University of Pretoria
Department of Economics Working Paper Series

Oil Tail Risks and the Forecastability of the Realized Variance of Oil-Price: Evidence from Over 150 Years of Data

Afees A. Salisu
University of Ibadan
Christian Pierdzioch
Helmut Schmidt University
Rangan Gupta
University of Pretoria
Working Paper: 2021-46
June 2021

Department of Economics
University of Pretoria
0002, Pretoria
South Africa
Tel: +27 12 420 2413

Oil Tail Risks and the Forecastability of the Realized Variance of Oil-Price: Evidence from Over 150 Years of Data

Afees A. Salisu[a], Christian Pierdzioch[b], Rangan Gupta[c]

June 2021

Abstract

We examine the predictive value of tail risks of oil returns for the realized variance of oil returns using monthly data for the modern oil industry (1859:10-2020:10). The Conditional Autoregressive Value at Risk (CAViaR) framework is employed to generate the tail risks for both 1% and 5% VaRs across four variants of the CAViaR framework. We find evidence of both in-sample and out-of-sample predictability emanating from both 1% and 5% tail risks. Given the importance of real-time oil-price volatility forecasts, our results have important implications for investors and policymakers.

JEL classification: C22, C53, Q02

Keywords: Oil Tail Risks, Realized Variance of Oil-Price, Forecasting

[a] Centre for Econometric & Allied Research, University of Ibadan, Ibadan, Nigeria. Email address: adebare1@yahoo.com.

[b] Corresponding author. Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany; Email address: macroeconmoics@hsu-hh.de.

[c] Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa; Email address: rangan.gupta@up.ac.za.

1 Introduction

Accurate forecasting of the volatility of oil-price movements (in the following, oil-price volatility for short) has central implications for both investors and policymakers. This is because the recent financialization of the oil (and, in general, the overall commodity) market (Bampinas and Panagiotidis, 2017) has led to a substantial increase in the participation of key players such as hedge funds, pension funds, and insurance companies in the market, thus making oil a potentially profitable alternative investment in the asset-allocation decisions of financial institutions (Degiannakis and Filis, 2017). Moreover, given that volatility, when interpreted as uncertainty, is a core input to investment decisions and portfolio choices (Poon and Granger 2003), precise forecasts of oil-price volatility are of vital importance to participants in the oil market. Importantly, the second-moment movements in crude oil prices have been associated with a negative impact on global economic activity (van Eyden et al., 2019), implying that predicting the future path of oil-price volatility is of paramount importance to policymakers as well for the design of macroeconomic (monetary and fiscal) policies.

Against this backdrop, the objective of our paper is to analyse the forecastability of oil-price volatility, as measured by monthly realized variance (RV), based on the information content of oil market lower tail risks, where we use data for the West Texas Intermediate (WTI) oil price over the extended historical period from 1859:10 to 2020:10. Our data set basically covers the entire modern era of the petroleum industry (drilling of the first oil well was done in the United States (US) at Titusville, Pennsylvania in 1859). It should be noted that tail risk is the additional risk which, as has been commonly observed, fat-tailed asset returns exhibit relative to a normal distributions (Li and Rose, 2009). Intuitively, lower tail risks, i.e., those risks associated with extreme negative returns of the oil price, are likely to originate from massive reductions in aggregate demand (and also to some extent due to excessive supply) due to risks associated with rare disasters (Demirer et al., 2018), and as has been witnessed currently during the ongoing COVID-19 pandemic (Bouri et al., 2020), as well as geopolitical acts and threats (Cunado et al., 2020). This, in turn, is likely to result in an increase in the volatility of the oil-price via the well-established presence of the leverage effect in the oil market (Asai et al., 2020).

