

Shortages and Machine-Learning Forecasting of Oil Returns Volatility: 1900-2024[#]

Onur Polat ^{*}, Dhanashree Somani ^{**}, Rangan Gupta ^{***} and Sayar Karmakar ^{****}

Abstract

The objective of this paper is to forecast the volatility of the West Texas Intermediate (WTI) oil returns over the monthly period of January 1900 to June 2024 by utilizing the information content of newspapers articles-based indexes shortages for the United States (US). We measure volatility as the inter-quantile range by fitting a Bayesian time-varying parameter quantile regression (TVP-QR) on oil returns. The TVP-QR is also used to estimate skewness, kurtosis, lower- and upper-tail risks, and we control for them in our forecasting model along with leverage. Based on the Lasso estimator to control for overparameterization, we find that the model with moments outperform the benchmark autoregressive model involving 12 lags of volatility. More importantly, the performance of the moments-based model improves further when we incorporate the aggregate metric of shortages and its sub-indexes, particularly those related to the industry and labor sectors. These findings carry significant implications for investors.

Keywords: Oil market volatility; Shortages; Bayesian Time-Varying Parameter Quantile Regressions; Lasso Estimator; Forecasting

JEL Codes: C32; C53; E23; Q40; Q43; Q47

[#] We would like to thank the Editor-in-Chief, Professor Tony Klein, and two anonymous referee for many helpful comments. However, any remaining errors are solely ours. The fourth author's research is partially sponsored by NSF DMS 2124222.

^{*} Department of Public Finance, Bilecik Şeyh Edebali University, Bilecik, Türkiye. Email address: onur.polat@bilecik.edu.tr.

^{**} Department of Statistics, University of Florida, 230 Newell Drive, Gainesville, FL, 32601, USA. Email address: dhanashreesomani@ufl.edu.

^{***} Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa. Email address: rangan.gupta@up.ac.za.

^{****} Department of Statistics, University of Florida, 230 Newell Drive, Gainesville, FL, 32601, USA. Email address: sayarkarmakar@ufl.edu.

1. Introduction

The last two decades have witnessed a process of financialization of the oil market (and the commodity sector in general) which, in turn, has resulted in increased market participation of institutional investors like hedge funds, pension funds, and insurance companies, to the extent that crude oil is now considered a profitable “alternative” investment in the portfolios of financial institutions (Bampinas and Panagiotidis, 2015, 2017). Accurate forecasts of oil returns volatility are crucial for oil investors, as they serve as essential inputs for investment decisions and portfolio management.

Recently, there has been a notable increase in supply chain disruptions due to so-called “rare disaster events” associated with heightened geopolitical and climate risks, and, of course, the COVID-19 pandemic. Against this backdrop, and given the importance of determining the future path of price volatility from the perspective of oil market players, the objective of this paper is to forecast the volatility of the West Texas Intermediate (WTI) oil returns over the monthly period of January 1900 to June 2024 by utilizing the information content of newspapers articles-based indexes of shortages for the United States (US), as developed by Caldara et al. (2024). It is worthwhile noting that, by studying the longest possible sample of data available for shortages, we are able to avoid any sample-selection-bias, and track, in a robust manner, the historical effects on the forecast ability of the oil market due to events such as major coal mines strikes during the turn of the 20th century, two World Wars, the Suez Crisis in 1956, the oil shocks during 1970s, the Iraqi invasion of Kuwait in 1990, besides the recent COVID-19 pandemic.

Theoretically speaking, supply chain disruptions, serving as a proxy for disaster events, besides emanating from strikes and price controls, are prone to increase oil returns volatility by contributing to jump risk in oil prices, which, in turn, constitutes a large part of the variation in crude oil prices (Demirer et al., 2018, 2022a; Bouri et al., 2021). Moreover, supply chain constraints have been associated with negative effects on output (Ascari et al., 2024; Burriel et al., 2024; Diaz et al., 2024; Ginn, 2024), which is likely to translate into higher macroeconomic uncertainty (Ludvigson et al., 2021), and enhance the volatility in the oil market through the “theory of storage” (Cepni et al., 2022; Gupta and Pierdzioch, 2022; Wen et al., 2024).¹ Finally, a

¹ As discussed in detail by Fama and French (1987), this theory related to the determination of commodity market volatility states that, rising macroeconomic uncertainty causes risk-averse commodity producers to increase holding

contraction in output due to the supply restrictions can also adversely impact demand for oil, leading to lower trading volumes and, hence, reduced oil returns volatility (Demirer et al., 2020; Salisu et al., 2022a).² In other words, there are multiple theories influencing the impact of shortages on oil returns volatility, though the sign of the effect can be ambiguous depending on the strength of the channels.

