, 2005; Kang, Kang and Yoon, 2009; Mason and Wilmot, 2014), while the other is the stochastic volatility (SV) model (Vo, 2009; Larsson and Nossman, 2011). However, forecasting of the real oil price itself using time-varying volatility is rare. This paper aims to fill this gap by using both GARCH and SV models to forecast the real oil price and compares the forecasting results with those from constant volatility models.

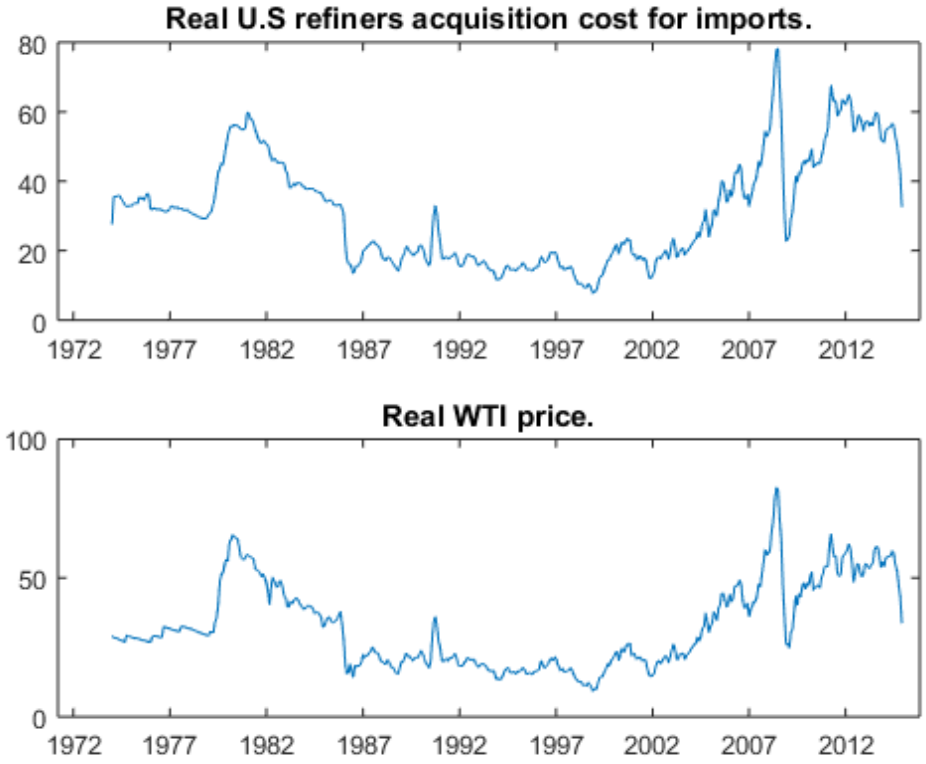

Figure 1: The two types of ex-post revised oil prices from January 1974 to December 2014.

Chan (2013) shows that moving average (MA) stochastic volatility models provide better out of sample forecast performance than standard variants which only include stochastic volatility for U.S. inflation. This paper investigates the potential forecasting improvement for the real oil price by adding the Chan (2013) MA components to both the SV and GARCH models. When it comes to the stochastic volatility in mean (SVM) model as developed by Koopman and Uspensky (2002), volatility is added to both the conditional mean and the conditional variance. The feature of the volatility entering the conditional mean as a covariate may be important for improving the accuracy of real oil price forecasts since volatility clustering is a prominent feature in the real oil price data. However, the SVM model has not been used to date to forecast the real oil price. This paper tests the forecast performance of the SVM model for the real-time oil price.

Most of the papers forecasting the real oil price compare alternative models based on point forecasts and measure the accuracy of the point forecasts using the mean square forecast error (MSFE), which assesses the ability of different models to correctly predict the central tendency (Yu, Wang, and Lai, 2008; Alquist, Kilian and Vigfusson, 2011;

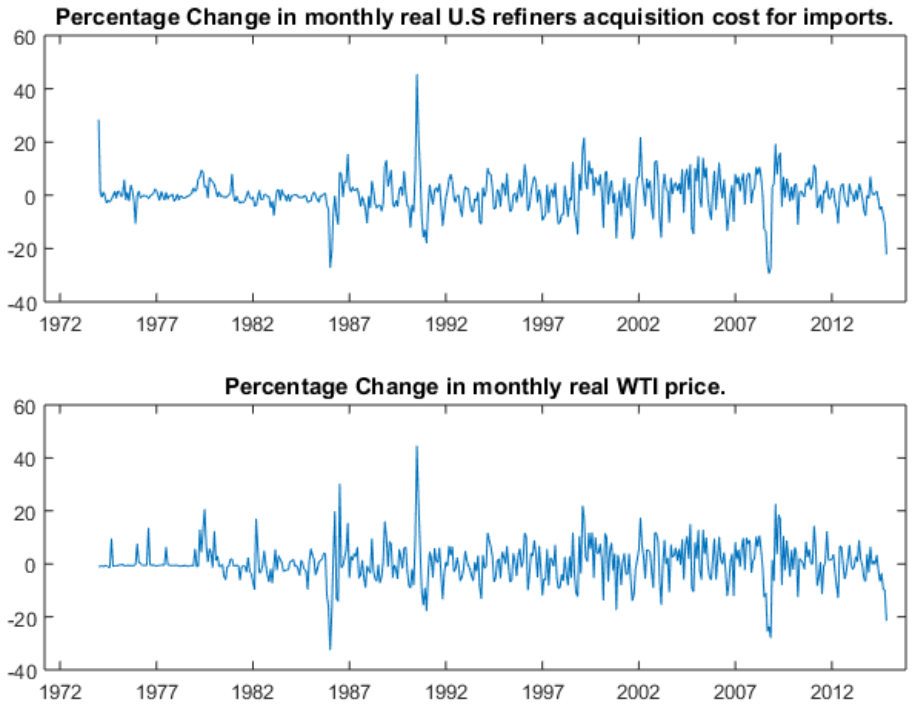

Figure 2: The two types of ex-post revised oil prices from January 1974 to December 2014.

Baumeister and Kilian, 2012; Baumeister and Kilian, 2014a). However, many recent papers forecasting macroeconomic and financial market data compute the predictive likelihoods as the density forecast to assess and compare the forecast performance of models (Geweke and Amisano, 2011; Chan, 2013; Chan, 2017). Using predictive likelihoods is the most natural way to assess and compare the forecast performance of models in a Bayesian approach (Geweke and Amisano, 2011). This paper follows this literature and also uses the predictive likelihood as an alternative forecast evaluation metric to the MSFE to evaluate and compare the forecasting performance of alternative models.

Two sorts of oil price data are used in literature on oil price forecasting. One is ex-post revised data (Knestch, 2007; Alquist, Kilian and Vigfusson, 2011) and the other is real-time data (Baumeister and Kilian, 2012; Baumeister and Kilian, 2015). Economic activity is determined by real time data. However, many forecasts of the oil price and hence the macroeconomy are formed using ex-post revised data due to the unavailability of real-time oil price data (Fair and Shiller, 1990; Baumeister and Kilian, 2012). A real time oil price series as constructed by Baumeister and Kilian (2012) exists from 1974 to 2010. However, the time series is not publicly available. Following the real-time oil price

construction process proposed by Baumeister and Kilian (2012), the real time data in this paper is constructed and extended to February 2015.

The Point and density forecast performance of eight models of the oil price are compared relative to the random walk. The models used here are three commonly used constant volatility models, including the Random Walk (RW) model, the autoregressive (AR) model, and the autoregressive-moving-average (ARMA) model; the remaining five models account for time-varying volatility, including stochastic volatility with a stationary AR process, stochastic volatility with an ARMA Process, stochastic volatility in mean with constant parameters, the standard GARCH model, and the GARCH model with moving average innovations.

The results show that for the forecasts using real time data, the time-varying volatility models dominate their counterparts with constant volatility, in terms of point forecasting at the longer horizon and density forecasting at all horizons. In addition, the moving average component provides a substantial improvement to the point and density forecasting performance for both types of time-varying models while the stochastic volatility in mean model is superfluous for forecasting oil prices. The main results for the ex-post data are slightly different to those using the real-time data. The results using the ex-post data provide strong evidence that including time-varying volatility consistently improves the accuracy of point forecasts relative to models with constant volatility at all horizons.

The paper proceeds as follows. Section 2 describes the process of collecting and constructing the real-time and ex-post oil datasets. Section 3 introduces the constant volatility and time-varying volatility models. Section 4 gives an overview of the forecast metrics and compares not only the point but also the density forecasting performance of candidate models for both the real-time and the ex-post real oil prices. Lastly, a conclusion is given in Section 5.