To the best of our knowledge, ours is the first paper to use estimates of tail risks of the oil market using over 150 years of monthly data for forecasting the RV of oil returns. In this regard, it should be emphasized that measuring volatility using RV, which in our case is captured by the sum of daily squared returns over a month (following Andersen and Bollerslev, 1998), provides an observable and unconditional metric of volatility (unlike, for example, generalized autoregressive conditional heteroscedasticity (GARCH) and stochastic volatility (SV) models), which is otherwise a latent process. The only related (working) paper that we could find is the work of Ellwanger (2017), who used (options-based) estimates of tail risks to forecast oil market volatility over the period 1986-2013. In this regard, by covering the longest possible data available on the evolution of the tail risks over the extended historical period from 1859 to 2020 (covering events such as, the US Civil War, the two World Wars, West coast gas famine, the Great Depression, the oil gluts in the early 1980s and multiple times over 2010-2019, global recession of 2001, the global financial crises of 2007-2009, and, of course, the current outbreak of the Coronavirus pandemic in 2020), and studying the corresponding forecastability of oil-price volatility, we avoid the issue of sample selection bias in our analysis. In the process, our paper, adds to the already existing large literature on the forecastability of oil-price volatility based on a wide array of models and macroeconomic, financial, behavioural, and climate patterns-related predictors (see, Gkillas et al., (2020), Bouri et al., (2021), and Salisu et al., (2021) for detailed reviews), by considering the role of tail risks in the oil market.

We organize the remainder of our paper as follows: Section 2 describes the data and methodologies, Section 3 discusses the forecasting results, and Section 4 concludes.

2 Data and Methodologies

We study daily and monthly WTI crude oil prices covering the period from 1859:09 to 2020:M10. The data are obtained from Global Financial Data[1], with the sample period being driven entirely by data availability at the time of compiling this paper. Both frequencies of the oil price data

[1] https://globalfinancialdata.com/.

are converted into log returns in percentage, i.e., the first-difference of the natural logarithm of the price multiplied by 100. Finally, RV_t is computed as the sum of daily squared returns over a month.[2] We conduct the forecasting analysis using $\tilde{RV}_t = \ln(1+RV_t)$ rather than modeling directly RV_t because the latter exhibits a very large peak at the very end of the sample period. However, our basic results, discussed in Section 3 below, continue to hold for RV_t as well (complete details of these results are available upon request from the authors). The forecasting model we use for predicting $\tilde{RV}_t$ is given by:

$$\tilde{RV}_{t+h} = c + \theta \tilde{RV}_t + \gamma TR_t + \eta_{t+h}, \tag{1}$$

where c, θ, γ are coefficients to be estimated, $\tilde{RV}_t$ is the (log) realized volatility at time t, $\tilde{RV}_{t+h}$ is the average value of $\tilde{RV}_t$ over t to $t+h$, where $h = 1, 3, 6, 12$ months ahead are the four forecasting horizons we consider in our paper, TR_t is the tail risk at time t , estimation of which we discuss below, and η_{t+h} is a disturbance term.[3] We construct the data such that the number of forecasts is the same for all four forecast horizons. Our benchmark model is nested in Equation (1) and obtained by setting $\gamma = 0$.

There are primarily two approaches for estimating tail risks: one relies on option-implied measures, while the other uses returns data (Kelly and Jiang, 2014). Historical options data are unavailable for the very long sample period that we study in this paper and, thus, we use the second approach, whereby we estimate tail risk using Value at Risk (VaR). To this end, we implement the conditional autoregressive VaR specifications as outlined by Engle and Manganelli (2004). In this regard, we consider the following four models: (i) the adaptive model; (ii) the symmetric slope model, (iii) the asymmetric slope model, and (iv) the indirect generalized autoregressive conditional heteroscedasticity (GARCH) model with an autoregressive mean.

[2]Oil price data is available at a monthly frequency for the period 1896-1899 and 1920-1976, and hence, we measure RV_t as the monthly squared log-returns over these time intervals. Also note that, since on the 20th of April, 2020, the WTI oil price became negative, we computed the daily returns for that date and the day after, i.e., 21st of April, 2020, by using the standard formula for the calculation of growth rate, instead of log-returns since, understandably, the natural logarithmic value of a negative number is undefined.

[3]Following Corsi (2009), we initially specified a heterogenous autoregressive (HAR)-RV model, which included the one lag each of the averages of $\tilde{RV}_t$ over a quarter and a year, but the corresponding coefficients were insignificant, and hence we used the specification in Equation (1).

