



Predicting the volatility of major energy commodity prices: The dynamic persistence model[☆]

Jozef Baruník^{a,b,*}, Lukáš Vácha^b

^a Institute of Economic Studies, Charles University, Czech Republic

^b The Czech Academy of Sciences, Institute of Information Theory and Automation, Czech Republic

ARTICLE INFO

JEL classification:

C14
C18
C22
C50

Keywords:

Persistence heterogeneity
Wold decomposition
Local stationarity
Time-varying parameters

ABSTRACT

Time variation and persistence are crucial properties of volatility that are often studied separately in energy volatility forecasting models. Here, we propose a novel approach that allows shocks with heterogeneous persistence to vary smoothly over time, and thus model the two together. We argue that this is important because such dynamics arise naturally from the dynamic nature of shocks in energy commodities. We identify such dynamics from the data using localised regressions and build a model that significantly improves volatility forecasts. Such forecasting models, based on a rich persistence structure that varies smoothly over time, outperform state-of-the-art benchmark models and are particularly useful for forecasting over longer horizons.

1. Introduction

Energy commodities are of paramount interest to global economic prosperity because they are the most widely used source of energy.¹ They are constantly changing due to the constant influx of new technologies, environmental pressures and the geopolitical importance of controlling oil supplies. In recent years, large fluctuations in energy prices have been a major concern for both market participants and regulators, as price uncertainty has a significant impact on the economy (Elder and Serletis, 2010) and measuring and predicting the variability of energy prices is essential for pricing derivatives, asset allocation or risk management, but is also a key factor in understanding fluctuations in stock prices, growth rates, inflation, employment and exchange rates (Hamilton, 1983; Kilian, 2009; Kang and Yoon, 2013). While an increasing number of authors pay considerable attention to forecasting volatility (Haugom et al., 2014; Sévi, 2014; Lu et al., 2022; Zhang et al., 2019), all approaches are challenged by severe

limitations imposed by model assumptions. Determining the true data-generating process of volatility dynamics becomes a challenging and open question of great priority to market participants, financial analysts and policymakers.

In particular, it is well documented that energy volatility has exhibited a very high degree of temporal variation over the past decades, as both stable and uncertain periods have been driven by different shocks (Le et al., 2023). These can be associated with different states of an economy, as well as supply-side shocks, endogenous shocks driven by inventory and short-term price changes, or shocks due to the financialisation of commodities (and portfolio rebalancing). At the same time, an increasing number of authors have recently argued that a number of economic variables are driven by shocks that influence their future value with heterogeneous levels of persistence (Bandi and Tamoni, 2022). A possibly non-linear combination of transitory and persistent responses to shocks will produce volatility with heterogeneous persistence structures that remain hidden to the observer using

[☆] We are grateful to the editor, Lance Bachmeier, and two anonymous referees for their useful comments and suggestions, which have greatly improved the paper. We are indebted to Rainer von Sachs, Federico Severino, Roman Liesenfeld, Nikolas Hautsch, Christian Hafner, Wolfgang Härdle, Lubos Hanus for invaluable discussions and comments. We acknowledge insightful comments from numerous seminar presentations, such as: the Recent Advances in Econometrics (2023, Louvain), the 2021 and 2022 STAT of ML conference in Prague; the 15th International Conference on Computational and Financial Econometrics. The support from the Czech Science Foundation under the 19-28231X (EXPRO) project is gratefully acknowledged. For estimation of the quantities proposed, we provide a package `tvPersistence.jl` in JULIA. The package is available at <https://github.com/barunik/tvPersistence.jl>.

* Corresponding author at: Institute of Economic Studies, Charles University, Czech Republic.

E-mail address: barunik@fsv.cuni.cz (J. Baruník).

¹ To illustrate, more than 30% of US energy consumption used petroleum-based fuels and another 30% used natural gas. Most of this energy was used in transportation, industry and electricity generation. These figures are based on reports from the Lawrence Livermore National Laboratory: <https://www.llnl.gov>.

traditional methods. In turn, research on energy volatility is limited to models that aggregate one of these features separately.

Inferring such time-varying persistence from data on energy prices and various measures of uncertainty has crucial implications for policy making, modelling or forecasting. However, despite progress in exploring structural breaks (Wen et al., 2016; Aroui et al., 2012), regime switching models (Ma et al., 2017), popular heterogeneous autoregressive models (Zhang et al., 2019) or more complicated long memory structures (Wang and Wu, 2012; Wang et al., 2016; Ozdemir et al., 2013; Charfeddine, 2014; Herrera et al., 2018), which can exhibit large amounts of time persistence without being non-stationary (Baillie et al., 1996), there is still no clear consensus on how to explore such dynamic nature of the data. The inability to identify the dependence from the data alone leads to a tendency to rely on assumptions that are difficult, if not impossible, to validate. To better understand and forecast energy time series, we need an approach that can precisely locate the horizons and time periods in which the critical information occurs.

This paper proposes a novel representation for the non-stationary volatility of energy commodities that allows a researcher to identify and explore its rich time-varying heterogeneous persistence structures. We aim to identify localised persistence of volatility in major energy commodities that will be useful for modelling and forecasting purposes. Our work is closely related to the recent work of Baruník and Vácha (2023), who propose to localise the persistence structure in time series and open new avenues for modelling and forecasting. In this paper, we develop such a model and open new avenues for modelling and forecasting in the energy economics literature.

Different degrees of persistence, which also change over time, are natural in energy commodity data. External shocks due to changing geopolitical risks and economic conditions affecting energy markets have different duration and persistence, depending on the behaviour and sentiment of agents. The combination of shocks determines the persistence structure of energy commodity prices and, more importantly, their volatility. When analysing volatility time series using classical aggregate measures, we cannot see this complicated persistence structure with different shock durations (Ortu et al., 2020; Bandi et al., 2021; Baruník and Vácha, 2023). Wold decomposition, the cornerstone of time series modelling that underlies the vast majority of models used, assumes that volatility is driven by a linear combination of shocks. For example, if the process has little dependence, then each shock will have a very short duration and we call this behaviour a transient. On the other hand, if the process is very persistent, close to a random walk, then a shock will remain in the time series for a long time. In a real time series we usually observe a mixture of transient and persistent effects of shocks. However, it is difficult to identify the different effects of shocks because the Wold decomposition aggregates the entire persistence structure. To see the full persistence structure, Ortu et al. (2020) proposed the extended Wold decomposition, which allows to identify the effect of shocks at different horizons.