To econometrically conduct our analysis, realizing the possibility of structural changes over 125 (1900-2024) years of data, we first employ a Bayesian time-varying parameter quantile regression (TVP-QR) model on oil returns, following Pfarrhofer (2022), to obtain robust calculations of the corresponding volatility as an inter-quantile range from the conditional quantiles of the univariate framework. A further advantage of this approach is that we are also able to compute, from the estimated conditional quantiles, additional oil market moments such as, skewness, kurtosis, lower- and upper-tail risks, which, along with leverage (i.e., negative only oil returns) have been shown to play important roles in forecasting historical monthly estimates of oil returns volatility (Salisu et al., 2022b; Gupta et al., 2023).³ Then secondly, we utilized a linear predictive regression model, but estimated using the popular least absolute shrinkage and selection operator (Lasso) estimator (Tibshirani, 1996), given that our forecasting models, over rolling-windows of 120 months (i.e., 10 years), can contain between 12 to 21 predictors, associated with lags of volatility, oil market moments, and various indexes of shortages.

To the best of our knowledge, this is the first paper to analyze the role of monthly indexes of shortages in forecasting the corresponding volatility of WTI real oil returns based on a machine-learning approach covering more than a century of historical data. By doing so, we contribute to the already existing large literature on forecasting of oil returns volatility by studying the role of supply-side constraints, with extant studies having considered a plethora of economic variables by

of physical inventory, as future cash flows are expected to be negatively impacted causing a rise in the convenience yield to produce increased commodity-price volatility.

² Fall in oil prices can also lead to increase or reduction of oil market volatility through the traditional “leverage effect” signalling bad-news (Geman and Shih, 2009), or due to fall in fear among oil consumers (Aboura and Chevallier, 2013), respectively.

³ See, for example, Asai et al. (2020), Gkillas et al. (2020), Degiannakis and Filis (2022), Demirer et al. (2022b), Luo et al. (2022), who have stressed on the relevance of these moments in the context of realized volatility derived from intraday oil price data.

utilizing a wide array of econometric frameworks (see, Degiannakis and Filis (2017), Salisu et al. (2022c) and Gupta et al. (2024) for detailed reviews of this literature).

The remainder of this paper is organized as follows: Section 2 describes the data, while Section 3 lays out the basic details of the forecasting framework and its estimation process. Section 4 presents the results, with Section 5 concluding the paper.

2. Data

As far as the monthly WTI spot oil price is concerned, from which we compute log-returns in percentages (i.e., first-differences of the natural logarithms of prices multiplied by 100), the series is obtained from the Global Financial Data.⁴ The log-returns are subsequently fitted to the Bayesian TVP-QR model, with detailed specifications provided in the Appendix at the end of the paper.. Once we obtain the fitted values oil log-returns (\widehat{y}_{pt}) at the conditional quantiles, i.e., $p = 0.01, 0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95,$ and 0.99 , we, along with leverage (LEV) — a time series involving the months that correspond to only negative real raw (i.e., unfitted) oil log-returns, we obtain our estimates of lower (LTR)- and upper (UTR)-tail risks, skewness (SKEW) and kurtosis (KURT), to forecast our metric of oil log-returns volatility, namely the inter-quantile range (IQR). Note that, following Foglia et al. (2025), $IQR = \widehat{y}_{0.90} - \widehat{y}_{0.10}$; $LTR = \widehat{y}_{0.05}$; $UTR = \widehat{y}_{0.95}$; $SKEW = (\widehat{y}_{0.90} + \widehat{y}_{0.10} - 2\widehat{y}_{0.50}) / (\widehat{y}_{0.90} - \widehat{y}_{0.10})$, and; $KURT = (\widehat{y}_{0.99} - \widehat{y}_{0.01}) / (\widehat{y}_{0.75} - \widehat{y}_{0.25})$.

We now turn our attention to the main predictors, i.e., the shortage indexes, which are basically monthly newspapers-based indicator that measures the intensity of shortages of materials, goods, labor, and energy in the U.S., with the individual indexes (for energy, food, industry, and labor shortages) adding up to the overall index.⁵ Caldara et al. (2024) outline the construction of these indexes using a sample of approximately 20,000 news articles per month, starting in 1900 and continuing to the present. This dataset comprises around 25 million articles published in six major U.S. newspapers.: The Boston Globe, The Chicago Tribune, The Los Angeles Times, The New York Times, The Wall Street Journal, and The Washington Post. Each month during the sample period, the shortage indexes track the number of articles discussing shortages in energy, food,

⁴ See: <https://globalfinancialdata.com/>.

⁵ The data can be publicly downloaded from: <https://www.matteoiacoviello.com/shortages.html>.

industry, or labor. Due to the extensive sample period and the broad text corpus analyzed by Caldara et al. (2024), these indexes capture a wide range of historical domestic and global events.