Given an observable vector of returns, $\{y_t\}_{t=1}^T$, we let ϕ denote the associated probability with VaR, X_t the vector of observable variables at time t, and α_ϕ the p-vector of unknown parameters, and for ease of notation we write $f_t(\alpha) \equiv f_t(X_{t-1}, \alpha_\phi)$ to denote the time-t ϕ-quantile of the distribution of y_t formed at time $t-1$. We express a generic representation of the CAViaR as

$$f_t(\alpha) = \alpha_0 + \sum_{i=1}^{k} \alpha_i f_{t-i}(\alpha) + \sum_{j=1}^{r} \alpha_j l(X_{t-j}), \tag{2}$$

where $p = k + r + 1$ denotes the dimension of α and l is a function of a finite number of lagged values of observables (X_{t-j}). Note that we suppress the ϕ subscript from α_ϕ in Equation (2) for notational convenience. The quantile changes "smoothly" over time via the autoregressive terms, $\alpha_i f_{t-i}(\alpha)$, $i = 1, ..., k$, while $l(X_{t-j})$ is used to link $f_t(\alpha)$ to observable variables that belong to the information set. Consequently, the four variants of the CAViaR estimated are derived from Equation (2) as follows:

Adaptive:

$$f_t(\alpha_1) = f_{t-1}(\alpha_1) + \alpha_1 \left(\{1 + \exp\left(G[y_{t-1} - f_{t-1}(\alpha_1)]\right)\}^{-1} - \phi \right). \tag{3}$$

Symmetric absolute value:

$$f_t(\alpha) = \alpha_1 + \alpha_2 f_{t-1}(\alpha) + \alpha_3 |y_{t-1}|. \tag{4}$$

Asymmetric slope:

$$f_t(\alpha) = \alpha_1 + \alpha_2 f_{t-1}(\alpha) + \alpha_3 (y_{t-1})^+ + \alpha_4 (y_{t-1})^-. \tag{5}$$

Indirect GARCH (1,1):

$$f_t(\alpha) = \left(\alpha_1 + \alpha_2 f_{t-1}^2(\alpha) + \alpha_3 y_{t-1}^2\right)^{0.5}. \tag{6}$$

It is understood that the last term in Equation (3) converges to $\alpha_1 \left[I(y_{t-1} \leq f_{t-1}(\alpha_1)) - \phi \right]$ almost surely if $G \to \infty$, with $I(.)$ representing the indicator function, while G is some positive finite number which makes the model a smoothed version of a step function.[4]

[4]The computational procedure of these models is detailed in Engle and Manganelli (2004) and we refer a technically-minded readers their paper for further exposition on the tail risk estimation.

3 Empirical Results

We estimate all the four CAViaR variants and produce tail risks for both 1% and 5% VaRs. We then apply the Dynamic Quantile test (DQ) test, the Regression Quantile (RQ) statistic, and the %Hits statistic to identify the model that best fits the data (see also Salisu et al., (forthcoming)). We expect the %Hits to be 1% for 1% VaR and 5% for 5% VaR, while the DQ test statistic is not expected to be significant, and the parameters are expected to minimize the RQ loss function, so the smaller the RQ statistic, the better. The DQ test, however, takes prominence over the %Hits and RQ statistics. In cases where more than one tail risk is statistically insignificant in terms of the DQ test, we consider both the %Hits and the RQ test statistic to determine the model with the best fit. In the same vein, where all the tail risks are statistically significant, then the %Hits and the RQ test statistic becomes a major criteria except where some distinctions can still be made with the significant DQ test statistic.

– Please include Table 1 about here. –

As evidenced by the estimation results for the CAViaR models given in Table 1, the asymmetric model fits the oil returns data best in terms of the statistics mentioned above in capturing 1% and 5% tail risks, with us using both in our forecasting exercise (see also Salisu et al., (forthcoming)).

A strong test of any predictive model (in terms of the econometric methodologies and predictors employed) is its out-of-sample performance, and for this reason we focus on the predictive analysis from an out-of-sample perspective. For the sake of completeness, however, the full-sample estimation results for Equation (1) for $h = 1, 3, 6, 12$ months, with heteroscedasticity and autocorrelation adjusted standard errors, are reported in Table 2. In line with the intuition that returns at lower tails serve as bad news and enhance volatility via leverage, both 1% and 5% tail risks consistently increase $\tilde{RV}$ in a statistically significant manner for all four forecasting horizons.