Our paper contributes to the literature exploring the persistence structure of volatility through time-varying changes in unconditional volatility, persistence and the occurrence of jumps (Granger and Ding, 1996; Granger and Hyung, 2004; Bollerslev and Engle, 1993; Baruník and Vácha, 2023). We argue that energy volatility has a rich time-varying persistence structure of shocks. To exploit these stochastic properties, we use a model that can capture the time variation of the persistence structure of volatility, namely the Time-Varying Extended Wold Decomposition (TV-EWD) model of Baruník and Vácha (2023). To identify the time variation of the stochastic properties of volatility time series, the model uses the locally stationary process proposed by Dahlhaus (1996). The localisation allows to obtain the localised extended Wold decomposition, which allows a precise identification of the time-varying persistence structure. In the empirical part, we show that the time-varying persistence structure reveals significant

changes in the way shocks propagate through the volatility of energy commodities.

In the empirical part, we argue that the TV-EWD model is a useful tool for understanding and forecasting the volatility of energy commodities. Specifically, we use highly liquid crude oil (CL), natural gas (NG) and regular gasoline (HU/RBOB) futures contracts, which account for more than 75% of the trading volume in the energy commodities market. Note that crude oil is the raw material used to produce heating oil, gasoline and other petroleum-based products. First, we identify the evolution of the time-varying persistence structure of volatility. Second, we build a model based on the identified structure and use it in a forecasting exercise. Our results indicate that forecasts based on the TV-EWD model are superior to other commonly used models, especially for longer horizon forecasts. As energy commodity prices and volatility are strongly related to the geopolitical situation, global economic conditions and agents' sentiment, it is natural to expect that the degree of persistence of shocks will vary over time. The empirical part discusses these persistence dynamics for events such as the annexation of Crimea in 2014, the Covid pandemic in 2020 and the aggression against Ukraine in February 2022. In contrast to these turbulent times, we also consider the calm period of 1990's when the volatility was driven by very different structure of shocks.

The rest of the paper is structured as follows. Section 2 proposes a model for the time-varying persistence of energy commodities based on a locally stationary process, Section 3 examines the time-varying persistence of volatility in oil-based markets, and Section 4 concludes.

2. Dynamically persistent volatility

Here we present a model that captures the time-varying persistence structure of volatility. We begin by modifying the key idea of classical time series, which associates any covariance stationary time series with a linear combination of its own past shocks and moving average components of finite order (Wold, 1938; Hamilton, 2020). While stationarity plays a key role in volatility analysis over decades due to the availability of natural linear Gaussian modelling frameworks, volatility turns out to be non-stationary in the longer run (Stărică and Granger, 2005). The state of the world, as well as the behaviour of agents, is highly dynamic and the assumption of time-invariant mechanisms generating energy volatility is unrealistic. A more general non-stationary process can be one that is locally close to a stationary process at any point in time, but whose properties (covariances, parameters, etc.) gradually change in a non-specific way over time. The localisation of the classical representation is essential to allow time variation that captures unstable behaviour. Finally, we introduce the decomposition of shocks in a locally stationary volatility into a heterogeneous degree of persistence and propose a forecasting model that captures both the time variation and the persistence of shocks.

2.1. Locally stationary volatility

The idea that volatility can only be stationary for a limited period of time and that this is still valid for estimation was introduced by Stărică and Granger (2005). Locally stationary processes that allow for slow variation of the stochastic properties of the process were formally introduced by Dahlhaus (1996).

Assume that a nonstationary volatility process v_t depends on a time-varying parameter model. Following Dahlhaus (1996), we replace v_t by a triangular array of observations $(v_{t,T}; t = 1, \dots, T)$ where T denotes the sample size. We interpret the process $v_{t,T}$ as a local stationary approximation around a fixed point t/T . As a consequence, the process can change its stochastic properties smoothly over time. An important feature of this local approximation is the possibility to represent a locally stationary process $v_{t,T}$ as a time-varying MA(∞):

$$v_{t,T} = \sum_{h=-\infty}^{+\infty} \alpha_{t,T}(h) \epsilon_{t-h}, \quad (1)$$

where the coefficients $\alpha_{i,T}(h)$ can be approximated under certain smoothness assumptions as $\alpha_{i,T}(h) \approx \alpha(t/T, h)$, see [Dahlhaus \(1996\)](#). The innovations ϵ_t are independent random variables with $\mathbb{E}\epsilon_t = 0$, $\mathbb{E}|\epsilon_t| < \infty$ and $\mathbb{E}\epsilon_s \epsilon_t = 0$ for $s \neq t$. The Wold decomposition of a locally stationary process in Eq. (1) is a linear combination of uncorrelated innovations with time-varying impulse response functions $\alpha(t, T)$ at a fixed point t/T .

2.2. Time-Varying Extended Wold Decomposition (TV-EWD)

Next, to identify the persistence structure of volatility, we use ([Baruník and Vácha, 2023](#); [Ortu et al., 2020](#)), which decompose the (localised) time series into a collection of independent components. Each of these components represents a horizon (scale) with its own impulse response function that defines the degree of persistence at a given horizon and time. This allows us to decompose the responses of volatility to a unit shock into transient and persistent effects of these shocks in a time-varying manner. This variety of responses allows us to model the heterogeneity in the persistence structure that would otherwise remain hidden in an aggregate model.