2 Data

Although the real oil price is one of the most important variables in model-based macroeconomic projections generated by private industry forecasters, central banks, and international institutions, most forecasting studies are done on the basis of the ex-post revised oil prices (Morana, 2001; Knetsch, 2007). One reason for this is that real-time datasets of variables required in forecasting the real price of oil are not readily available and are not easily constructed (Baumeister and Kilian, 2012). However, macroeconomic forecasts are affected by the use of real-time data (Fair and Shiller, 1990). Stark and Croushore (2002) find that the use of real-time macroeconomic data may positively affect forecast performance. Therefore, in recent years, there has been increased interest in real-time forecasts of the real oil price at least for horizons up to one or two years. Baumeister

and Kilian (2012) construct a monthly real-time dataset by using backcasting and nowcasting techniques over the period 1974 to 2010 then compare the forecasting results with those based on ex-post revised data. In a later paper they then analyse the forecast performance using the real-time real oil price through forecast scenarios (Baumeister and Kilian, 2014b), and product spreads (Baumeister, Kilian and Zhou, 2015). These three papers discuss the process of constructing the real-time dataset for the real oil price. However, the authors do not share the real-time dataset of the variables required for forecasting the oil price with the public. This paper follows the construction process of Baumeister and Kilian (2012) to extend the data to February 2015. This section documents the process of constructing the real-time data and then also describes the ex-post data.

2.1 The construction of the real-time oil price data

Two types of oil price series are collected and analysed. The monthly averages of the West Texas Intermediate (WTI) spot price is obtained from the database of the Federal Reserve Bank of St. Louis and the Energy Information Administration (EIA) database. The WTI spot price was collected by the Federal Reserve Bank of St.Louis until July 2013 when it was discontinued. This paper uses the monthly WTI prices until July 2013 from the FRED database. Data after July 2013 is obtained from the EIA. The WTI spot price of crude oil is available in real time and is not subject to data revisions.

An alternative oil price series to the WTI is the monthly U.S. refiners' acquisition cost for crude oil imports which is also collected for analysis in this paper. According to Baumeister, Kilian and Zhou (2015), the U.S. refiners' acquisition cost for crude oil imports is a better proxy than the WTI price for the global price of crude oil that is published by the Monthly Energy Review of the EIA. However unlike the WTI price, the U.S. refiners' acquisition cost for crude oil imports is not available in real-time and is available with a delay and subject to revisions.

This paper first constructs the real-time dataset of the monthly U.S. refiners' acquisition cost for crude oil imports consisting of monthly vintages from 1994.2 through 2015.2, each covering data extending back to 1974.1. Issues of the Monthly Energy Review are only available from January 1993 in the EIA database. Each vintage as reported by the EIA only includes data for a maximum of 3 years. For example, the 1993.1 issue includes data as far back as 1991.1, which means that the pre-1991.1 data for all vintages needs to be approximated. As data moves from one Monthly Energy Review to the next, some of the earlier observations cease to be reported. Following the real-time dataset construction process used by Baumeister and Kilian (2012), a backcasting approach is introduced to fill the gaps. Specifically, the pre-1991.1 data for all vintages are approximated using the ex-post revised data. What's more, the gaps in the current vintages due to the fact that the earlier observations are no longer reported after 3 years

are filled with the most recent data from earlier vintages. See Baumeister and Kilian (2012) for a more detailed discussion of backcasting.

There are some missing observations in the real-time variables owing to the fact that the monthly U.S. refiners' acquisition cost for crude oil imports becomes available only with a delay. For example, the vintages from 1994.2 to 2005.7 report observations for up to $t-3$, and vintages from 2005.8 to 2015.2 report observations up to $t-2$.[1] Baumeister and Kilian (2012) use nowcasting to fill the gaps. However, in many studies of forecasting with real-time macroeconomic variables (Stark and Croushore, 2002; Clark and Ravazzolo, 2014), researchers use real-time data with a delay to examine the forecasting performance of models instead of extrapolating the missing real-time observations by nowcasting. The reason why these studies do not use nowcasting is that the missing real-time observations are extrapolated using a growth rate of a related macroeconomics time series and the extrapolated results may not be accurate. The extrapolated time series may affect the accuracy of the forecasting performance. This paper follows the studies of Stark and Croushore (2002) and Clark and Ravazzolo (2014) and uses real-time U.S. refiners' acquisition costs for crude oil imports with a lag to compare forecasting performance.

The real price of oil is constructed by deflating the nominal price of oil by the U.S. consumer price index. Real-time data for the seasonally adjusted monthly U.S. consumer price index for all urban consumers are obtain from the Real-time Data Set for Macroeconomists complied by the Federal Reserve Bank of Philadelphia and the Economic Indicators published by the Council of Economic Advisers.[2] The real-time CPI data is also available with a lag. In vintage t, the available data runs through period $t-1$.

As discussed by Romer and Romer (2000), Stark and Croushore (2002) and Clark and Ravazzolo (2014), a key question in real-time forecasting exercise is, which vintage of data should be used to represent the 'actual' true data from which forecasting errors are calculated? This paper follows Stark and Croushore (2002) and Baumeister and Kilian (2012) and uses the final release of the real oil price data as the 'actual' data in evaluating forecasting accuracy.

For the real-time U.S. refiners' acquisition cost for imports variable, forecasts are evaluated from 1993.12 to 2014.12, which requires real-time data vintages from 1994.2 to 2015.2, each covering data extending back to 1974.1. The simulation exercise begins with vintage 1994.2, and the parameters of the forecasting model are estimated using the sample 1974.1-1993.11. The paper computes the forecast horizons ranging from 1

[1] Usually when moving from one vintage to the next, the later vintage contains additional observations. However, there are no new observations available in vintage 1994.5 and vintage 2006.9

[2] Vintages for 1998.11 to 2015.2 are obtained from Federal Reserve Bank of Philadelphia. The pre-1998.11 vintages are obtained from the Economic Indicators published by the Council of Economic Advisers

to 12 months ahead outside the estimation window. Then it repeats the above step in a rolling procedure, going forward one month each step, adding one more observation to the sample used for estimation.[3] When the time $t-3$ information is actually incorporated into the models used for forecasting at t, the 1-month-ahead forecast is the month $t-2$ forecast, while the 3-month-ahead forecast is the month t forecast. For the real-time WTI price, the forecasts are evaluated from 1994.2 to 2015.2, which also requires real-time data vintages from 1994.2 to 2015.2, each covering data extending back to 1974.1.

2.2 Ex-post revised data

The most recent data set available at the time of the study is to April 2015. The paper discards the last 4 months of data since the most recent data is still preliminary and may be revised in the following months (Baumeister and Kilian, 2012). The remaining data for 1974.1 – 2014.12 as reported in vintage 2015.4 are treated as the ex-post revised data when evaluating the forecasting accuracy of the candidate models. The evaluation window is 1994.1 to 2014.12. The ex-post revised monthly U.S. refiners' acquisition cost for imports is constructed from the latest available data, and the data downloaded from the data set of the EIA. The ex-post revised seasonally adjusted CPI data is downloaded from Federal Reserve Bank of Philadelphia. As for the real-time oil data, the ex-post real price of oil is constructed by deflating the nominal price of oil by the U.S. consumer price index.

3 The forecasting models

Random walk (RW) and standard AR and ARMA models with constant volatility are easy to use for forecasting (Baumeister and Kilian, 2012; Alquist, Kilian and Vigfusson, 2011). However, a voluminous literature has demonstrated that the volatility of a wide variety of energy series tends to change over time (Narayan and Narayan, 2007; Kang, Kang and Yoon, 2009; Larsson and Nossman, 2011). Allowing for time-variation in the volatility of energy data is important for energy data estimation and forecasting. Therefore, in addition to the standard RW, AR and ARMA models, this paper considers the forecasting performance of models with two classes of time-varying volatility specifications: GARCH models and stochastic volatility models. This section discusses the forecasting models and provides an overview of the Bayesian estimation methods used. The three constant volatility models are discussed in section 3.1, and the five

[3]For vintage 2005.8, the EIA added two more observations instead of one. But for vintage 1994.5 and vintage 2006.9, there are no more new observations

time-varying volatility models are discussed in section 3.2. The three constant volatility models and the five time-varying volatility models are summarized in Table 1.

3.1 Constant volatility models

Constant volatility models are widely used to forecast energy prices and are described in this section. The first model is a random walk model, which is denoted as RW:

$$y_t = y_{y-1} + u_t, \qquad u_t \ \sim \mathcal{N}(0, \sigma^2). \tag{1}$$

In equation (1), y_t is the oil prices, and u_t is the error term which is an independent and identically normally distributed random variable with mean $E(u_t) = 0$ and variance $var(u_t) = \sigma^2$. The RW model is the benchmark model for the analysis.

The second model is an Autoregressive model with p lags:

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t, \qquad u_t \ \sim \mathcal{N}(0, \sigma^2). \tag{2}$$

In equation (2), the y_t and u_t are the same as in the equation (1). β_0 is a constant intercept and the scalar parameters β_i capture the magnitudes of effects of the lagged oil prices $y_{t-i}, i = 1, 2, ..., p,$.