– Table 2 about here. –

In order to conduct the out-of-sample forecasting exercise, we first determine the in- and out-of-sample split, based on the Bai and Perron (2003) tests of multiple structural breaks applied to Equation (1) under 1% and 5% tail risks for $h = 1, 3, 6, 12$ months. The tests detected between 3 to 5 breaks among the 8 cases, with the earliest structural change detected under $h = 6$ at 1888:12 for 1% tail risks.[5] Given this, our in-sample covers the period from 1859:10 to 1888:12, and the out-of-sample period, therefore, starts in 1888:12. We estimate our forecasting models with and without the tail risks recursively over the out-of-sample period.

– Table 3 about here. –

Table 3 presents the ratio of root-mean-squared-forecasting errors (RMSFE), with both 1% and 5% tail risks considered separately. The RMSFE ratio is computed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. The model without tail risk is the benchmark model, and the model extended to include tail risk is the rival model. As can be seen, the relative MSFEs are less than one under both 1% and 5% tail risks cases at all horizons. The results show that including information on tail risks in the forecasting model produces lower forecast errors associated with $\tilde{RV}$ compared to the case when this information is excluded from the model.

– Tables 4 and 5 about here. –

Table 4 reports results of the Clark and West (2007) test for an equal mean-squared prediction error (MSPE). Corroborating the the evidence of in-sample predictability, the more powerful test based on-out-of-sample forecasting picks up statistically significant prediction of $\tilde{RV}$ based on tail risks at all horizons. In Table 5, we report as a robustness check results of the Clark-West test when we use a model that features a leverage effect as the benchmark model. The leverage effect is equal to returns when the latter are negative, and zero otherwise. The leverage effect, thereby, is a simple metric of downside risk and, thus, an alternative metric of bad news. The results show that the model that features tail risk as a predictor has a better out-of-sample forecasting performance than the leverage model for both the 1% and the 5% VaR.

[5]Complete details of the structural break dates are available upon request from the authors.

4 Concluding Remarks

We have analysed the predictive value of tail risks of oil returns for the realized oil-price variance over 150 years of monthly data, i.e., 1859:10-2020:10. We have estimated tail risks based on the popular Conditional Autoregressive Value at Risk (CAViaR) model, where we have considered four variants of this model for both 1% and 5% VaRs. Upon using the “best" fitting model selected for each VaR, our analysis of in-sample predictability has shown that tail risks tends to increase realized variance in a statistically significant manner at all four forecasting horizons that we have studied. We have then turned to an out-of-sample forecasting analysis, which is well-established as a more powerful test of predictability, where we have used structural break tests to determine the split of the data into a training and an out-of-sample period. Given this set-up, we have found that statistically significant forecasting gains from both 1% and 5% tail risks for realized variance, also when we compared the predicrtive performance of tail risks with that of a leverage effect.

In any event, given the importance of real-time forecasts of oil-price volatility for both investors and policy authorities, our results have important implications for these two groups of economic agents. Specifically speaking, our findings can be used by policymakers to obtain information on the future path of oil-price volatility due to oil returns tail risks in the short- to medium-runs. This knowledge, in turn, may be useful to predict economic recessions gathering steam, given that oil-price volatility is known to negatively impact real economic activity, and two implement appropriate macroeconomic policies that help to prevent the macroeconomy to fall into a recession. Moreover, with volatility being a key input in portfolio decisions, the forecastability of oil-price volatility due to oil tail risks may be of vital importance to oil investors, especially over a semi-yearly investment horizon.

As an extension of this study, it would be interesting to investigate the role of tail risks in predicting volatility of other commodity markets, especially from a historical perspective, contingent on the availability of data.

References

Andersen T.G., and Bollerslev T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. International Economic Review, 39(4), 885−905.

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.

Bai, J., and Perron, P. (2003). Computation and analysis of multiple structural change models. Journal of Applied Econometrics, 18, 1−22.

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., Gkillas, K., Gupta, R., and Pierdzioch, C. (2020). Infectious diseases, market uncertainty and realized volatility of oil. Energies, 13(16), 4090.

Bouri, E., Gupta, R., Pierdzioch, C., and Salisu, A.A. (2021). El Niño and forecastability of oil-price realized volatility. Theoretical and Applied Climatology, 144, 1173−1180.

Corsi, F. (2009) A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics, 7, 174−196.

Clark, T.E., and West, K.D. (2007). Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics, 138(1), 291−311.

Cunado, J., Gupta, R., Lau, C.K.M., and Sheng, X. (2020). Time-varying impact of geopolitical risks on oil prices. Defence and Peace Economics, 31(6), 692−706.