In other words, we want to build a model in which the decomposed persistence structure changes smoothly over time. Such a Time-Varying Extended Wold Decomposition (TV-EWD) of volatility can significantly improve forecasting. The main assumption of the model is that we have a locally stationary (zero mean) process $v_{i,T}$, which has the representation $v_{i,T} = \sum_{h=-\infty}^{+\infty} \alpha_{i,T}(h)\epsilon_{i-h}$. Then, following Proposition 1 in [Baruník and Vácha \(2023\)](#), for any $j \in \mathbb{N}$, $k \in \mathbb{N}$ we can define the decomposition of the persistence structure as:

$$v_{i,T} = \sum_{j=1}^{+\infty} \sum_{k=0}^{+\infty} \beta_{i,T}^{(j)}(k) \epsilon_{i-k2^j}^{(j)}, \quad (2)$$

where coefficients $\beta_{i,T}^{(j)}(k) = \frac{1}{\sqrt{2^j}} \left[\sum_{i=0}^{2^j-1} \alpha_{i,T}(k2^j + i) - \sum_{i=0}^{2^j-1} \alpha_{i,T}(k2^j + 2^{j-1} + i) \right]$ denote the time-varying impulse response functions associated with scale j and time-shift $k2^j$ at a fixed point t/T , and $\sum_{k=0}^{\infty} \left(\beta_{i,T}^{(j)}(k) \right)^2 < \infty$ for all j . The innovations $\epsilon_{i-k2^j}^{(j)} = \frac{1}{\sqrt{2^j}} \left(\sum_{i=0}^{2^j-1} \epsilon_{i-k2^j} - \sum_{i=0}^{2^j-1} \epsilon_{i-2^{j-1}-k2^j} \right)$ are independent random variables with $\mathbb{E}\epsilon_t = 0$, $\mathbb{E}|\epsilon_t| < \infty$ and $\mathbb{E}\epsilon_s \epsilon_t = 0$ for $s \neq t$.

The TV-EWD model allows the time series to be decomposed into j uncorrelated persistence components that can vary smoothly over time. Thus, it is a stationary approximation of the process $v_{i,T}$ with time-varying uncorrelated persistent components $v_{i,T}^{(j)}$:

$$v_{i,T} = \sum_{j=1}^{+\infty} v_{i,T}^{(j)}. \quad (3)$$

Each component at scale j can be written as: $v_{i,T}^{(j)} = \sum_{k=0}^{+\infty} \beta_{i,T}^{(j)}(k) \epsilon_{i-k2^j}^{(j)}$. In other words, the decomposition allows us to study the time-varying impulse responses at different scales j . The shape of the impulse response provides information about how shocks propagate at a given scale at a given time. For example, with daily data, the first scale, $j = 1$, shows how a unit shock propagates in 2 days, for the second scale, $j = 2$, in 4 days, and so on.

2.3. Identification of the dynamic persistence in volatility

Having obtained the locally persistent representation of the volatility time series, we next estimate the parameters. Specifically, we start with the time-varying autoregressive (TVP-AR) coefficients model to compute the localised Wold decomposition.

The TVP-AR(p) model of locally stationary process $v_{i,T}$ is for every fixed point t/T defined as:

$$v_{i,T} = \phi_0(t/T) + \phi_1(t/T) v_{i-1,T} + \dots + \phi_p(t/T) v_{i-p,T} + \epsilon_t. \quad (4)$$

To obtain the time-varying coefficient estimates $\hat{\Phi}(t/T) = \left(\hat{\phi}_1(t/T), \dots, \hat{\phi}_p(t/T) \right)'$ we centre the locally stationary process such that $\tilde{v}_{i,T} =$

$v_{i,T} - \hat{\phi}_0(t/T)$. This setting is robust against a possible time trend. The coefficient functions $\phi_i(t/T)$ are estimated by the local linear method. This nonparametric regression approach has several advantages such as efficiency, bias reduction, and adaptation of boundary effects. For more details see [Fan and Gijbels \(1996\)](#) and [Baruník and Vácha \(2023\)](#).

Having the time-varying coefficients, we can compute the local Wold decomposition (Eq. (1)) and then further proceed to decompose the time-varying persistence structure using the TV-EWD (Eq. (2)). Since we have a finite number of T observations, we limit the depth of the decomposition to a finite number of scales J .

$$\hat{v}_{i,T} = \sum_{j=1}^J \hat{v}_{i,T}^{(j)} + \pi_t^{(J)} = \sum_{j=1}^J \sum_{k=0}^{N-1} \hat{\beta}_{i,T}^{(j)}(k) \hat{\epsilon}_{i-k2^j}^{(j)} + \pi_t^{(J)}, \quad (5)$$

where the estimate of (scale specific impulse response coefficients) $\hat{\beta}_{i,T}^{(j)}(k)$ is computed as $\hat{\beta}_{i,T}^{(j)}(k) = \frac{1}{\sqrt{2^j}} \left(\sum_{i=0}^{2^j-1} \hat{\alpha}_{i,T}(k2^j + i) - \sum_{i=0}^{2^j-1} \hat{\alpha}_{i,T}(k2^j + 2^{j-1} + i) \right)$, the scale dependent innovation process is obtained as $\hat{\epsilon}_{i-k2^j}^{(j)} = \frac{1}{\sqrt{2^j}} \left(\sum_{i=0}^{2^j-1} \hat{\epsilon}_{i-k2^j} - \sum_{i=0}^{2^j-1} \hat{\epsilon}_{i-2^{j-1}-k2^j} \right)$ and $\pi_t^{(J)}$ is the residual component defined that is usually very small hence we do not take it into account in the estimation. For more details, see [Ortu et al. \(2020\)](#).

2.4. Forecasting procedure (model)

We want to find a h step-ahead forecast of the energy time series $v_{T+h,T}$. We have the localised version of the original, possibly non-stationary, time series $v_{1,T}, \dots, v_{T,T}$ and estimated values of $\hat{\phi}_0(t/T)$, $\hat{\beta}_{i,T}^{(j)}$ and $\hat{\epsilon}_{i-k2^j}^{(j)}$ we can proceed with the forecasting procedure. Since the scale components $j \in 1, \dots, J$ have different information values, it is important to determine the importance of each component at horizon j . We do this using weights $w^{(j)}$, so we can write:

$$v_{i,T} = \hat{\phi}_0(t/T) + \sum_{j=1}^J w^{(j)} \hat{v}_{i,T}^{(j)} + \eta_{i,T}. \quad (6)$$

Following [Baruník and Vácha \(2023\)](#), we construct the conditional h -step-ahead forecasts as a sum of the conditional expected trend forecast $\mathbb{E}_t[\hat{v}_{T+h,T}^{(0)}]$ and the weighted sum of scale components² as:

$$\mathbb{E}_t[v_{T+h,T}] = \mathbb{E}_t[\hat{v}_{T+h,T}^{(0)}] + \sum_{j=1}^J \hat{w}^{(j)} \mathbb{E}_t[v_{T+h,T}^{(j)}] \quad (7)$$

3. Persistence structure of volatility in major energy commodities

In this section, we demonstrate the importance of identifying the persistence structure in the volatility time series of energy commodities. Namely, we use the daily realised volatilities of three major energy commodities: crude oil, natural gas and regular gasoline. From the high frequency, irregularly spaced data,³ we extract 5-minute intraday prices of futures contracts, using the last available price in the 5-minute window. Contracts are automatically rolled over to provide continuous price records. We compute realised variance as the sum of squared intraday returns ([Andersen et al., 2003](#)) computed from log prices and annualise it to volatility as $100 \times \sqrt{252 \times RV}$.