According to Baumeister and Kilian (2012), an AR model with $p = 12$ is common in the literature on oil market regression-based models, so this value is utilized in this model. In addition, the condition that the roots of the characteristic polynomial associated with the AR coefficients all lie inside the unit circle is imposed, so that the AR process is stationary. This model is denoted as AR.

Finally, the third constant volatility model is the autoregressive moving average model:

$$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t, \tag{3}$$

$$\xi_t = u_t + u_{t-1}, \qquad u_t \ \sim \mathcal{N}(0, \sigma^2). \tag{4}$$

In equation (3), the error term ξ_t follows a first-order moving average process and the y_t and u_t are the same as in equation (1).

This model is denoted as ARMA and has a lag order of 1 in both the AR and MA components.

3.2 Time-varying volatility models

In this subsection, two classes of time-varying volatility models are discussed: One is the stochastic volatility model firstly introduced by Taylor (1994) while the other is the generalized autoregressive conditional heteroskedasticity (GARCH) model (Bollerslev, 1986) that is an extension of the pioneering work done on ARCH models by Engle (1982).

3.2.1 Stochastic volatility models

The first specification for the time varying volatility models is the autoregressive model with stochastic volatility:

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t, \qquad u_t \ \sim \mathcal{N}(0, e^{h_t}), \tag{5}$$

for $t = 1, 2, ..., T$. The state h_t is the log-volatility and is assumed to evolve according to a stationary process:

$$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \ \sim \mathcal{N}(0, \sigma_h^2), \tag{6}$$

for $t = 2, ..., T$. Here the paper assumes that $|\phi_h| < 1$ and the states are initialized with $h_1 \ \sim \mathcal{N}(\mu_h, \frac{\sigma_h^2}{(1-\phi_h^2)})$.

In equation (5), the term u_t is an independent and identically normally distributed random variable with mean $E(u_t = 0)$ and variance $var(u_t) = e^{h_t}$. In addition, the volatility process is described by a stationary autoregressive process of order one, given in equation (6). The parameter $|\phi_h| < 1$ (for stationarity) measures the volatility persistence and the error variance σ_h^2 is the volatility of the log-volatility h_t. This model is denoted as AR-SV.

The second specification is the stochastic volatility in mean model developed by Koopman and Uspensky (2002). In this model, volatility enters the conditional mean equation and the volatility appears in both the conditional mean and the conditional variance equations:

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + \alpha e^{h_t} + u_t, \qquad u_t \ \sim \mathcal{N}(0, e^{h_t}), \tag{7}$$

where u_t has the same stochastic volatility specification as before. This model is denoted as AR-SVM. In equation (7), the exponential of log-volatility enters the conditional mean of the observation equation (7) as an additional explanatory variable and the scalar parameter α captures the magnitude of the volatility feedback.

The third specification is the moving average stochastic volatility model by Chan (2013). This model includes both the moving average and stochastic volatility components:

$$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t, \tag{8}$$

$$\xi_t = u_t + \psi u_{t-1}, \qquad u_t \sim \mathcal{N}(0, e^{h_t}), \tag{9}$$

where $u_0 = 0$ and $|\psi| < 1$. The error term ξ_t follows a first-order moving average process. The term u_t is an independent and identically normally distributed random variable with mean $E(u_t) = 0$ and variance $var(u_t) = e^{h_t}$. Again the log-volatility h_t is assumed to follow an AR (1) process as in equation (6). This version of the stochastic volatility model is denoted as ARMA-SV.

3.2.2 GARCH Models

This subsection introduces two types of GARCH models. While volatility is a random variable in stochastic volatility models, the time-varying conditional variance in GARCH models is a deterministic function of past observations and past variances.

The first is the autoregressive model with GARCH (1,1) errors. In this setting, the conditional variance σ_t^2 is a linear function of the squared past shock u_{t-1}^2 and the past variance σ_{t-1}^2:

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_t^2) \tag{10}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \gamma \sigma_{t-1}^2, \tag{11}$$

where $\alpha_0 > 0, \alpha_1 \geq 0$ and $\gamma \geq 0$ ensure a positive conditional variance and $\sigma_0^2 = u_0 = 0$ is assumed for convenience. The model is referred to as the AR-GARCH model.

The second model combines an autoregressive-moving-average model with GARCH(1,1) errors:

$$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t, \tag{12}$$

$$\xi_t = u_t + \psi u_{t-1}, \qquad u_t \sim \mathcal{N}(0, \sigma_t^2). \tag{13}$$

The condition for invertibility is imposed, i.e., $|\psi| < 1$. The variance σ_t^2 follows the same GARCH process as above. This version of the GARCH model is referred to as ARMA-GARCH model.

4 Empirical forecasting results

The objective of the paper is to forecast the level of the real price of oil rather than the log oil price since policymakers are more concerned about the real price (Hamilton,

Table 1: List of models

Model	Equation	Variance	
Constant volatility models			
RW	$y_t = y_{y-1} + u_t,$	$u_t \sim \mathcal{N}(0, \sigma^2)$	(1)
AR	$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t,$	$u_t \sim \mathcal{N}(0, \sigma^2)$	(2)
ARMA	$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t,$		(3)
	$\xi_t = u_t + u_{t-1},$	$u_t \sim \mathcal{N}(0, \sigma^2)$	(4)
Time-varying volatility models			
(1) Stochastic volatility models			
AR-SV	$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t,$	$u_t \sim \mathcal{N}(0, e^{h_t})$	(5)
	$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h,$	$\varepsilon_t^h \sim \mathcal{N}(0, \sigma_h^2)$	(6)
AR-SVM	$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + \alpha e^{h_t} + u_t,$	$u_t \sim \mathcal{N}(0, e^{h_t})$	(7)
	$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h,$	$\varepsilon_t^h \sim \mathcal{N}(0, \sigma_h^2)$	(8)
ARMA-SV	$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t,$		(9)
	$\xi_t = u_t + \psi u_{t-1},$	$u_t \sim \mathcal{N}(0, e^{h_t})$	(10)
	$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h,$	$\varepsilon_t^h \sim \mathcal{N}(0, \sigma_h^2)$	(11)
(2) GARCH models			
AR-GARCH	$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t,$	$u_t \sim \mathcal{N}(0, \sigma_t^2)$	(12)
	$\sigma_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \gamma \sigma_{t-1}^2,$		(13)
ARMA-GARCH	$y_t = \beta_0 + \beta_1 y_{t-1} + \xi_t,$		(14)
	$\xi_t = u_t + \psi u_{t-1},$	$u_t \sim \mathcal{N}(0, \sigma_t^2)$	(15)
	$\sigma_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \gamma \sigma_{t-1}^2,$		(16)

1996; Alquist, Kilian and Vigfusson, 2011). Following Baumeister and Kilian (2012), the regression models in the paper are specified in logs, and the forecasts are transformed back to the level form. The forecast horizons are 1, 3, 6, 9 and 12 months. The paper broadly sets similar priors for both the constant and time-varying volatility models. Specifically, it chooses the same prior for common parameters. All priors are proper but relatively non-informative. The details of the Bayesian estimation methods are included in the appendix. The following subsections introduce the forecast evaluation metrics and show the forecasting results for the real-time and for the ex-post revised data.

4.1 Forecast evaluation metrics

A recursive out of sample forecasting scheme is used to evaluate the performance of the models listed in Table 1 for forecasting both the real-time and ex-post revised oil price at various horizons. Both point and density k-step-ahead iterated forecasts, with $k =$ 1, 3, 6, 9, 12, are computed to measure the forecast performance. The accuracy of the point forecast is measured by the mean square forecasting errors (MSFE), while the accuracy of density forecast is measured by the sum of the predictive likelihood. It is worth mentioning that the predictive likelihood motivated and described by Geweke and Amisano (2010) and Geweke and Amisano (2011) is commonly viewed as the broadest measure of forecasting.

Given data up to time t, which is denoted as $\mathbf{y}_{1:t}$, the MCMC sampler described in the appendix is implemented to obtain the posterior draws. Then the predictive mean $\mathrm{E}(y_{t+k}|\mathbf{y}_{1:t})$ is computed as the point forecast and the predictive density $p(y_{t+k}|\mathbf{y}_{1:t})$ as the density forecast. In the next step, the forecast moves one period ahead and the exercise is repeated in a rolling procedure, going forward one month by adding one more observation to the sample used for estimation. These forecasts are then evaluated for $t = t_0, ..., T - k$ where t_0 is 1993.12 for the real-time data and 1994.1 for the ex-post revised data.