Note that, each month, the index counts articles that discuss shortages in energy, food, industry or labor markets normalized by the total number of articles. A higher index value indicates more intense shortages. The steps to followed to create the search query is as follows: First, Caldara et al. (2024) construct a broader set of articles, which must mention at least one shortage term namely, ‘shortage’, ‘scarcity’, ‘bottleneck’, or ‘rationing’, along with an economics-related term, such as ‘economy’, ‘market’, or ‘commerce’. Second, these authors draw a random sample of about 3300 articles from the above set of articles and extract the 1000 most frequent collocates, words appearing within five words of the shortage terms. From these collocates, Caldara et al. (2024) select those indicative of shortages in specific sectors. The most common collocates, excluding stopwords, involve ‘oil’, ‘water’, ‘war’, ‘time’, ‘coal’, ‘days’, ‘food’, ‘cars’, ‘people’, ‘government’, ‘million’, ‘labor’, ‘state’, ‘home’, ‘steel’, and ‘fuel’. Generic words which convey little information about shortages, such as ‘time’, ‘days’, ‘people’, and ‘government’, are subsequently removed. The remaining collocates are grouped into four topics: food, industry, labor, and energy. Third, using the lists of shortage terms, economics terms, and topic-specific terms, the search query is constructed. An article is included in the shortage index if it meets two conditions: (a) a shortage term must appear within five words of a topic-specific term, and; (b) the article must contain at least one economics term. If an article meets the first condition for multiple topic-specific categories, it is counted once for each, such the total number of shortage articles is the sum across categories. This approach gives greater weight to articles discussing multiple types of shortages, thus enhancing the informational content of the index.⁶

At the time of writing this paper, our sample spans the monthly period from January 1900 to June 2024. Figure 1 illustrates the WTI log-returns, the corresponding interquartile range (IQR), and various predictors, including LTR, UTR, SKEW, KURT, and the shortage indexes. The figure clearly demonstrates that the shortage indexes tend to rise significantly during periods of economic

⁶ It must be pointed out that, the article classification into four topics is further corroborated by an ex-post Latent Dirichlet Allocation (LDA) analysis on a sample of articles meeting the inclusion criteria. Specifically, for each article, LDA provides a probability vector indicating the degree to which each of the four topics is present, with these probabilities allowing Caldara et al. (2024) to validate the topic classification used in the searches, and also enable them to track trends in topic prevalence over time.

turmoil, such as the World Wars and the oil crises of the 1970s (energy). Notably, they also spiked during the COVID-19 pandemic, reaching their highest level in the past 40 years. However, there are some other past peaks, especially associated with the World Wars I and II, and the oil crises, which are of comparable or larger in size. One must note that, in general, the IQR of oil returns also tends to track quite well the fluctuations of the shortages indexes, barring the post-World II period until the effective emergence of Organization of Petroleum Exporting Countries (OPEC) in the early 1970s, which was characterized by a various types of government controls in the oil market. Moreover, given that there is also considerable variation across the shortages sub-indexes, it is warranted to utilize these sub-indexes of energy (Shortages_Energy), food (Shortages_Food), industry (Shortages_Industry), and labor (Shortages_Labor) in our forecasting experiment, besides the overall shortages (Shortages_Aggregate) index.

[INSERT FIGURE 1]

3. Forecasting Model and Results

Our linear predictive regression model is given by:

$$IQR_{t+h} = a + X_t' b + \epsilon_{t+h} \quad (1)$$

where h denotes the forecast horizon (in months); IQR_{t+h} denotes the average of volatility between periods of time t and $t + h$, with volatility being captured by the inter-quantile range (IQR) obtained from the estimated quantiles of the Bayesian TVP-QR under the dhs-TVS prior setting; X_t is the vector of our predictors, which vary according to the models under consideration, and are described below; ϵ_{t+h} denotes the usual disturbance term, and; a denote the constant, i.e., the conditional mean of IQR_{t+h} , and b is a vector of coefficients in \mathbf{R}^n , corresponding to X_t involving n predictors, that needs to be estimated.

As far as our benchmark model (M1) is concerned, X_t includes 12 lags of IQR_t , chosen based on the Akaike Information Criterion (AIC). Model 2 (M2) builds on M1 by including LEV, SKEW, KURT, LTR and UTR, and Model 3 (M3) adds Shortages_Aggregate to the predictors in M2. M4 extends the covariate set of M2 by including the four sub-indexes of shortages, i.e., Shortages_Energy, Shortages_Food, Shortages_Industry, and Shortages_Labor.