It should be noted that the character of high-frequency futures data has changed dramatically over the decades, with major changes in the trading system, the introduction of near continuous trading on the CME Globex(R) electronic trading platform in December 2006, and the collection of data with millisecond time stamps in July 2011. In addition, the energy data have experienced both calm and turbulent periods where we can expect very different shocks to drive the time series. We therefore examine the performance of our model over two different periods that reflect these changes in the data.

² We follow [Ortu et al. \(2020\)](#) to forecast $\mathbb{E}_t[v_{i+1,T}^{(j)}]$.

³ Obtained from Tick Data, Inc., which uses data from Globex.

The first period covers January 2010 to December 2022, when energy prices were influenced by several turbulent sub-periods, such as COVID-19, and when millisecond timestamp data were collected almost all day on the electronic trading platform, resulting in large volumes. In contrast, the second period, from January 1993 to December 1999, was a calm and stable period for energy data after the decades of high volatility of the Great Inflation, and high frequency data were only available during open floor trading hours with a second time stamp. Futures were only available to customers through a broker, resulting in much lower average daily volumes compared to the upcoming periods.

3.1. Energy price volatility over the period 2010–2022

During this period, futures contracts are traded on the New York Mercantile Exchange (NYMEX) and transactions are recorded with a millisecond timestamp on a 23-hour basis with a one-hour gap in trading. We redefine the day according to the electronic trading system. In addition, we exclude trades executed on US federal holidays, 24–26 December and 31 December to 2 January, due to low liquidity on these days, which could lead to an estimation bias. The data cover the period from January 2010 to December 2022. The left column of the Fig. 1 shows the time series for all three commodities. Following Baruník and Vácha (2023), we use the multiscale impulse response functions from the TV-EWD model. Looking at the evolution of the volatility persistence structure gives us detailed information about the persistence of the shocks that generate the volatility time series of the energy commodities under study.

We begin by illustrating the persistence structure derived by our model. Fig. 1 shows the evolution of the persistence structure of volatility. We report the persistence at a given scale j as the persistence share of that scale relative to the persistence at other scales. More specifically, we compute the ratio of the first coefficient of the multiscale impulse response function $\hat{\beta}^j(t/T, 1)$ to the sum of the first coefficients of all scales $\sum_j \hat{\beta}^j(t/T, 1)$. The scales j represent 2, 4, 8, 16, 32, 64, 128 days persistence of shocks and are represented in the figure from cold (scale 2) to warm (scale 128) colours (1). As the persistence structure is highly heterogeneous and time-varying, it motivates the use of the TV-EWD model for modelling and forecasting volatility. There are four notable sub-periods in which the persistence structure of volatility changes significantly:

- **Before 2014:** The volatility of oil is dominated by components calculated from shocks of medium persistence (16–64 days), with transitory shocks (2–8 days) of rather low importance. In contrast, natural gas and gasoline volatilities are dominated by transitory shocks.
- **Feb 2014 (annexation of Crimea):** The importance of the more persistent components of oil and gasoline has increased significantly, with gasoline spiking only for very short periods. In December 2014, the oil price was 50% lower than in April 2014, and oil volatility is mostly driven by persistent shocks of 64–128 days, which generate about 60% of the shocks in this period. Gas and petrol also show higher persistence.
- **Onset of the Covid-19 pandemic:** The beginning of 2020 is characterised by the oil price spike followed by a sharp drop, leading to very high volatility. Interestingly, components driven by shocks with a persistence of 64 and 128+ days generate more than 50% of the volatility. Gasoline shows similar behaviour to oil. However, natural gas does not react as strongly to the pandemic, as we observe a mixture of transient and persistent shocks (almost homogeneous structure with equal persistence).
- **Feb 2022:** After the start of the aggression in February 2022, we observe a strong dynamic with the dominance of persistent shocks being overtaken by transitory shocks (2 and 4 day persistence accounts for 40%). For oil, transitory shocks also dominate in this period, but they are not as strong (2 and 4 day persistence

accounts for 35%). This is a consequence of a very unstable and rapidly changing situation, in which transitory shocks to market participants' expectations generated a large proportion of price fluctuations. After February 2022, we clearly observe an increase in overall volatility for all three commodities with mixed persistence structures.

Note that although the level of volatility in these periods is very different, especially Covid-19 is characterised by very high volatility, all turbulent sub-periods are driven by shocks with higher persistence due to the increased uncertainty of market participants. Therefore, compared to other calmer periods, we can observe a higher proportion of higher persistence components driving the series. We will return to this discussion in a later section that examines the calm period of the 1990s. It is also interesting to note that this is not the case for natural gas, which is driven by transitory components after February 2022. With the immense impact of the Russian invasion on the importance of gas, we observe large fluctuations in volatility after February 2022, showing that a lot of transitory uncertainty is driving the time series.

3.1.1. Volatility forecasts

Having identified a rich dynamic in the volatility persistence of the three major energy commodities, we use it to build the forecasting model. The time-varying persistence structure shows that we often witness changes from the transient to the persistent nature of the shocks that dominate volatility. This motivates us to use the TV-EWD model, which can capture the non-trivial, possibly non-stationary dynamics of the volatility series.

To compare the forecasting performance of the TV-EWD model with other commonly used models for estimating the volatility of energy commodities, we use the models that best approximate persistence or long memory behaviour, such as the Heterogeneous Autoregressive Model (HAR) of Corsi (2009) and the EWD model of Ortu et al. (2020). This model also uses the extended Wold decomposition, but faces the same difficulties in capturing the time-varying nature of volatility time series.⁴ While both models capture the unconditional long-run dependence well, they assume stationarity of volatility. To assess the dynamic behaviour, we use both the simple TVP-AR(3) model and the TV-HAR model. The latter model is powerful in capturing the time-varying structure as well as in exploiting the heterogeneous horizon structure (persistence) as it allows for daily, weekly and monthly horizons. This makes the model very versatile and usually hard to beat as it is both highly effective and parsimonious. We save the first 700 observations for the in-sample fit and obtain 2600 forecasts from a rolling window starting the out-of-sample forecasting period with September 15, 2012.