According to Geweke and Amisano (2010) and Chan (2013), in practice, neither the predictive mean $\mathrm{E}(y_{t+k}|\mathbf{y}_{1:t})$ nor the predictive density $p(y_{t+k}|\mathbf{y}_{1:t})$ can be analytically computed. Instead, predictive simulation is used to obtain them. To be more precise, for each MCMC iteration, given the model parameters and states (up to time t), future states from time $t+1$ until $t+k$ can be simulated using the relevant transition equations. Meanwhile, future errors $u_s \sim \mathcal{N}(0, \sigma_y^2)$, $u_s \sim \mathcal{N}(0, e^{h_s})$ for $s = t+1, ..., t+k-1$ can be simulated. As mentioned earlier, the regression models in the paper are specified in logs and the forecasts are exponentiated. Therefore, given these draws, y_{t+k} follows the log normal distribution instead of the normal distribution and the point and density forecasting for y_{t+k} can be easily calculated at each MCMC iteration. These forecasts are then averaged over all the posterior draws to produce estimates for the predictive mean $\mathrm{E}(y_{t+k}|\mathbf{y}_{1:t})$ and the predictive density $p(y_{t+k}|\mathbf{y}_{1:t})$. The procedure then moves forward

to use data at t+1 and the process is repeated recursively to obtain $\mathrm{E}(y_{t+1+k}|\mathbf{y}_{1:t+1})$ and $p(y_{t+1+k}|\mathbf{y}_{1:t+1})$, and so forth.

To measure the accuracy of the point forecasts, the root mean squared forecast error (MSFE) is used, and is defined as

$$MSFE = \frac{\sum_{t=t_0}^{T-k} (y_{t+k}^o - \mathrm{E}(y_{t+k}|\mathbf{y}_{1:t}))^2}{T-k-t_0+1},$$

where y_{t+k}^o denotes the observed outcome of y_{t+k} that is known at time $t+k$. For this metric, a smaller value indicates better forecast performance.

The metric used to evaluate the density forecasts $p(y_{t+k}|\mathbf{y}_{1:t})$ is the log predictive likelihood $p(y_{t+k}^o|\mathbf{y}_{1:t})$. This is the predictive density for y_{t+k} formed at time t using the data from period $1, ..., t$, evaluated at the realization y_{t+k}^o. The value of the predictive likelihood will be large if the actual observation y_{t+k}^o is likely under the density forecast. The sum of log predictive likelihoods is used to evaluate the density forecasts:

$$\sum_{t=t_0}^{T-k} logp(y_{t+k} = y_{t+k}^o|\mathbf{y}_{1:t}).$$

For this metric, a larger value indicates better forecast performance.

4.2 Real-time forecasting results

Tables 2 and 3 summarize the results of the real-time forecasts based on the MSFE for the two oil price series: the real U.S. refiners' acquisition cost for imports, and the real WTI price. The forecast horizons are 1 month, 3 months, 6 months, 9 months and 12 months. In order to facilitate the comparison between different models, the results are reported in terms of the relative MSFE, i.e., the ratio between the MSFE of a specific model and the MSFE of the benchmark model. If the relative MSFE is lower than unity, the forecasts of that given model will be on average more accurate than those generated by the benchmark.

Overall, the ARMA models with time-varying volatility consistently outperform the RW and AR models at all forecast horizons for both types of the real oil price data. For example, the MSFEs for ARMA-SV and ARMA-GARCH are less than 95% of the value for RW at all forecast horizons. However, the improvements by adding the MA components in the constant volatility model does not exist at the long horizons. The ARMA model only improves on the RW baseline at short horizons. For example, for

Table 2: Recursive MSFE ratio for forecasts of the monthly real U.S. refiners' acquisition cost for imports relative to the RW using real-time data

Model	Forecast horizon				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	1	1	1	1	1
AR	0.942	0.943	1.027	1.148	1.224
ARMA	0.945	0.910	0.935	0.993	1.047
Time-varying volatility models					
AR-SV	0.961	0.948	1.019	1.099	1.102
AR-SVM	0.959	0.948	1.004	1.057	1.048
ARMA-SV	0.923	0.912	0.928	0.932	0.927
AR-GARCH	0.923	0.905	0.961	1.020	1.019
ARMA-GARCH	0.949	0.938	0.940	0.927	0.913

Table 3: Recursive MSFE ratio for forecasts of the monthly real WTI price relative to the RW using real-time data

Model	Forecast horizon				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	1	1	1	1	1
AR	0.924	0.966	1.013	1.089	1.153
ARMA	0.927	0.962	0.968	1.003	1.040
Time-varying volatility models					
AR-SV	0.919	0.968	1.032	1.082	1.088
ARSV-M	0.934	0.993	1.060	1.037	1.038
ARMA-SV	0.897	0.936	0.947	0.953	0.960
AR-GARCH	0.944	0.998	1.044	1.093	1.101
ARMA-GARCH	0.893	0.932	0.938	0.934	0.928

the real U.S. refiners' acquisition cost for imports, the ratio of the ARMA MSFE to the RW MSFE is 0.961 at the 1-month-ahead horizon, while the corresponding ratio is 1.047 at the 12-months-ahead horizon. Comparing the ARMA models with the RW and AR models shows that the MA component typically yields consistent improvements in forecast accuracy.

The results in Tables 2 and 3 also indicate that all AR models with and without time-varying volatility produce lower MSFEs than the RW forecast at short horizons (1-month-ahead and 3-month-ahead) for both types of the real oil price. The magnitude of the forecast accuracy gains is typically somewhat small, and the AR models reduce the MSFE by less than 10% at the short horizons. However, none of the AR specifications, even with time-varying volatility, yields any consistent advantages over the RW baseline. For longer horizon forecasts ($k \geq 6$), forecasts from the AR models are slightly less accurate than the RW model forecast.

Among the AR models, it is of interest to note that the specifications with time-varying volatility perform slightly better than the corresponding models with constant volatility at the longer-horizon forecast periods. These results show that including time-varying volatility in conditional variances yields gains in longer-horizon period forecast accuracy. However, including stochastic volatility in mean does not have much effect on the MSFE-based point accuracy. Specifically, compared to the AR-SV models, the extension improves forecast accuracy by a small amount in the case of real U.S. refiners' acquisition cost for imports but slightly reduces it at the shorter horizons for the real WTI price. In terms of the ARMA models, the model with stochastic volatility yields point forecasts that are in general more accurate than forecasts from the counterpart with constant variance.

As mentioned earlier, there are two classes of time-varying volatility models, namely GARCH models and SV models. The paper also investigates the performance of these two types of models in terms of the real-time oil prices point forecast. Comparing the AR-SV with the AR-GARCH, and the ARMA-SV with the ARMA-GARCH show that the SV models are more accurate in some cases and reduce accuracy in others. For example, the forecast performance of the AR-SV specification for the real WTI price is better than the forecast performance of the AR-GARCH model, while the performance of the AR-SV models in the real U.S. refiners' acquisition cost for imports is worse than their counterpart. However, compared to the forecasting performance of GARCH models, the improvements and reductions of forecasting performance using SV models are not considerable.

The second method to assess the accuracy of the models is to look at the entire predictive density. The results of the density forecasts for the two types of real-time oil price data are presented in Tables 4 and 5 respectively. To simplify the process of comparison, the sum of the log predictive likelihoods of a given model are computed, from which the corresponding RW baseline value is subtracted. Therefore, entries greater than zero indicate that forecasts from the indicated model are more accurate than forecasts from the associated baseline model.

Overall, the results in Tables 4 and 5 indicate that the specifications with constant volatility are always worse than the benchmark model for density forecasting at all horizons for both types of real-time oil price. Specifically, for both the ARMA and AR

models with constant volatility, all the relative sum of log predictive likelihood values are negative and far from zero. Moreover, the models that allow for time-varying variance yield density forecasts that, in general, significantly improve the accuracy of the density forecast relative to the RW baseline models and the models with constant volatility. The conclusion is that including time-varying volatility in AR models and ARMA models yields sizeable gains in density accuracy as measured by the sum of the log predictive likelihood.

Table 4: Sum of log predictive likelihood for forecasts of the monthly real U.S. refiners' acquisition cost for imports relative to the RW using real-time data

Model	**Forecast horizon**				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	0	0	0	0	0
AR	-156.81	-162.68	-149.36	-172.48	-173.48
ARMA	-148.05	-140.97	-120.48	-147.99	-147.65
Time-varying volatility models					
AR-SV	6.11	38.99	48.17	14.61	2.49
AR-SVM	15.32	43.36	51.50	17.06	1.88
ARMA-SV	23.21	47.84	58.79	31.05	17.25
AR-GARCH	8.69	33.96	37.89	1.69	-14.21
ARMA-GARCH	-14.12	30.96	43.93	17.76	6.83

Among the models with time-varying volatility, only the ARMA-SV model offers consistent advantages over the RW baseline model at all horizons for the two real-time oil price series. The ARMA-GARCH is the other specification that always dominates the baseline in the real WTI price forecasts. To further investigate the relevance of the MA component, this paper compares the AR model with the ARMA, the AR-SV with the ARMA-SV, the AR-GARCH with the ARMA-GARCH. All classes of models, models with the MA component show considerable improvement in the density forecasting accuracy. For example, the ARMA-SV model dominates the AR-SV model in forecasting the real WTI price at all horizons with the sum of the log predictive likelihoods increasing by factors as much as 10.81 at horizon 1, 19.55 at horizon 3, 35.91 at horizon 6, 14.52 at horizon 9 and 4.38 at horizon 12.