Given the above set-up, M1, M2, M3, and M4 contain 12, 17, 18 and 21 predictors, and with us using a rolling-window prediction structure involving 120 monthly observation each time, we use the popular Lasso shrinkage estimator to accommodate for the possibility of overparameterization and associated poor out-of-sample forecasting performances. The idea underlying this shrinkage estimator is to reduce the dimension of a regression model in a data-driven manner to improve the accuracy of predictions derived from the penalized model as follows:

$$\hat{b} = \underset{b}{\operatorname{argmin}} \left(\sum_{t=1}^T (IQR_{t+h} - a - X_t' b)^2 + \lambda \sum_{j=1}^n |b_j| \right) \quad (2)$$

where T denotes the number of observations used to estimate the forecasting model; λ is a shrinkage parameter, and n corresponds to the number of coefficients that are subject to the shrinkage process. Hence, the Lasso estimator adds to the standard quadratic loss function in ordinary least squares (OLS) estimator a penalty term that increases in the absolute value of the coefficients. The Lasso estimator, thereby, shrinks a few co-ordinates towards zero, where the effect of this shrinking must be balanced against the resulting effect on the quadratic loss function. The final non-zero coefficients indicate the corresponding predictors are significant.

As far as our forecasting set-up is concerned, we use the first 10 years, i.e., 1900:01 to 1909:12 as our in-sample and then roll this window of 120 months forward by leaving out one initial observation till 2024:06- h to produce $h = 1$ -, 3-, 6-, 12-, and 24-month-ahead forecasts. Note that, to prevent any possibility of a look-ahead-bias in the derived values of SKEW, KURT, LTR and UTR, the Bayesian TVP-QR was re-estimated across the various quantiles using a rolling-window of 120 months as well.

Table 1 presents the Root Mean Square Errors (RMSEs) of models M1, M2, M3 and M4. As can be seen, M2, M3 and M4 outperforms the benchmark M1 model by massive margins, with the gains ranging from 39% (at $h = 24$) to 86% (at $h = 6$), confirming the findings in the literature (see, for example, Asai et al. (2020), Gkillas et al. (2020), Degiannakis and Filis (2022), Demirer et al. (2022), Luo et al. (2022)) underlining the importance of moments in forecasting IQR_{t+h} , i.e., oil returns volatility across various forecasting horizons, and in the process, justifies our decision to control for them when analyzing the predictive role of shortages. The RMSEs across M2, M3 and M4 are, however, quite close to each other, though the latter two produce slightly lower RMSEs than M2, suggesting the role played by aggregate and sector-specific shortages. Between M3 and

M4, the former records slightly higher RMSEs than the latter, suggesting that disaggregated information on shortages tend to matter more. Having depicted the role of shortages in producing forecasting gain for the inter-quantile range of oil returns, i.e., its volatility, in terms of the RMSEs across short-, medium-, and long-run, it is important to determine whether these gains are statistically significant or not, even though they might seem marginal relative to the information content of the moments.

[INSERT TABLE 1]

In this regard, since M2, M3 and M4 nests the benchmark M1, as well as M3 and M4 nests M2, we utilize the Clark and West (2007; CW) test of statistical significance of forecasts involving two nested competing models. The null hypothesis posits that both models exhibit equal predictive performance, whereas the alternative hypothesis suggests that the competing model outperforms the benchmark model. Hence, the CW test is a one-sided test.

As can be seen from the p -values of the CW test reported in Table 2, M2, i.e., the model with the moments (LEV, SKEW, KURT, LTR and UTR) outperforms the benchmark model (M1) containing only 12 lags of the volatility at the 5% level of significance consistently over 1-, 3-, 6-, 12-, and 24-month-ahead forecasts of IQR_{t+h} . As outlined above, this finding is in line with the extant literature that has highlighted the important predictive role for moments in forecasting oil returns volatility. Not surprisingly, M3 and M4 also outperform M1 at the 1% level of significance for all the 5 forecast horizons considered, suggesting that moments and shortages, both at the aggregate- and sectoral-level, can produce forecasting gains relative the benchmark model involving only lags of oil returns volatility. More importantly, both M3 and M4, containing either the overall shortages index (Shortages_Aggregate), or the sub-indexes (Shortages_Energy, Shortages_Food, Shortages_Industry, and Shortages_Labor) respectively, strongly outperforms M2 at the 1% level of significance for $h = 1, 3, 6, 12, \text{ and } 24$. Additionally, we employ the Diebold and Mariano (1995) test to compare the forecast performance of the two non-nested models, M3 and M4, assessing whether disaggregated shortage information provides better predictive accuracy than the overall index. Accordingly, the null hypothesis states that the forecasts from M3 and M4 are equally accurate, while the alternative hypothesis posits that M4 outperforms M3.. We find that the null is strongly rejected at the 1% level of significance for $h = 1, 3, \text{ and } 6$, and at the 5% and 10% levels of significance respectively, for 12-, and 24-month-ahead forecasts. In other words, we

confirm that aggregate and, especially, sub-indexes of shortages have important and predictive content for movements of the inter-quantile range, i.e., volatility, of oil returns through the history covering 125 years of data, at short-, medium-, and long-term horizons.