Note that the kernel width is set to 0.3 and the decomposition depth to 7 scales, which is similar for all TV-EWD model settings. This means that the longest horizon considered is 128 days and more (128+). However, the TV-AR(p) model differs for both forecast horizons and commodities. The TV-AR(p) model for CL has $p = 2$ for the 1-day-ahead forecast and $p = 6$ for the 5- and 22-day-ahead forecasts. For NG the TV-AR(p) model has $p = 5$ for all forecasts and for RBOB the TV-AR(p) model has $p = 2$ for the 1-day-ahead forecast and $p = 5$ for the 5- and 22-day-ahead forecasts. All parameters, such as the number of lags of the approximating model, the kernel widths or the degree of decomposition, have been chosen with a view to improving the out-of-sample performance for a given time series and thus reducing losses. One of the key parameters of the model is the number of lags p in the TV-AR(p), which determines the decomposition. For longer horizon forecasts, it is crucial to choose a larger p as it helps to capture a richer structure in terms of persistence. On the other hand, for $h = 1$ forecasts, a more parsimonious model with fewer lags helps to reduce the error. Again, this parameter is selected for each forecast horizon and time series based on the largest reduction in losses.

⁴ In contrast to the paper by Ortu et al. (2020), we use a shorter in-sample period and depth of decomposition. This allows the model to partially account for changes in the persistence structure of volatility.

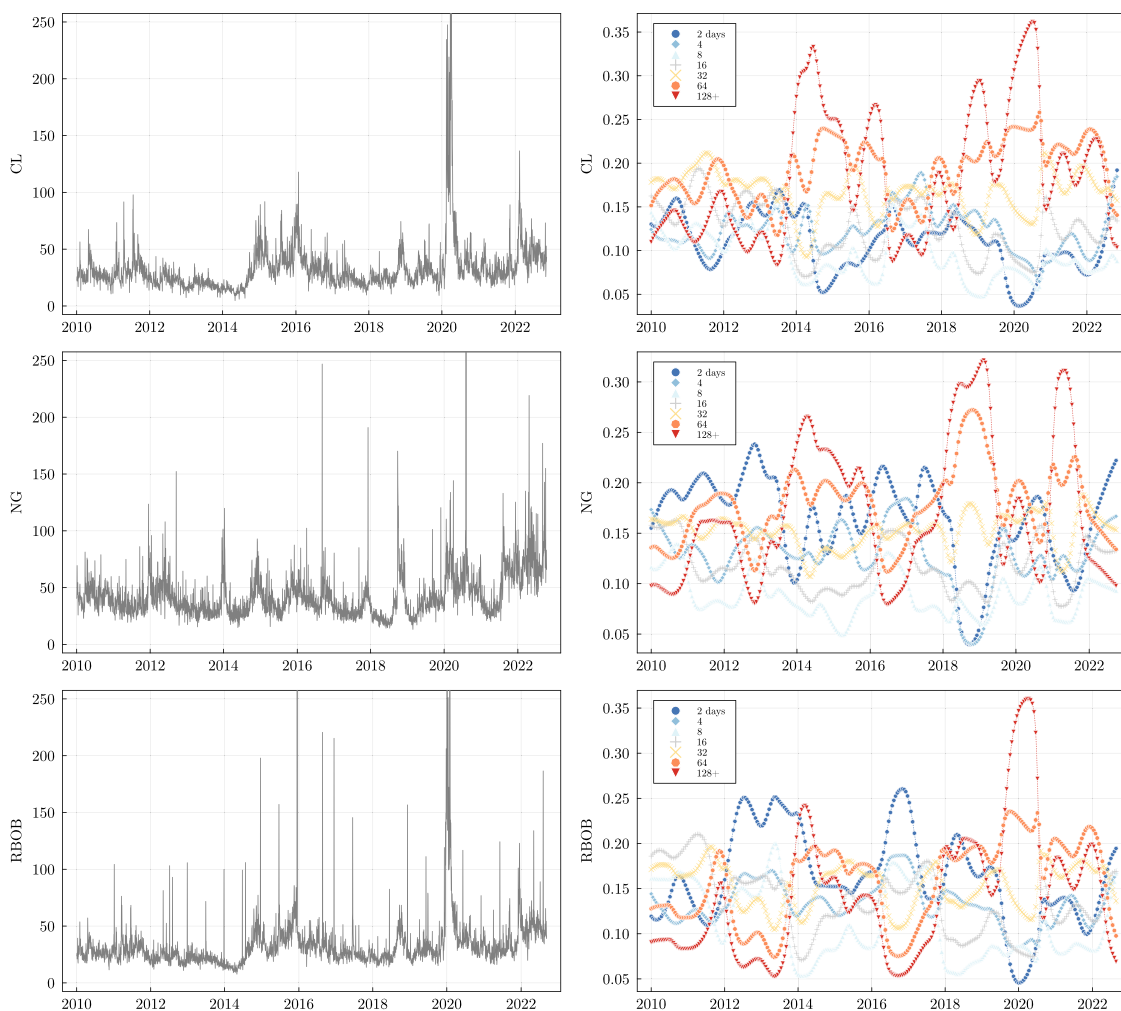


Fig. 1. Left: Realised volatilities for crude oil (CL), natural gas (NG) and gasoline (RBOB). Right: Time-varying persistence structure of CL (top), NG (middle), RBOB (bottom). The plot shows the ratios of $\hat{\beta}^j(t/T, 1) / \sum_j \hat{\beta}^j(t/T, 1)$ on the y-axis, with j corresponding to 2, 4, 8, 16, 32, 64, 128+ days persistence of the shocks represented by cold to warm colours, respectively, over the period from January 2010 to December 2022 on the x-axis. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1 summarises the results of the volatility forecasting performance of the main energy commodities. We report the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE) loss functions. The TV-HAR model is used as a benchmark and all reported losses are relative to the TV-HAR. Thus, a value less (greater) than one indicates that the corresponding model has a better (worse) forecasting performance relative to the benchmark.