The AR-SV and ARSV-M specifications dominate the RW, with the exception of the real WTI price at longer horizons. Compared to the AR-SV model, the ARSV-M model improves the density forecast accuracy in some cases particularly for the real U.S. refiners' acquisition price at shorter horizons but reduces the accuracy in others, including for the real WTI price. The comparison shows that accounting for stochastic volatility in

Table 5: Sum of log predictive likelihood for forecasts of the monthly real WTI price relative to the RW using real-time data

Model	Forecast horizon				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	0	0	0	0	0
AR	-158.75	-157.00	-133.68	-148.24	-150.34
ARMA	-149.95	-153.21	-128.63	-145.97	-144.32
Time-varying volatility models					
AR-SV	5.21	15.16	31.42	5.45	-5.02
ARSV-M	-0.69	10.75	25.90	3.15	-10.55
ARMA-SV	10.81	19.55	34.91	14.52	4.38
AR-GARCH	5.82	11.60	28.26	3.35	-9.27
ARMA-GARCH	11.34	16.79	32.77	15.51	6.29

mean is not important for real-time oil prices forecasting. In addition, the SV specifications are usually more accurate than the corresponding GARCH models, except in a few cases (For example, for WTI price forecasting, the ARMA-GARCH model performances slightly better than the ARMA-SV model for $k = 9, 12$).

4.3 Forecasting results for ex-post revised data

Many existing studies used the ex-post revised data to forecast the oil prices. For convenience of comparison with other studies, this paper also includes the forecasting performance of models using the ex-post revised data. The results in Tables 6 and 7 are the counterparts of the real-time analysis in Tables 2 and 3. Meanwhile, the results in Tables 8 and 9 are the counterparts of the real-time analysis in Tables 4 and 5. All real-time data is replaced by the ex-post revised data. The evaluation window is from 1994.1 to 2014.12. According to Baumeister and Kilian (2012), virtually all data are revised half a year after the first publication and the data for the last few months of the available sample are excluded for forecasting the oil prices. Therefore, the paper treats the data up to 2014.12 in the 2015.4 vintage as the proxy for ex-post revised data.

The point forecasting results are broadly similar to those for the real-time data. The best two models are the ARMA-GARCH and the ARMA-SV models, and these models are the only two that dominate the RW baseline at all forecast horizons for both oil price series. In addition to the best forecasting performance using the real-time data,

Table 6: Recursive MSFE ratio for forecasts of the monthly real U.S refiners' acquisition cost for imports relative to the RW using ex-post revised data

Model	**Forecast horizon**				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	1	1	1	1	1
AR	0.796	0.930	1.047	1.166	1.257
ARMA	0.826	0.929	0.965	1.011	1.059
Time-varying volatility models					
AR-SV	0.792	0.921	1.031	1.111	1.124
ARSV-M	0.784	0.888	0.999	1.055	1.026
ARMA-SV	0.795	0.903	0.941	0.952	0.946
AR-GARCH	0.773	0.876	0.974	1.039	1.043
ARMA-GARCH	0.798	0.927	0.941	0.934	0.910

Table 7: Recursive MSFE ratio for forecasts of the monthly real WTI relative to the RW using ex-post revised data

Model	**Forecast horizon**				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	1	1	1	1	1
AR	0.933	0.971	1.029	1.104	1.170
ARMA	0.918	0.951	0.960	1.001	1.049
Time-varying volatility models					
AR-SV	0.927	0.974	1.040	1.091	1.099
ARSVM	0.936	0.977	1.018	1.101	1.155
ARMA-SV	0.902	0.939	0.949	0.957	0.967
AR-GARCH	0.951	0.999	1.051	1.100	1.110
ARMA-GARCH	0.898	0.931	0.937	0.935	0.933

the ARMA model with time-varying volatility improves forecasting accuracy for the ex-post revised data as well. In addition, the AR models with and without volatility are more accurate than the RW baseline at the shorter horizons only. However, unlike the case for forecasting using the real-time data, the results for forecasting using the ex-

post revised data provide more evidence that including stochastic volatility consistently improves, often dramatically, the accuracy of point forecasts relative to models with constant volatility at all horizons. Specifically, the ARMA-SV model and the AR-SV model usually dominate their counterparts with constant variance. When it comes to the GARCH class of time-varying volatility models, these are usually more accurate than the specifications with constant volatility, except for the forecasts of the real WTI at shorter forecasting horizons.

Among the models with time-varying volatility, the AR-SVM model yields consistent improvement relative to the AR-SV model in forecasting the monthly real U.S . refiners' acquisition cost for imports. Hoverer, the AR-SVM model is always worse than AR-SV in monthly real WTI forecasting. In addition, for the point forecasting of the ex-post revised data, SV models is better than the GARCH models in some cases and is worse in some other cases. For example, AR-SV always dominates AR-GARCH in forecasting monthly real WTI while AR-SV is dominated by AR-GARCH in forecasting monthly real U.S. refiners' acquisition cost.

For the case of forecasting using the ex-post data, the results for the density forecasting are broadly similar to the results for the case using the real-time data. The AR and the ARMA models with constant volatility perform considerably worse than the baseline model for the two oil price series. When it comes to the time-varying volatility models compared to the RW model, the models with time-varying volatility improve the density accuracy in most cases. Among the models with time-varying volatility, the ARMA-SV model is still the best model in terms of density forecasting for the ex-post data series. The second best model is the ARMA-GARCH in most cases. The results also indicate that in general the forecast performance of SV models is better than the forecast performance of GARCH models.

5 Conclusion

The previous literature generates real-time forecasts for the real price of oil using constant volatility models, but does not allow for volatility clustering, which is a prominent feature in oil data. This paper fills the gap by comparing the forecasting performance of a variety of models with constant and time-varying volatility, including the AR and ARMA specifications for real oil prices. The set of models includes those based on constant volatility, AR stochastic volatility, ARMA stochastic volatility, stochastic volatility in mean with constant parameters, standard GARCH and GARCH with MA innovations. Specifically, this paper has highlighted the importance of MA components and stochastic volatility in mean for forecasting the oil price, which translates into better forecast performance in some domestic macroeconomics aggregates as well (Chan, 2013; Chan, 2017). While many studies measure the forecasting performance of the oil price on

Table 8: Sum of log predictive likelihood for forecasts of the monthly real U.S refiners' acquisition cost for imports relative to the RW using ex-post revised data

Model	**Forecast horizon**				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	0	0	0	0	0
AR	-160.69	-166.22	-146.66	-166.90	-172.16
ARMA	-152.20	-144.36	-117.60	-141.54	-145.35
Time-varying volatility models					
AR-SV	20.37	36.72	49.75	17.39	1.20
AR-SVM	26.56	40.23	52.13	19.63	1.15
ARMA-SV	31.02	43.09	59.36	34.63	16.56
AR-GARCH	21.18	34.28	41.21	6.94	-12.97
ARMA-GARCH	12.41	35.14	47.88	24.94	10.46

Table 9: Sum of log predictive likelihood for forecasts of the monthly real WTI price relative to the RW using ex-post revised data

Model	**Forecast horizon**				
	$k=1$	$k=3$	$k=6$	$k=9$	$k=12$
Constant volatility models					
RW	0	0	0	0	0
AR	-161.98	-162.00	-140.60	-152.50	-153.33
ARMA	-152.91	-158.38	-135.19	-150.35	-147.62
Time-varying volatility models					
AR-SV	3.22	16.15	29.38	5.08	-4.75
AR-SVM	-1.02	9.50	21.47	0.62	-14.24
ARMA-SV	8.95	19.83	33.20	14.58	4.27
AR-GARCH	3.69	12.82	27.25	3.11	-9.36
ARMA-GARCH	7.30	18.94	32.84	15.75	5.98

the basis of the accuracy of point forecasts, this paper has also compared and evaluated the Bayesian predictive distribution from alternative models to compare the forecasting performance. Additionally, the importance of real-time forecasting is well recognized

in the literature (Croushore 2011, Baumeister and Kilian 2012). In recent years, there have been some studies generating real-time forecasts for the real price of oil, which is widely considered one of the key global macroeconomic indicators. However, there is no readily available real-time dataset for the variables required to forecast the real price of oil. This paper constructed a real-time dataset for the real oil price using backcasting by following the construction process that was introduced by Baumeister and Kilian (2012) and updated the real-time dataset for the real oil price until 2015. The paper examined both the real-time and ex-post revised oil data including the real U.S. refiners' acquisition cost for imports and real WTI price.