As a complementary exercise in terms of the p -values of the CW-test, to check which sub-index of shortages among the four matters relatively more, we carried out our forecasting experiment with a modified version of M4, wherein, instead of considering all the sub-indexes together, we utilized each of them individually. In other words, we defined M4(I), M4(II), M4(III) and M4(IV) as models corresponding to a structure which involves M2 and Shortages_Energy or Shortages_Food or Shortages_Industry or Shortages_Labor, respectively. Interestingly, M4(I) fails to outperform M2 at all horizons, while M4(II) does better than M2 only at longer horizons of $h = 12$ and 24 , at the 5% level of significance. The strongest forecasting gains for oil returns volatility relative to M2, i.e., the model of moments, comes from M4(III) and M4(IV) at the 1% level of significance recorded by the CW test across all forecast horizons. **In other words, shortages related to industries, where role of oil is undeniable as an input, and the labor market, possibly governing aggregate demand in the economy, seems to matter the most among the four types of shortages concerned. In other words, shortages in industry and the labor market tends to signal negative effects on output and earnings due to reduced economy-wide demand and higher macroeconomic uncertainty, which in turn can increase the volatility in the oil market through the “theory of storage”. Alternatively, a contraction in output can also lead to lower trading volumes of oil and, hence, produce a reduction in future oil returns volatility. In other words, shortages associated with industries and the labor market can capture the up- and down-movements, and hence, fluctuations (volatility in the oil market) relatively accurately. This is more so, since, as can be seen from Figure 1, these two types of shortages have been comparatively more frequent with strong peaks. The fact that the sub-index shortages associated the energy sector cannot outperform the model for oil returns volatility involving its moments, is reflective of the fact that information involving own-sector-specific shortages are possibly already contained in the moments of its returns (Wen et al., 2021).**

[INSERT TABLE 2]

4. Conclusion

We forecast the volatility of the WTI oil returns over the monthly period of January 1900 to June 2024 by utilizing the information content of newspapers articles-based indexes of aggregate and sectoral shortages of the US. We measure volatility employing the interquantile range, estimated through a Bayesian time-varying parameter quantile regression on oil returns to derive the underlying fitted quantiles. This approach also allows us to estimate skewness, kurtosis, lower- and upper-tail risks, and incorporate them into our forecasting model alongside leverage.. Based on the shrinkage estimation using the Lasso estimator to control for overparameterization, we find that the model with moments outperform the benchmark autoregressive model, but the performance of the former, in turn, is improved further when we incorporate the role of the aggregate metric of shortages, as well as its sub-indexes considered all together, with highest gains emanating from supply constraints in the industry and labor sectors.

Our historical findings hold significant investment implications, as they demonstrate that incorporating information on shortages—particularly from the industrial and labor sectors—enable oil market investors to more accurately forecast oil return volatility. This, in turn, aids in the construction of optimal investment portfolios.⁷ Moreover, from a policy perspective, precise prediction of the future path of oil-market volatility by exploiting the informational content of shortages using machine learning could aid policymakers to propose and implement fiscal and monetary policies ahead of time so as to stabilize strong fluctuations in economic activity (originating initially from shortages), which are likely to be exacerbated, given its strong historical association with variability in oil prices (van Eyden et al., 2019).

As part of future research, it would be interesting to extend our analysis to other (energy and non-energy) commodity markets, possibly using intraday data, though over shorter sample periods, given the availability of newspapers-based daily supply bottlenecks indicators of major economies

⁷ An alternative way to assess the economic implications of our findings is to directly use an investor's utility function, as in Cenesizoglu and Timmermann (2012), to measure the net benefits from utilizing information on Shortages_Aggregate in predicting oil market volatility (Model M3), relative to not using the same and relying only on moments (Model M2). To this end, in Table A1 in the Appendix, we present the annualized net certainty equivalent returns (CER) in % for three levels of increasing risk aversion ($\gamma= 3, 5$ and 10). As can be seen, the net CER is always positive across $h = 1, 3, 6, 12$ and 24 , suggesting that oil investors are likely to make utility gains, especially at shorter horizons and/or higher degree of risk aversion, by utilizing the information content of shortages over and above its own moments.

around the world as developed by Burriel et al. (2024).⁸ Moreover, though here we focus only on the Lasso approach to forecasting oil returns volatility due to shortages, as part of further work, we can indeed investigate other linear and nonlinear machine learning approaches to confirm if the performance of forecasting can be improved further.⁹