In general, the ability of the TV-EWD model to outperform all alternatives increases with increasing forecast horizons. TV-EWD provides significantly better forecasts than the benchmark in all cases for forecasts of 22 days ahead. The superior performance in the MAE also suggests that the forecasts are not systematically biased relative to the other models.

For the 5-day-ahead forecasts, the results are mixed. In terms of RMSE, the TV-EWD again outperforms the benchmark significantly. In terms of MAE, our model still provides significantly better forecasts for natural gas, but the forecasts for the other two commodities are indistinguishable from the TV-HAR. Similarly, the forecast errors of our approach and TV-HAR are the same for the 1-day-ahead forecast for both the RMSE and MAE loss functions.

The EWD model of Ortu et al. (2020) performs well for the 22-day-ahead forecasts as measured by the RMSE loss. This implies that the model is immune to large forecast outliers. Conversely, the poor performance, as measured by the MAE, indicates a systematic bias in

the forecast. This is caused by a very slow adjustment to changes in the level of long-term volatility. This is a consequence of the model’s lack of time-varying capabilities.

3.2. Energy price volatility of the 1993–1999 period

While 2010–2022 period was characterised by more persistent shocks driving energy volatility, together with turbulent times and high volumes, we would like to challenge our model with the calm period of the 1990s, when energy futures enjoyed much less volatile prices and much lower volumes. The trades were recorded with a second time stamp during a regular hour when the floor was open. Note that we use unleaded gasoline (HU) futures instead of RBOB gasoline futures, which were introduced after December 2006 only. During this period, we also exclude trades executed on US federal holidays, 24–26 December and 31 December to 2 January, due to low liquidity on these days, which could lead to an estimation bias. Data cover the period from January 1993 to December 1999. The left column of the Fig. 2 shows the time series for all three commodities, and it is immediately apparent that the series are much more stable and less volatile.

Similarly to the previous period, we use the multiscale impulse response functions from the TV-EWD model to capture the evolution of the persistence structure of volatility and the persistence of the shocks that generate the volatility time series over this period. Similarly to

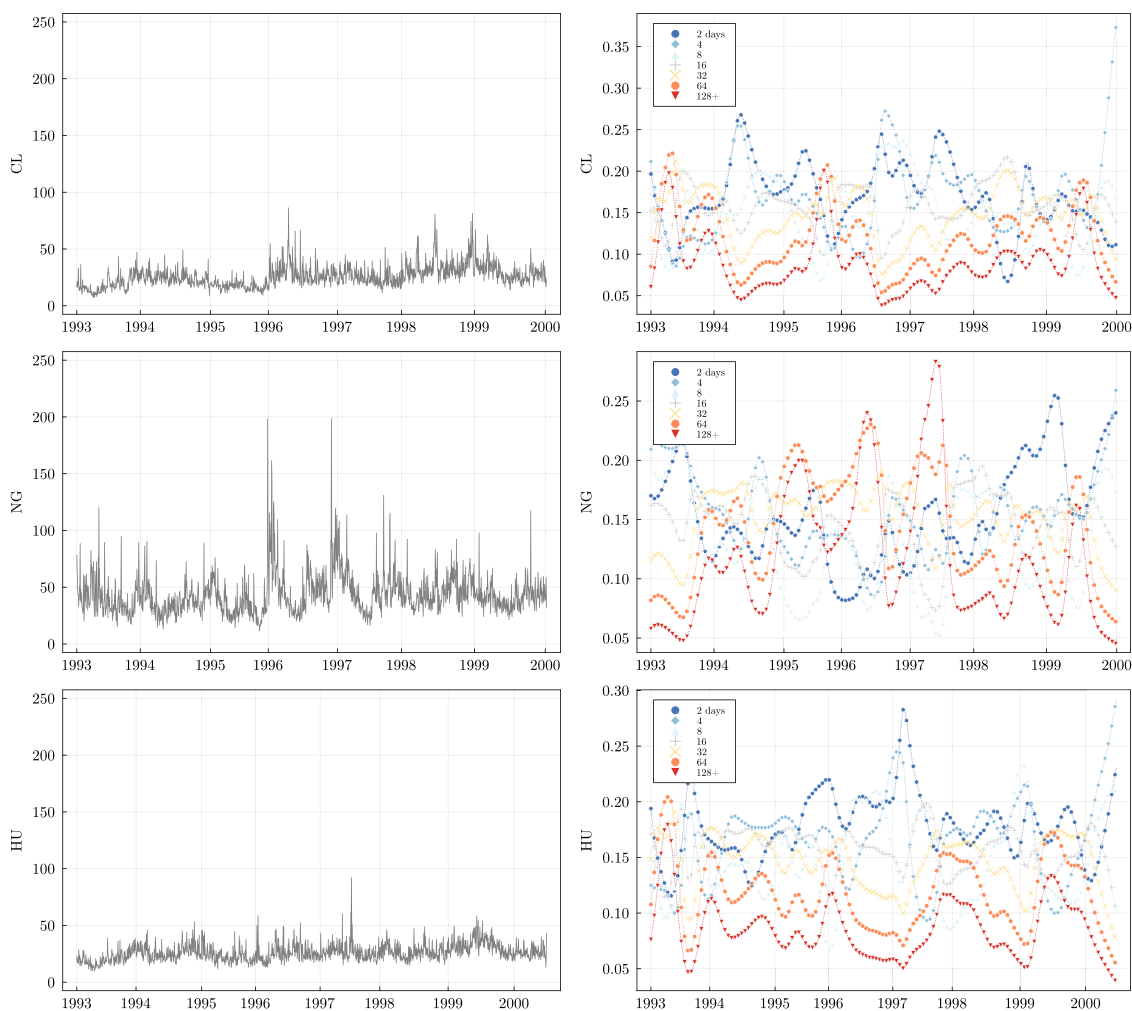


Fig. 2. Left: Realised volatilities for crude oil (CL), natural gas (NG) and gasoline (HU). Right: Time-varying persistence structure of CL (top), NG (middle), HU (bottom). The plot shows the ratios of $\hat{\beta}^j(t/T, 1) / \sum_j \hat{\beta}^j(t/T, 1)$ on the y-axis, with j corresponding to 2, 4, 8, 16, 32, 64, 128+ days persistence of the shocks represented by cold to warm colours, respectively, over the period from January 1993 to December 1999 on the x-axis. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the previous section, we report the persistence at a given scale j as the persistence share of that scale relative to the persistence at other scales.