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, 28−49.

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.

Ellwanger, R. (2017). On the tail risk premium in the oil market. Staff Working Papers 17-46, Bank of Canada.

Engle, R.F. and Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by regression quantiles. Journal of Business and Economic Statistics, 22(4), 367−381.

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.

Kelly, B., and Jiang, H. (2014). Tail risk and asset prices. Review of Financial Studies, 27(10), 2841−2871.

Li, X-M., and Rose, L.C. (2009). The tail risk of emerging stock markets. Emerging Markets Review, 10, 242−256.

Poon, S-H, and Granger, C.W.J. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41(2), 478−539.

Salisu, A.A., Pierdzioch, C., and Gupta, R. (Forthcoming). Geopolitical risk and forecastability of tail risk in the oil market: Evidence from over a century of monthly data. Energy.

Salisu, A.A., Gupta, R., Bouri, E., and Ji, Q. (2021). Mixed-frequency forecasting of crude oil volatility based on the information content of global economic conditions. Journal of Forecasting. DOI: https://doi.org/10.1002/for.2800.

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, 612−621.

Table 1: CAViaR Models

Parameter/statistic VaR level	Asymmetric slope 1% VaR	Asymmetric slope 5% VaR
α_1	1.3322***	0.2309***
s.e.	0.3324	0.0944
α_2	0.7977***	0.7416***
s.e.	0.0254	0.0377
α_3	0.2415***	0.2059***
s.e.	0.0675	0.0687
α_4	0.7180***	0.7376***
s.e.	0.1603	0.0892
RQ	357.0882	1235.0178
%Hits in sample	1.0101	5.0505
%Hits out-of sample	1.6000	10.4000
DQ in-sample (p-values)	0.1135	0.0001
DQ out-of-sample (p-values)	0.9625	0.0011

Note: s.e. is standard error and *** denotes significance at the 1% level. The table reports the estimation results for the CAViaR model that best "fits" the oil returns series. 75% (25%) of the data is used for in-sample (out-of-sample).

Table 2: Full-Sample Results

Forecast horizon	Estimate (1 %)	t-statistic	Estimate (5 %)	t-statistic
h=1	0.0257***	4.7166	0.0280***	3.5655
h=3	0.0354***	4.2481	0.0393***	3.9762
h=6	0.0385***	5.4766	0.0432***	5.0202
h=12	0.0356***	4.9428	0.0386***	5.2421

Note: Estimates of the the coefficeint γ in Equation (1), at forecasting horizons $h = 1, 3, 6, 12$ months ahead. Robust t-statistics are also reported; *** and significance at the 1% level.

Table 3: Ratio of Root-Mean-Squared-Forecasting Errors

VaR	h=1	h=3	h=6	h=12
Risk 1%	0.9737	0.9450	0.9312	0.9387
Risk 5%	0.9789	0.9542	0.9419	0.9514

Note: The ratio of the root-mean-squared-forecasting errors (RMSFE) is computed by dividing the RMSFE of the rival model by the RMSFE of the benchmark model. The model without tail risk is the benchmark model, and the model extended to include tail risk is the rival model. The parameter h denotes the forecast horizon (in months).

Table 4: Baseline Out-of-Sample Test Results

VaR	h=1	h=3	h=6	h=12
Risk 1%	6.5762	7.1063	6.5829	5.2419
Risk 5%	6.3948	6.8306	6.3568	4.7865

Note: t-statistics of the Clark-West tests for an equal mean-squared prediction error are based on robust standard errors. Critical values (one-sided test) are 1.282 (10%) and 1.645 (5%). The model without tail risk is the benchmark model, and the model extended to include tail risk is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter h denotes the forecast horizon (in months). The p-values are based on robust standard errors.

Table 5: Robustness Check of Out-of-Sample Test Results

VaR	h=1	h=3	h=6	h=12
Risk 1%	6.7486	7.8601	6.5916	5.1018
Risk 5%	6.3519	7.2910	6.3346	4.7280

Note: t-statistics of the Clark-West tests for an equal mean-squared prediction error (MSPE) are based on robust standard errors. Critical values (one-sided test) are 1.282 (10%) and 1.645 (5%). The model without tail risk but with a leverage effect is the benchmark model, and the model extended to include tail risk is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter h denotes the forecast horizon (in months). The p-values are based on robust standard errors.