⁸ Given the observation by Gozgor et al. (2023) that global supply chain constraints spills over to commodity market returns, we used the k -th order nonparametric causality-in-quantiles test of Balcilar et al. (2018) to analyse the causal effect on the entire conditional distributions of the log-returns and squared log-returns (i.e., volatility) of various S&P GSCI Total Return Indexes associated with the overall commodity market, and various sub-indexes (Precious Metals, Energy, Non-Energy, Industrial Metals, Agriculture, Livestock, Softs, and Grains), due to a world metric of supply bottlenecks over the daily period of 15th January 2010 to 7th June 2024. The S&P GSCI Total Return Indexes are derived from Refinitiv Datastream, while the global index of supply bottlenecks (GSBI) is obtained using the first principal component of the indexes of Burriel et al. (2024) for the US, the United Kingdom (UK), Germany, France, Italy, Spain and China, available at: <https://www.bde.es/wbe/en/areas-actuacion/analisis-e-investigacion/recursos/indices-de-cuellos-de-botella-en-la-oferta-basados-en-articulos-de-prensa.html>. As can be seen from the standard normal tests statistics reported in Table A2 in the Appendix of the paper, predictability from the GSBI holds over entire conditional distributions (conditional quantiles of 0.10 to 0.90) of both returns and squared returns for the overall and sub-sectors of the commodity market, but with a stronger influence on volatility. Interestingly, the effect on the returns and volatility of the energy-based commodities stands out relative to the other sectors. This serves as a preliminary motivation to investigate this question further using intraday data to forecast realized volatility of not only the aggregate commodity market but also its sub-sectors with these supply bottlenecks indexes at both country- and global-level, after controlling for realized moments.

⁹ It must be realized that the focus here is on a horserace across predictors of oil returns volatility, and in particular that of shortages, rather than models. Despite this, based on the suggestion of an anonymous referee, we utilized a Long Short-Term Memory model, which is a type of Recurrent Neural Network (RNN), to re-conduct our forecasting exercise. As with the Lasso, we found that the model with shortages outperformed the benchmark, and the one with moments in a statistically significant manner, thus ensuring the robustness of our findings. However, more importantly, the Lasso was found to statistically outperform the LSTM model across the various forecasting models, hence suggesting the superiority of the former. Complete details of these results are available upon request from the authors.

References

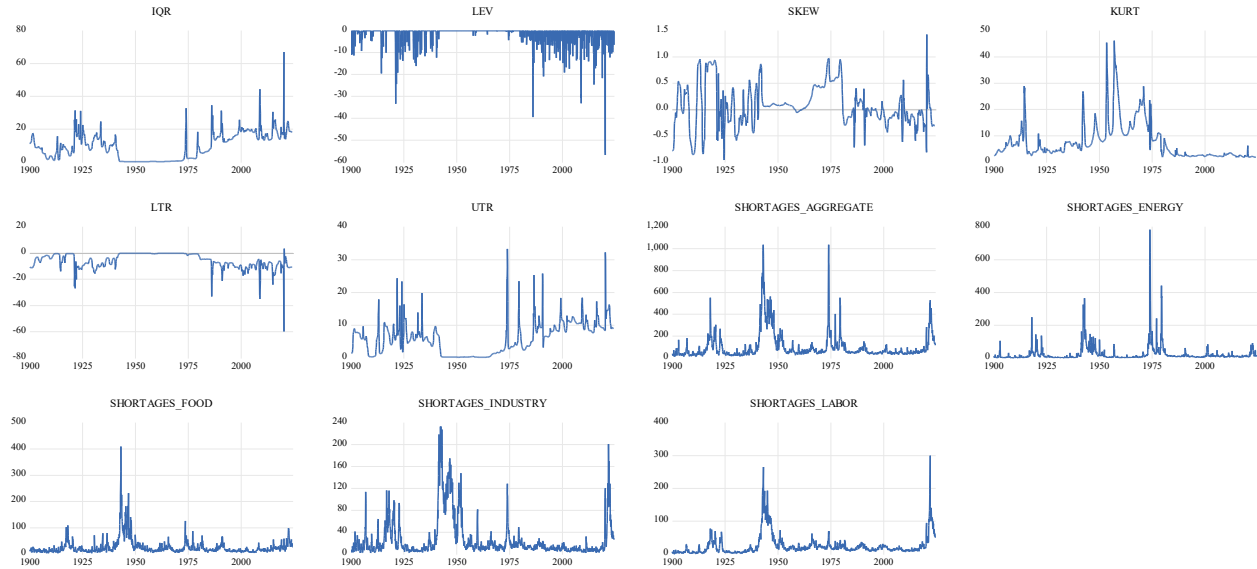
- Aboura, S., and Chevallier, J. (2013). Leverage vs. Feedback: Which effect drives the oil market? *Finance Research Letters*, 10(3), 131-141.
- Asai, M. Gupta, R., and McAleer, M. (2020). Forecasting volatility and co-volatility of crude oil and gold futures: Effects of leverage, jumps, spillovers, and geopolitical risks. *International Journal of Forecasting*, 36(3), 933-948.
- Ascari, G., Bonam, D., and Smadu, A. (2024). Global supply chain pressures, inflation, and implications for monetary policy. *Journal of International Money and Finance*, 142(C), 103029.
- Balcilar, M., Gupta R., Nguyen D.K., and Wohar, M.E. (2018). Causal effects of the United States and Japan on Pacific-Rim stock markets: nonparametric quantile causality approach. *Applied Economics*, 50(53), 5712-5727.
- Bampinas, G., and Panagiotidis, T. (2015). On the relationship between oil and gold before and after financial crisis: Linear, nonlinear and time-varying causality testing. *Studies in Nonlinear Dynamics & Econometrics*, 19(5), 657-668.
- Bampinas, G., and Panagiotidis, T. (2017). Oil and stock markets before and after financial crises: A local Gaussian correlation approach. *The Journal of Futures Markets*, 37(12), 1179-1204.
- Bouri, E., Gupta, R., Pierdzioch, C., and Salis, A.A. (2021). El Niño and Forecastability of Oil-Price Realized Volatility. *Theoretical and Applied Climatology*, 144(3-4), 1173-1180.
- Burriel, P., Kataryniuk, I., Moreno Pérez, C., and Viani, F. (2024). A New Supply Bottlenecks Index Based On Newspaper Data. *International Journal of Central Banking*, 20(2), 17-69.
- Caldara, D., Iacoviello, M., and Yu, D. (2024). Measuring Shortages since 1900. Working Paper: https://www.matteoiacoviello.com/research_files/SHORTAGE_PAPER.pdf.
- Cepni, O., Gupta, R., Pienaar, D., and Pierdzioch, C. (2022). Forecasting the realized variance of oil-price returns using machine learning: Is there a role for U.S. state-level uncertainty? *Energy Economics*, 114(C), 106229.
- Cenesizoglu, T., and Timmermann, S. (2012). Do return prediction models add economic value? *Journal of Banking and Finance*, 36(11), 2974-2987.
- Clark, T.D., and West, K.D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291-311.
- Degiannakis, S., and Filis, G. (2017). Forecasting oil price realized volatility using information channels from other asset classes. *Journal of International Money and Finance*, 76(C), 28-49.
- Degiannakis, S., and Filis, G. (2022). Oil price volatility forecasts: What do investors need to know? *Journal of International Money and Finance*, 123(C), 102594.
- Demirer, R., Gkillas, K., Gupta, R., and Pierdzioch, C. (2022b). Risk aversion and the predictability of crude oil market volatility: A forecasting experiment with random forests. *Journal of the Operational Research Society*, 73(8), 1755-1767.
- Demirer, R., Gupta, R., Nel, J., and Pierdzioch, C. (2022a). Effect of rare disaster risks on crude oil: evidence from El Niño from over 145 years of data. *Theoretical and Applied Climatology*, 147(1-2), 691-699.
- Demirer, R., Gupta, R., Pierdzioch, C., and Shahzad, S.J.H. (2020). The predictive power of oil price shocks on realized volatility of oil: A note. *Resources Policy*, 69(C), 101856.
- Demirer, R., Gupta, R., Suleman, M.T., and Wohar, M.E. (2018). Time-varying rare disaster risks, oil returns and volatility. *Energy Economics*, 75(C), 239-248.