Fig. 2 shows a stark contrast in the persistence structures compared to the previous turbulent period 2010–2022 discussed in the previous section. The main change is that components with higher persistence play a weaker role in the time series and, conversely, the series seem to be driven by transitory shocks. This is important information for modelling purposes.

3.2.1. Volatility forecasts

The quieter period 1993–1999 is mainly driven by transitory shocks, but at the same time we can observe some change in the structure. Therefore, our TV-EWD model, which is able to capture the non-trivial dynamics, can also be useful for forecasting, although with increasing complexity it may lose its forecasting power, especially because the model is useful for data with strong long-term persistence.

To compare the forecasting performance of the TV-EWD model, we again use the same models in the same setting, saving the first 700 observations for the in-sample fit and leaving the rest for the out-of-sample forecast period.

Note that the kernel width is set to 0.3 and the depth of decomposition to 7 scales, which is similar for all TV-EWD model settings. This means that the longest horizon considered is 128 and more days (128+). However, the TV-AR(p) model differs for both forecast horizons

and commodities. The TV-AR(p) model for CL has $p = 3$ for all forecasts. For NG, the TV-AR(p) model has $p = 5$ for the 1- and 5-day-ahead forecasts and $p = 3$ for the 22-day-ahead forecast, and for HU the TV-AR(p) model has $p = 2$ for the 1-day-ahead forecast and $p = 5$ for the 5- and 22-day-ahead forecasts.

Table 2 summarises the results of volatility forecasting performance over this period. Again, the TV-HAR model is used as the benchmark and all reported losses are relative to the TV-HAR. In general, the TV-EWD model’s ability to outperform all alternatives is weaker during this calm period compared to the turbulent period of 2010–2022. We attribute this mainly to the fact that the high persistence component is weak and therefore time series driven by transitory shocks are more difficult to forecast over multiple horizons. TV-EWD provides improved forecasts in all cases, but the forecast errors are not always significantly different from other models.

4. Conclusion

In this paper, we study volatility of major energy commodities and identify shocks with heterogeneous persistence that vary smoothly over time. Based on such dynamics, we construct a forecasting model that significantly outperforms alternative models, especially at longer horizons. The model provides valuable information about the fundamental behaviour of the volatility time series and is useful for both practitioners and policy makers.

Table 1

The 2010–2022 period: Root mean square error (RMSE) and mean absolute error (MAE) of the TV-EWD model compared to Ortu et al. (2020)'s extended wold decomposition (EWD), time-varying autoregression (TV-AR(3)). All errors are relative to the time-varying HAR (TV-HAR) model of Corsi (2009) over $h = 1$, $h = 5$ and $h = 22$. Thus, the value below one implies that the corresponding model outperforms the TV-HAR benchmark. We use *, ** and *** to denote that a competing model has significantly lower MSE and MAE (is better) in comparison to the TV-HAR model at 90%, 95% and 99% significance levels, respectively. We use †, †† and ††† to denote that a competing model has significantly higher MSE and MAE (is worse) in comparison to the TV-HAR model at 90%, 95% and 99% significance levels, respectively.

		RMSE			MAE		
		$h = 1$	$h = 5$	$h = 22$	$h = 1$	$h = 5$	$h = 22$
CL	EWD	1.373 ^{†††}	1.542 ^{†††}	0.962	1.457 ^{†††}	1.577 ^{†††}	1.185 ^{†††}
	TV-AR3	1.051 ^{††}	1.062 [†]	0.855 ^{***}	1.035 ^{†††}	1.108 ^{†††}	1.024
	TV-EWD	1.009	0.946*	0.807 ^{***}	1.008	0.982	0.960 ^{**}
NG	EWD	1.297 ^{†††}	1.381 ^{†††}	1.235 ^{†††}	1.437 ^{†††}	1.621 ^{†††}	1.358 ^{†††}
	TV-AR3	1.064 ^{†††}	1.134 ^{†††}	1.300 ^{†††}	1.089 ^{†††}	1.253 ^{†††}	1.234 ^{†††}
	TV-EWD	1.003	0.945 ^{**}	0.923 ^{***}	1.006	0.965 ^{***}	0.895 ^{***}
RBOB	EWD	1.249 ^{†††}	1.444 ^{†††}	1.125 ^{†††}	1.481 ^{†††}	1.535 ^{†††}	1.299 ^{†††}
	TV-AR3	1.052 ^{†††}	1.112 ^{†††}	0.997	1.084 ^{†††}	1.186 ^{†††}	1.105 ^{†††}
	TV-EWD	1.000	0.954*	0.880 ^{***}	1.017	1.003	0.951 ^{***}

Table 2

The 1993–1999 period: Root mean square error (RMSE) and mean absolute error (MAE) of the TV-EWD model compared to Ortu et al. (2020)'s extended wold decomposition (EWD), time-varying autoregression (TV-AR(3)). All errors are relative to the time-varying HAR (TV-HAR) model of Corsi (2009) over $h = 1$, $h = 5$ and $h = 22$. Thus, the value below one implies that the corresponding model outperforms the TV-HAR benchmark. We use *, ** and *** to denote that a competing model has significantly lower MSE and MAE (is better) in comparison to the TV-HAR model at 90%, 95% and 99% significance levels, respectively. We use †, †† and ††† to denote that a competing model has significantly higher MSE and MAE (is worse) in comparison to the TV-HAR model at 90%, 95% and 99% significance levels, respectively.

		RMSE			MAE		
		$h = 1$	$h = 5$	$h = 22$	$h = 1$	$h = 5$	$h = 22$
CL	EWD	1.122 ^{†††}	1.206 ^{†††}	1.231 ^{†††}	1.128 ^{†††}	1.301 ^{†††}	1.354 ^{†††}
	TV-AR3	1.013	1.029	0.972	1.024 ^{††}	1.062 ^{†††}	1.013
	TV-EWD	0.995	0.975*	0.978	0.990	0.998	1.002
NG	EWD	1.055 ^{†††}	1.005	0.946 ^{**}	1.093 ^{†††}	1.056 ^{††}	0.914 ^{***}
	TV-AR3	1.032 ^{†††}	1.081 ^{†††}	1.065 ^{†††}	1.050 ^{†††}	1.158 ^{†††}	1.063 ^{†††}
	TV-EWD	1.002	0.968	0.937 ^{***}	1.003	0.995	0.960 ^{**}
HU	EWD	1.080 ^{†††}	1.128 ^{†††}	1.247 ^{†††}	1.106 ^{†††}	1.181 ^{†††}	1.351 ^{†††}
	TV-AR3	1.014	1.042	1.047	1.025 ^{††}	1.093 ^{†††}	1.134 ^{†††}
	TV-EWD	0.998	0.964*	0.918*	0.998	0.979*	0.976

As the paper introduces a novel decomposition of the volatility time series, it also opens up new avenues for further research. It would be interesting to learn about the dynamic persistence structure of other important energy data and, in particular, to use them in other forecasting or economic models.