The results indicate that for both real-time and ex-post data, models with time-varying volatility dominate their counterparts with alternative volatility specifications in terms of point forecasting to a smaller degree in a longer horizon period and density forecasting to a large degree at all horizons. Among the models with time-varying volatility, the SV models perform better than the counterpart GARCH models, particularly in density forecasting. In addition, the model with stochastic volatility in mean does not show considerable improvement in the point and density forecasting accuracy compared to the standard SV model. Overall, the stochastic volatility models with moving average innovations are the best models for both real-time data and ex-post revised oil prices data point and density forecasting.

Appendix A

This section outlines the methods for estimating SV models including AR-SV, AR-SVM, ARMA-SV and the GARCH models including the AR-GARCH model, ARMA-GARCH model, which are introduced in Subsection 3.2.

Stochastic volatility models

This subsection firstly restates the AR-SV model. This model is an autoregressive model with SV errors. The conditional variance of y_t is $\text{Var}(y_t|h_t) = e^{h_t}$, and the state h_t is often called the log-volatility.

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t, \qquad u_t \ \sim \mathcal{N}(0, e^{h_t}),$$

$$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \ \sim \mathcal{N}(0, \sigma_h^2),$$

where $|\phi_h| < 1$ and the states are initialized with $h_1 \sim \mathcal{N}(\mu_h, \sigma_h^2/(1-\phi_h))$. The stochastic volatility model is an example of non-linear state space models in which the measurement equation is non-linear in the state. The joint conditional density of the state $\mathbf{h} = (h_1, h_2, ...h_t)$ given the model parameters and the data is high-dimensional and non-standard, while the the conditional density of the states of a linear Gaussian state space model is Gaussian. Chan and Hsiao (2014) use the auxiliary mixture sampler to estimate the stochastic volatility model. The nonlinear stochastic volatility model is approximated using a mixture of linear Gaussian models, where the latter models is standard. More specifically, given the mixture component indicator $\mathbf{s} = (s_1, ..., s_T)'$, the Gaussian approximation of the conditional posterior distribution $p(\mathbf{h}|\mathbf{y}, \beta, \mathbf{s}, \mu_h, \phi_h, \sigma_h^2)$ can be obtained. Bayesian analysis can be performed using a sample from the joint posterior distribution $p(\mathbf{h}, \mathbf{s}, \mu_h, \phi_h, \sigma_h^2, \beta|\mathbf{y})$. Posterior draws can be obtained via a Gibb sampler or M-H process that cycles through the conditional posterior distributions of the parameters. Refers to Chan and Hsiao (2014) for the estimation details.

The second SV specification is AR-SVM model, where the volatility appears in both the conditional mean and the conditional variance equations.

$$y_t = \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + \alpha e^{h_t} + u_t, \qquad u_t \ \sim \mathcal{N}(0, e^{h_t}),$$

$$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \ \sim \mathcal{N}(0, \sigma_h^2),$$

where $|\phi_h| < 1$ and the states are initialized with $h_1 \sim \mathcal{N}(\mu_h, \sigma_h^2/(1-\phi_h))$. For the stochastic volatility models, an important step is to jointly sample the log-volatility

$\mathbf{h}$ and the conditional density $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$ needs to be sampled. Unlike the auxiliary mixture sampler used in AR-SV model, for the ARMA-SV model this paper uses the accept-reject Metropolis-Hasting algorithm described in Chan (2017) to sample $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$.

One of the key ingredients to implement the accept-reject Metropolis-Hasting algorithm is finding an appropriate proposal density which can approximate the target density $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$ well. The basic idea is to approximate the target density using a Gaussian density. The target density can be represented as $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2) \propto p(\mathbf{y}|, \beta, \alpha, \mathbf{h})p(\mathbf{h}|\mu_h, \phi_h, \sigma_h^2)$, and the latter density $p(\mathbf{h}|\mu_h, \phi_h, \sigma_h^2)$ is Gaussian.

The log likelihood $p(\mathbf{y}|, \beta, \alpha, \mathbf{h})$ is not Gaussian in $\mathbf{h}$ and needs to be approximated by a Gaussian density in $\mathbf{h}$. To approximate $p(\mathbf{y}|\beta, \alpha, \mathbf{h})$ by a Gaussian density in $\mathbf{h}$, this paper expands $\log p(\mathbf{y}|\beta, \alpha, \mathbf{h}) = \sum_{t=1}^{T} logp(y_t|\beta, \alpha, h_t)$ around the point $\widetilde{h}$, which is the mode of $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$, by a second-order Taylor expansion:

$$logp(\mathbf{y}|\beta, \alpha, \mathbf{h}) \approx logp(\mathbf{y}|\beta, \alpha, \widetilde{\mathbf{h}}) + (\mathbf{h} - \widetilde{\mathbf{h}})'\mathbf{f} - \frac{1}{2}(\mathbf{h} - \widetilde{\mathbf{h}})'\mathbf{G}(\mathbf{h} - \widetilde{\mathbf{h}}) = -\frac{1}{2}(\mathbf{h}'\mathbf{G}\mathbf{h} - 2\mathbf{h}'(\mathbf{f} + \mathbf{G}\widetilde{\mathbf{h}})) + c_2,$$

where c_2 is a constant and independent of $\mathbf{h}$, $\mathbf{f} = (f_1, f_2, ...f_T)'$ and $G = \text{diag}\ (G_1, ...G_T)$ with

$$f_t = \frac{\partial}{\partial h_t} logp(y_t|\beta, \alpha, h_t)|_{h_t = \widetilde{h_t},}$$
$$G_t = \frac{\partial^2}{\partial h_t^2} logp(y_t|\beta, \alpha, h_t)|_{h_t = \widetilde{h_t},}$$

where $\mathbf{G}$ is the negative Hessian of the log-density evaluated at $\widetilde{\mathbf{h}}$. Finally combining the $p(\mathbf{h}|\mu_h, \phi_h, \sigma_h^2)$ and the Gaussian approximation of $p(\mathbf{y}|\beta, \alpha, \mathbf{h})$, the Gaussian density $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$ can be obtained as the proposal density in the accept-reject Metropolis-Hasting algorithm. Refers to Chan (2017) for the estimation details.

The ARMA-SV model allows the errors in the measurement equation for following an MA(1)process

$$\begin{aligned}
y_t &= \beta_0 + \beta_1 y_{t-1} + \xi_t, \\
\xi_t &= u_t + \psi u_{t-1}, \qquad u_t \ \sim \mathcal{N}(0, e^{h_t}), \\
h_t &= \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \qquad \varepsilon_t^h \ \sim \mathcal{N}(0, \sigma_h^2),
\end{aligned}$$

where $u_0 = 0$, $|\psi| < 1$, $|\phi_h| < 1$ and the states are initialized with $h_1 \sim \mathcal{N}(\mu_h, \sigma_h^2/(1 - \phi_h))$.

This paper defines $\widetilde{\mathbf{y}} = \mathbf{H}_\varphi(\mathbf{y} - \mathbf{x}\beta)$, where $\mathbf{y} = [y_1, y_2, ..., y_T]'$, $x_t = [1, y_{t-1}]$, $\mathbf{x} = [x_1, x_2, ..., x_T]'$, $\beta = [\beta_0, \beta_1]'$ and $\mathbf{H}_\varphi$ is a $T \times T$ lower triangular matrix with ones on the main diagonal, φ on first lower diagonal

$$H_\varphi = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ \varphi & 1 & 0 & \dots & 0 & 0 \\ 0 & \varphi & 1 & \dots & 0 & 0 \\ \vdots & & \ddots & \ddots & & \vdots \\ 0 & 0 & 0 & \dots & \varphi & 1 \end{bmatrix}.$$

Then $(\widetilde{\mathbf{y}}|\beta, \varphi, \mathbf{h})$ follows the multivariate normal distribution with $(\widetilde{\mathbf{y}}|\beta, \varphi, \mathbf{h}) \sim \mathcal{N}(\mathbf{0}, S_\mathbf{y})$, where $S_\mathbf{y} = diag(e^{h_1}, ...e^{h_T})$. Hence, the auxiliary mixture sampler discussed in ARSV model can be directly applied to sample from $p(\mathbf{h}|\mathbf{y}, \beta, \alpha, \mu_h, \phi_h, \sigma_h^2)$. Refers to Chan (2013) for the estimation details.