- Diaz, E.M., Cunado, J., and de Gracia, F.P. (2023). Commodity price shocks, supply chain disruptions and US inflation. *Finance Research Letters*, 58, 104495.
- Diebold, F.X., and Mariano, R.S. (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statistics*, 13(3), 253-263.
- Fama, E.F., and French, K.R. (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1), 55-73.
- Foglia, M., Plakandaras, V., Gupta, R., and Ji, Q. (2024). Long-span multi-layer spillovers between moments of advanced equity markets: The role of climate risks. *Research in International Business and Finance*, 74(C), 102667.
- Geman, H., and Shih, Y.F. (2009). Modeling Commodity Prices under the CEV Model. *Journal of Alternative Investments*, 11(3), 65-84
- Ginn, W. (2024). Global supply chain disruptions and financial conditions. *Economics Letters*, 239, 111739.
- Gkillas, K., Gupta, R., and Pierdzioch, C. (2020). Forecasting realized oil-price volatility: The role of financial stress and asymmetric loss. *Journal of International Money and Finance*, 104(C), 102137.
- Gozgor, G., Khalfaoui, R., and Yarovaya, L. (2023). Global supply chain pressure and commodity markets: Evidence from multiple wavelet and quantile connectedness analyses. *Finance Research Letters*, 54(C), 103791.
- Gupta, R., Ji, Q., Pierdzioch, C., and Plakandaras, V. (2023). Forecasting the conditional distribution of realized volatility of oil price returns: The role of skewness over 1859 to 2023. *Finance Research Letters*, 58(Part C), 194501.
- Gupta, R., Nielsen, J., and Pierdzioch, C. (2024). Stock market bubbles and the realized volatility of oil price returns. *Energy Economics*, 132(C), 107432.
- Gupta, R., and Pierdzioch, C. (2022). Forecasting the realized variance of oil-price returns: a disaggregated analysis of the role of uncertainty and geopolitical risk. *Environmental Science and Pollution Research*, 29(29), 52070-52082.
- Ludvigson, S.C., Ma, S., and Ng, S. (2021). Uncertainty and business cycles: Exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics*, 13(4), 369-410.
- Luo, J., Demirer, R., Gupta, R., and Ji, Q. (2022). Forecasting oil and gold volatilities with sentiment indicators under structural breaks. *Energy Economics*, 105(C), 105751.
- Pfarrhofer, M. (2022). Modeling tail risks of inflation using unobserved component quantile regressions. *Journal of Economic Dynamics and Control*, 143(C), 104493.
- Salisu, A.A., Gupta, R., Bouri, E., and Ji, Q. (2022c). Mixed-frequency forecasting of crude oil volatility based on the information content of global economic conditions. *Journal of Forecasting* 41(1), 134-157.
- Salisu, A.A., Gupta, R., and Demirer, R. (2022c). Global financial cycle and the predictability of oil market volatility: Evidence from a GARCH-MIDAS model. *Energy Economics*, 108(C), 105934.
- Salisu, A.A., Pierdzioch, C., and Gupta, R. (2022b). Oil Tail Risks and the Forecastability of the Realized Variance of Oil-Price: Evidence from Over 150 Years of Data. *Finance Research Letters*, 46(Part B), 102378.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1), 267-288.