Our results can be useful to market participants, financial analysts and policymakers trying to understand energy economics data in several ways, as our approach opens up new avenues for modelling and forecasting in the energy economics literature. Researchers can use our model to understand the dynamically heterogeneous structure of shocks driving different time series, and then adjust the models according to the data. Especially interesting venue of the research is linking the identified components in the data with economic models. One can also use our model to identify transitory and persistent components in the data. Finally, one can use the model to improve forecasts in a number of problems where the data contain rich persistence structures. It would be interesting to use our approach in combination with machine learning.

CRedit authorship contribution statement

Jozef Baruník: Writing – original draft, Supervision, Methodology, Conceptualization. **Lukáš Vácha:** Writing – original draft, Methodology, Investigation, Formal analysis, Data curation.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.eneco.2024.107982>.

References

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71 (2), 579–625.

Arouri, M.E.H., Lahiani, A., Lévy, A., Nguyen, D.K., 2012. Forecasting the conditional volatility of oil spot and futures prices with structural breaks and long memory models. *Energy Econ.* 34 (1), 283–293.

Baillie, R.T., Chung, C.-F., Tieslau, M.A., 1996. Analysing inflation by the fractionally integrated ARFIMA–GARCH model. *J. Appl. Econom.* 11 (1), 23–40.

Bandi, F.M., Chaudhuri, S.E., Lo, A.W., Tamoni, A., 2021. Spectral factor models. *J. Financ. Econ.* 142 (1), 214–238.

Bandi, F.M., Tamoni, A., 2022. Spectral financial econometrics. *Econometric Theory* 38 (6), 1175–1220.

Baruník, J., Vácha, L., 2023. The dynamic persistence of economic shocks. Available at SSRN 4467110.

Bollerslev, T., Engle, R.F., 1993. Common persistence in conditional variances. *Econometrica* 61 (1), 167–186.

Charfeddine, L., 2014. True or spurious long memory in volatility: Further evidence on the energy futures markets. *Energy Policy* 71, 76–93.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. *J. Financ. Econ.* 7 (2), 174–196.

Dahlhaus, R., 1996. On the Kullback–Leibler information divergence of locally stationary processes. *Stoch. Process. Appl.* 62 (1), 139–168.

Elder, J., Serletis, A., 2010. Oil price uncertainty. *J. Money Credit Bank.* 42 (6), 1137–1159.

Fan, J., Gijbels, I., 1996. Local Polynomial Modelling and Its Applications: Monographs on Statistics and Applied Probability 66, vol. 66, CRC Press.

Granger, C.W., Ding, Z., 1996. Varieties of long memory models. *J. Econometrics* 73 (1), 61–77.

Granger, C.W., Hyung, N., 2004. Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *J. Empir. Financ.* 11 (3), 399–421.

Hamilton, J.D., 1983. Oil and the macroeconomy since World War II. *J. Polit. Econ.* 91 (2), 228–248.

Hamilton, J.D., 2020. Time Series Analysis. Princeton University Press.

Haugom, E., Langeland, H., Molnár, P., Westgaard, S., 2014. Forecasting volatility of the US oil market. *J. Bank. Financ.* 47, 1–14.

Herrera, A.M., Hu, L., Pastor, D., 2018. Forecasting crude oil price volatility. *Int. J. Forecast.* 34 (4), 622–635.

Kang, S.H., Yoon, S.-M., 2013. Modeling and forecasting the volatility of petroleum futures prices. *Energy Econ.* 36, 354–362.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *Amer. Econ. Rev.* 99 (3), 1053–1069.

Le, T.-H., Boubaker, S., Bui, M.T., Park, D., 2023. On the volatility of WTI crude oil prices: A time-varying approach with stochastic volatility. *Energy Econ.* 117, 106474.

Lu, F., Ma, F., Li, P., Huang, D., 2022. Natural gas volatility predictability in a data-rich world. *Int. Rev. Financ. Anal.* 83, 102218.

Ma, F., Wahab, M.I.M., Huang, D., Xu, W., 2017. Forecasting the realized volatility of the oil futures market: A regime switching approach. *Energy Econ.* 67, 136–145.

Ortu, F., Severino, F., Tamoni, A., Tebaldi, C., 2020. A persistence-based wold-type decomposition for stationary time series. *Quant. Econ.* 11 (1), 203–230.

Ozdemir, Z.A., Gokmenoglu, K., Ekinci, C., 2013. Persistence in crude oil spot and futures prices. *Energy* 59, 29–37.

Sévi, B., 2014. Forecasting the volatility of crude oil futures using intraday data. *European J. Oper. Res.* 235 (3), 643–659.

Stáricá, C., Granger, C., 2005. Nonstationarities in stock returns. *Rev. Econ. Stat.* 87 (3), 503–522.

Wang, Y., Wu, C., 2012. Long memory in energy futures markets: Further evidence. *Resour. Policy* 37 (3), 261–272.

Wang, Y., Wu, C., Yang, L., 2016. Forecasting crude oil market volatility: A Markov switching multifractal volatility approach. *Int. J. Forecast.* 32 (1), 1–9.

Wen, F., Gong, X., Cai, S., 2016. Forecasting the volatility of crude oil futures using HAR-type models with structural breaks. *Energy Econ.* 59, 400–413.

Wold, H., 1938. A Study in the Analysis of Stationary Time Series. Almqvist & Wiksells Boktryckeri.

Zhang, Y., Wei, Y., Zhang, Y., Jin, D., 2019. Forecasting oil price volatility: Forecast combination versus shrinkage method. *Energy Econ.* 80, 423–433.