GARCH models

Firstly, this subsection restates the AR-GARCH model. This model is an autoregressive model with GARCH (1,1) errors. The conditional variance σ_t^2 is a linear function of the squared past shock and the past variance,

$$\begin{aligned} y_t &= \beta_0 + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + u_t, \qquad u_t \ \sim \mathcal{N}(0, \sigma_t^2), \\ \sigma_t^2 &= \alpha_0 + \alpha_1 u_{t-1}^2 + \gamma \sigma_{t-1}^2, \end{aligned}$$

where $\alpha_0 > 0, \alpha_1 \geq 0$ and $\gamma \geq 0$ ensure a positive conditional variance and $\sigma_0^2 = u_0 = 0$ for convenience.

For the GARCH models, the paper samples the joint posterior distribution of the parameters using the approach suggested by Nakatsuma (1999). For example, for the AR-GARCH model, the recursive nature of the variance equation does not allow for conjugacy between the likelihood function. Therefore, the paper relies on the Metropolis-Hasting (M-H) algorithm to draw sample from the joint posterior distribution.

In order to express the likelihood, define the vector $\mathbf{y} = (y_1, ..., y_T)'$, $x_t = (y_{t-1}, ..., y_{t-12})$, $\mathbf{x} = (x_1, ..., x_T)'$, $u = (u_1, ..., u_T)'$. The likelihood function in matrix notation is:

$$p(\mathbf{y}|\beta, \alpha, \gamma) \propto |\Sigma|^{-\frac{1}{2}} \exp[-\frac{1}{2}(\mathbf{y} - \mathbf{x}\beta)'\Sigma^{-1}(\mathbf{y} - \mathbf{x}\beta)], \tag{14}$$

with $\Sigma = diag(h_1, ...h_T)$. Since the full conditional density is unknown analytically, parameters β , α and γ are sampled from three proposal densities. The proposal density

to sample the $p \times 1$ vector β is obtained by combining the likelihood function (14) and the prior density $p(\beta) \sim \mathcal{N}(u_\beta, \Sigma_\beta)$ by the use Bayes update:

$$p_\beta(\beta|\tilde{\beta}, \alpha, \gamma, \mathbf{y}) \sim \mathcal{N}(\widehat{u_\beta}, \widehat{\Sigma_\beta}),$$

with

$$\widehat{\Sigma_\beta}^{-1} = (\mathbf{x}\tilde{\Sigma}^{-1}\mathbf{x} + \Sigma_\beta^{-1}),$$
$$\widehat{u_\beta} = \widehat{\Sigma_\beta}(\mathbf{x}'\tilde{\Sigma}^{-1}\mathbf{y} + \Sigma_\beta^{-1}u_\beta),$$

where the $T \times T$ diagonal matrix $\tilde{\Sigma} = diag(h_1(\tilde{\beta}, \alpha, \gamma), ..., h_T(\tilde{\beta}, \alpha, \gamma))$ and $\tilde{\beta}$ is the previous draw of β in the M-H sampler. A candidate β^* is sampled from this proposal density and accepted with probability:

$$\min\left(\frac{p(\beta^*|\alpha,\gamma,\mathbf{y})p_\beta(\tilde{\beta}|\beta^*,\alpha,\gamma,\mathbf{y})}{p(\tilde{\beta}|\alpha,\gamma,\mathbf{y})p_\beta(\beta^*|\tilde{\beta},\alpha,\gamma,\mathbf{y})}, 1\right).$$

In terms of parameters α and γ, these densities are obtained by noting that the GARCH(1,1) models can be written as an ARMA(1,1) model for y_t^2. The conditional variance can be expressed as

$$u_t^2 = \alpha_0 + (\alpha_1 + \gamma)u_{t-1}^2 - \gamma\omega_{t-1} + \omega_t,$$

by defining $\omega_t = u_t^2 - h_t$.

According to Nakatsuma (2000), the variable ω_t which is a martingale difference process with a conditional mean of zero and a conditional variance of $2h_t^2$, can be approximated by a variable z_t which is normally distributed with a mean zero and a variance of $2h_t^2$. Then an auxiliary model is generated:

$$z_t = u_t^2 - \alpha_0 - (\alpha_1 + \gamma)u_{t-1}^2 + \gamma z_{t-1}, \tag{15}$$

by noting that z_t is a function of parameters α, β, γ. By defining the $T \times 1$ vector $\mathbf{z} = (z_1, ..., z_T)'$ as well as the $T \times T$ matrix: $\Lambda = diag(2h_1^2, ..., 2h_T^2)$, it can approximate the likelihood function from the auxiliary model (15) as:

$$p(\mathbf{y}|\beta, \alpha, \gamma) \propto (det\Lambda)^{-\frac{1}{2}} \exp(-\frac{1}{2}\mathbf{z}'\Lambda^{-1}\mathbf{z}). \tag{16}$$

This likelihood function is used to construct the proposal density for parameter α and γ.

Defines $v_t = y_t^2$ for notational convenience. To generate vector α, Nakatsuma (2000) uses the recursive transformation, initially proposed by Chib and Greenbe (1994), to express the function z_t in (15) as a linear function of a 2×1 vector α. However, instead of recursive transformations, this paper rewrites the function (15) as the linear function of α

$$H\mathbf{z} = H\mathbf{v} - \bar{\delta}\alpha, \tag{17}$$

where $\mathbf{v} = (v_1, ..., v_T)'$, $\bar{\delta}$ is a $T \times 2$ matrix:

$$\bar{\delta} = \begin{bmatrix} 1 & 0 \\ 1 & v_1 \\ \vdots & \vdots \\ 1 & v_T \end{bmatrix},$$

and H is a $T \times T$ lower diagonal matrix with ones on the main diagonal

$$H = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ -\gamma & 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \\ 0 & 0 & 0 & \dots & -\gamma & 1 \end{bmatrix}.$$

By pre-multiplying (17) by H^{-1},

$$Z = \mathbf{v} - \delta\alpha, \tag{18}$$

where $\delta = H^{-1}\bar{\delta}$. Therefore the approximate likelihood function of parameter α is expressed as follows:

$$p(\mathbf{y}|\beta, \alpha, \gamma) \propto (det\Lambda)^{-\frac{1}{2}} \exp[-\frac{1}{2}(\mathbf{v} - \delta\alpha)'\Lambda^{-1}(\mathbf{v} - \delta\alpha)].$$

The proposal density to sample vector α is obtained by combining this likelihood function and the prior density $p(\alpha) \sim \mathcal{N}(u_\alpha, \Sigma_\alpha)$, by the usual Bayes update:

$$p_\alpha(\alpha|\tilde{\alpha}, \beta, \gamma, \mathbf{y}) \sim \mathcal{N}_2(\widehat{u_\alpha}, \widehat{\Sigma_\alpha}),$$

with:

$$\widehat{\Sigma_\alpha}^{-1} = \delta'\tilde{\Lambda}^{-1}\delta + \Sigma_\alpha^{-1},$$
$$\widehat{u_\alpha} = \widehat{\Sigma_\alpha}(\Sigma_\alpha u_\alpha + \delta'\tilde{\Lambda}\mathbf{v}),$$

where the $T \times T$ diagonal matrix $\tilde{\Lambda} = diag(2h_1^2(\tilde{\alpha}, \beta, \gamma), ..., 2h_T^2(\tilde{\alpha}, \beta, \gamma))$ and $\tilde{\alpha}$ is the previous draw of α in the M-H sampler. A candidate α^* is sampled from this proposal density and accepted with probability

$$\min\ (\frac{p(\alpha^*|\beta,\gamma,\mathbf{y})p_\alpha(\tilde{\alpha}|\alpha^*,\beta,\gamma,\mathbf{y})}{p(\tilde{\alpha}|\beta,\gamma,\mathbf{y})p_\alpha(\alpha^*|\tilde{\alpha},\beta,\gamma,\mathbf{y})}, 1).$$

As mentioned in Nakatsuma (2000), the function in (15) can be expressed as a linear function of parameter α but cannot be expressed as a linear function of γ. To solve this problem, $z_t(\gamma)$ is linearised by a first order Taylor expansion at point $\tilde{\gamma}$, which is the previous draw in the M-H sampler

$$z_t(\gamma) = z_t(\tilde{\gamma}) + \frac{\partial z_t}{\partial \gamma}|_{\gamma=\tilde{\gamma}} \times (\gamma - \tilde{\gamma}).$$