- van Eyden R., Difeto, M., Gupta, R., and Wohar, M.E. (2019). Oil price volatility and economic growth: Evidence from advanced economies using more than a century of data. *Applied Energy*, 233(C), 612-621.
- Wen, D., He, M., Wang, Y., and Zhang, Y. (2024). Forecasting crude oil market volatility: A comprehensive look at uncertainty variables. *International Journal of Forecasting*, 40(3), 1022-1041.
- Wen, J., Zhao, X-X., and Chang, C-P. (2021). The impact of extreme events on energy price risk. *Energy Economics*, 99(C), 105308.

FIGURES AND TABLES:

Figure 1. Data Plots



Note: IQR: Inter-Quantile Range; LEV: Leverage; SKEW: Skewness; KURT: Kurtosis; LTR: Lower Tail Risk; UTR: Upper Tail Risk; Shortages correspond to the various indexes of aggregate shortages, which is the sum of shortages due to energy, food, industry and labor sectors.

Table 1. Root Mean Square Errors (RMSEs)

Models	Horizon (h)				
	1	3	6	12	24
M1	2.8751	10.8909	25.0669	19.7698	20.7861
M2	0.6379	1.8267	3.4784	6.7759	12.8031
M3	0.6375	1.8252	3.4743	6.7608	12.7835
M4	0.6374	1.8249	3.4735	6.7595	12.7821

Note: M1 is the benchmark model with 12 lags of the inter-quantile range of oil log-returns; M2 is M1+oil log-returns moments (leverage, skewness, kurtosis, lower- and upper-tail risks); M3 is M2+aggregate shortages index; M4 is M2+ all four sub-indexes of energy, food, industry and labor.

Table 2. Forecast comparison tests p -values

Models	Horizon (h)				
	1	3	6	12	24
M2 versus M1	0.0031	0.0274	0.0278	0.0175	0.0388
M3 versus M1	0.0010	0.0012	0.0235	0.0025	0.0000
M4 versus M1	0.0010	0.0012	0.0234	0.0025	0.0000
M3 versus M2	0.0005	0.0005	0.0005	0.0001	0.0001
M4 versus M2	0.0003	0.0002	0.0002	0.0001	0.0001
M4 versus M3	0.0043	0.0008	0.0001	0.0128	0.0924
M4(I) versus M2	0.9644	0.9978	0.9900	0.4730	0.7615
M4(II) versus M2	0.1486	0.2635	0.6085	0.0439	0.0443
M4(III) versus M2	0.0002	0.0001	0.0001	0.0000	0.0001
M4(IV) versus M2	0.0007	0.0005	0.0003	0.0001	0.0001

Note: The entries in all rows, except M4 versus M3, are p -values of the Clark and West (2007) test of forecast comparison across two nested models, with the null being forecast equality, and the alternative is that the rival model outperforms the benchmark. For M4 versus M3, we report the corresponding p -values for the Diebold and Mariano (1995) test of forecast comparison across two non-nested models, with the same null and alternative hypotheses as above. M1 is the benchmark model with 12 lags of the inter-quantile range of oil log-returns; M2 is M1+oil log-returns moments (leverage, skewness, kurtosis, lower- and upper-tail risks); M3 is M2+aggregate shortages index; M4 is M2+all four sub-indexes of energy, food, industry and labor; M4(I) is M2+sub-index of energy shortages; M4(II) is M2+sub-index of food shortages; M4(III) is M2+sub-index of industry shortages, and; M4(IV) is M2+sub-index of labor shortages.