Furthermore, the paper defines the following: $\nu_t = z_t(\tilde{\gamma}) + \nabla_t\tilde{\gamma}$ and $\nabla_t = -\frac{\partial z_t}{\partial \gamma}|_{\gamma=\tilde{\gamma}}$. Therefore $z_t = \nu_t - \gamma\nabla_t$. By differentiating (15) with respect to γ, the ∇_t can be computed by the following recursion: $\nabla_t = y_{t-1}^2 - z_{t-1}(\tilde{\gamma}) + \nabla_{t-1}\tilde{\gamma}$ with $\nabla_0 = 0$. By defining $\boldsymbol{\nu} = (\nu_1, ..., \nu_T)'$ and $\boldsymbol{\nabla} = (\nabla_1, ...\nabla_T)'$, the term with exponential in (16) can be approximated by $\mathbf{z} = \boldsymbol{\nu} - \gamma\boldsymbol{\nabla}$. This generates the approximate likelihood function for parameter γ

$$p(\mathbf{y}|\beta,\alpha,\gamma) \propto (det\Lambda)^{-\frac{1}{2}} \exp[-\frac{1}{2}(\boldsymbol{\nu} - \gamma\boldsymbol{\nabla})'\Lambda^{-1}(\boldsymbol{\nu} - \gamma\boldsymbol{\nabla})]. \tag{19}$$

The proposal density to sample γ can be obtained by combining the prior density $p(\gamma) \sim \mathcal{N}(u_\gamma, \Sigma_\gamma)$ and the approximate likelihood:

$$\begin{aligned} p_\gamma &\sim \mathcal{N}(\widehat{u_\gamma}, \widehat{\Sigma_\gamma}), \\ \widehat{\Sigma_\gamma}^{-1} &= \boldsymbol{\nabla}'\tilde{\Lambda}^{-1}\boldsymbol{\nabla} + \Sigma_\gamma^{-1}, \\ \widehat{u_\gamma} &= \widehat{\Sigma_\gamma}(\boldsymbol{\nabla}'\tilde{\Lambda}^{-1}\boldsymbol{\nu} + \Sigma_\gamma^{-1}u_\gamma), \end{aligned}$$

where the $T \times T$ diagonal matrix $= diag(2h_1^2(\beta,\alpha,\tilde{\gamma}), ..., 2h_T^2(\beta,\alpha,\tilde{\gamma})$. A candidate γ^* is sampled from this proposal density and accepted with probability:

$$\min\ (\frac{p(\gamma^*|\alpha,\beta,\mathbf{y})p_\gamma(\tilde{\gamma}|\gamma^*,\alpha,\beta,\mathbf{y})}{p(\tilde{\gamma}|\alpha,\beta,\mathbf{y})p_\gamma(\gamma^*|\tilde{\gamma},\alpha,\beta,\mathbf{y})}, 1).$$

The ARMA-GARCH model is a model which combines an autoregressive-moving-average model with GARCH innovations

$$\begin{aligned} y_t &= \beta_0 + \beta_1 y_{t-1} + \xi_t, \\ \xi_t &= u_t + \psi u_{t-1}, \qquad u_t \sim \mathcal{N}(0, \sigma_t^2), \end{aligned}$$

under which the condition for invertibility is imposed, i.e., $|\psi| < 1$. The variance σ_t^2 follows the same GARCH process as above.

The basic sampler remains the same. For the additional parameter ψ, the generation of ψ is similar to γ by using the first-order Taylor expansion.

References

[1] Alquist, R., and Kilian, L. (2010). What do we learn from the price of crude oil futures?. *Journal of Applied Econometrics*, 25(4), 539-573.

[2] Alquist, R., Kilian, L., and Vigfusson, R. J. (2013). Forecasting the price of oil. *Handbook of Economic Forecasting*, 2, 427-507.

[3] Baumeister, C., and Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business and Economic Statistics*, 30(2), 326-336.

[4] Baumeister, C., and Kilian, L. (2014). What central bankers need to know about forecasting oil prices. *International Economic Review*, 55(3), 869-889.

[5] Baumeister, C., and Kilian, L. (2014). Real-time analysis of oil price risks using forecast scenarios. *IMF Economic Review*, 62(1), 119-145.

[6] Baumeister, C., Kilian, L., and Zhou, X. (2013). Are product spreads useful for forecasting? An empirical evaluation of the Verleger hypothesis. CEPR Discussion Paper No. DP9572.

[7] Baumeister, C., and Kilian, L. (2015). Forecasting the real price of oil in a changing world: a forecast combination approach. *Journal of Business and Economic Statistics*, 33(3), 338-351.

[8] Berument, H., Yalcin, Y., and Yildirim, J. (2009). The effect of inflation uncertainty on inflation: Stochastic volatility in mean model within a dynamic framework. *Economic Modelling*, 26(6), 1201-1207.

[9] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3), 307-327.

[10] Chan, J. C. C., and Jeliazkov, I. (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2), 101-120.

[11] Chan, J. C. C. (2013). Moving average stochastic volatility models with application to inflation forecast. *Journal of Econometrics*, 176(2), 162-172.

[12] Chan, J. C. C., and Hsiao, C. Y. L. (2013). Estimation of stochastic volatility models with heavy tails and serial dependence, In I. Jeliazkov and X.-S. Yang(Eds.), *Bayesian Inference in the Social Sciences*. John Wiley & Sons, Inc., Hoboken, NJ, USA.

[13] Chan, J. C. C. (2017). The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business and Economic Statistics*, 35(1), 17-28.

[14] Chernenko, S., Schwarz, K., and Wright, J. H. (2004). The information content of forward and futures prices: Market expectations and the price of risk. No 808, International Finance Discussion Papers, Board of Governors of the Federal Reserve System (U.S.).

[15] Chib, S., and Greenberg, E. (1994). Bayes inference in regression models with ARMA (p, q) errors. *Journal of Econometrics*, 64(1), 183-206.

[16] Clark, T. E., and Ravazzolo, F. (2015). Macroeconomic Forecasting Performance under Alternative Specifications of TimeVarying Volatility. *Journal of Applied Econometrics*, 30(4), 551-575.

[17] Croushore, D. (2011). Forecasting with real-time data vintages. *The Oxford Handbook of Economic Forecasting*, 247-267.

[18] D'Agostino, A., Gambetti, L., and Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, 28(1), 82-101.

[19] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987-1007

[20] Fair, R. C., and Shiller, R. J. (1990). Comparing information in forecasts from econometric models. *The American Economic Review*, 80(3), 375-89. 375-389.

[21] Geweke, J., and Amisano, G. (2010). Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting*, 26(2), 216-230.

[22] Geweke, J., and Amisano, G. (2011). Hierarchical Markov normal mixture models with applications to financial asset returns. *Journal of Applied Econometrics*, 26(1), 1-29.

[23] Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics*, 38(2), 215-220.

[24] Kang, S. H., Kang, S. M., and Yoon, S. M. (2009). Forecasting volatility of crude oil markets. *Energy Economics*, 31(1), 119-125.

[25] Knetsch, T. A. (2007). Forecasting the price of crude oil via convenience yield predictions. *Journal of Forecasting*, 26(7), 527-549.

[26] Koopman, S. J., and Hol Uspensky, E. (2002). The stochastic volatility in mean model: empirical evidence from international stock markets. *Journal of applied Econometrics*, 17(6), 667-689.

[27] Larsson, K., and Nossman, M. (2011). Jumps and stochastic volatility in oil prices: Time series evidence. *Energy Economics*, 33(3), 504-514.

[28] Morana, C. (2001). A semiparametric approach to short-term oil price forecasting. *Energy Economics*, 23(3), 325-338.

[29] Mason, C. F., and Wilmot, N. A. (2014). Jump processes in natural gas markets. *Energy Economics*, 46(S1), S69-S79.

[30] Nakatsuma, T. (2000). Bayesian analysis of ARMAGARCH models: A Markov chain sampling approach. *Journal of Econometrics*, 95(1), 57-69.

[31] Narayan, P. K., and Narayan, S. (2007). Modelling oil price volatility. *Energy Policy*, 35(12), 6549-6553.

[32] Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3), 821-852.

[33] Romer, C. D., and Romer, D. H. (2000). Federal Reserve information and the behavior of interest rates. *American Economic Review*, 90(3), 429-457.

[34] Sadorsky, P. (2005). Stochastic volatility forecasting and risk management. *Applied Financial Economics*, 15(2), 121-135.

[35] Sanders, D. R., Manfredo, M. R., and Boris, K. (2009). Evaluating information in multiple horizon forecasts: The DOE's energy price forecasts. *Energy economics*, 31(2), 189-196.

[36] Stark, T., and Croushore, D. (2002). Forecasting with a real-time data set for macroeconomists. *Journal of Macroeconomics*, 24(4), 507-531.

[37] Taylor, S. J. (1994). Modeling stochastic volatility: A review and comparative study. *Mathematical finance*, 4(2), 183-204.

[38] Vo, M. T. (2009). Regime-switching stochastic volatility: evidence from the crude oil market. *Energy Economics*, 31(5), 779-788.

[39] Ye, M., Zyren, J., and Shore, J. (2002). Forecasting crude oil spot price using OECD petroleum inventory levels. *International Advances in Economic Research*, 8(4), 324-333.

[40] Yu, L., Wang, S., and Lai, K. K. (2008). Forecasting crude oil price with an EMD-based neural network ensemble learning paradigm. *Energy Economics*, 30(5), 2623-2635.