Università degli Studi di Padova

Padua Research Archive - Institutional Repository

The dynamic impact of uncertainty in causing and forecasting the distribution of oil returns and risk

Original Citation:

Availability:

This version is available at: 11577/3278342 since: 2020-01-04T17:05:50Z

Publisher:

Elsevier B.V.

Published version:

DOI: 10.1016/j.physa.2018.05.061

Terms of use:

Open Access

This article is made available under terms and conditions applicable to Open Access Guidelines, as described at <http://www.unipd.it/download/file/fid/55401> (Italian only)

(Article begins on next page)

The Dynamic Impact of Uncertainty in Causing and Forecasting the Distribution of Oil Returns and Risk*

G. Bonaccolto¹, M. Caporin², and R. Gupta³

¹University of Enna “Kore”, viale delle Olimpiadi, 94100 Enna, Italy. *Email : bonaccolto@unikore.it*

²Department of Statistical Sciences, University of Padova, via C. Battisti 241, 35121 Padova, Italy.

Email : massimiliano.caporin@unipd.it

³Department of Economics, University of Pretoria, 0002 Pretoria, South Africa.

Email : Rangan.Gupta@up.ac.za

Pre-print version of the paper published in Physica A

10.1016/j.physa.2018.05.061

Abstract

The aim of this study is to analyze the relevance of recently developed news-based measures of economic policy and equity market uncertainty in causing and predicting the conditional quantiles of crude oil returns and risk. For this purpose, we studied both the causality relationships in quantiles through a non-parametric testing method and, building on a collection of quantiles forecasts, we estimated the conditional density of oil returns and volatility, the out-of-sample performance of which was evaluated by using suitable tests. A dynamic analysis shows that the uncertainty indexes are not always relevant in causing and forecasting oil movements. Nevertheless, the informative content of the uncertainty indexes turns out to be relevant during periods of market distress, when the role of oil risk is the predominant interest, with heterogeneous effects over the different quantiles levels.

Keywords: Granger Causality in Quantiles; Quantile Regression; Forecast of Oil Distribution; Forecast Evaluation.

JEL codes: C58, C32, C53, Q02, Q35.

1 Introduction

Following the seminal work of Hamilton (1983), a large literature connects movements in oil returns and its volatility with recessions and inflationary episodes in the US economy (see, e.g., Elder and Serletis

*We would like to thank three anonymous referees for many helpful comments. However, any remaining errors are solely ours.

(2010), Kang and Ratti (2013b,a) and Antonakakis et al. (2014) for detailed reviews). Hamilton (2008) indicates that nine out of ten recessions in the US since World War II have been preceded by an increase in oil prices. Interestingly, Hamilton (2009) even goes as far as arguing that a large proportion of the recent downturn in the US GDP during the ‘Great Recession’ can also be attributed to the oil price shock in the period 2007—2008.

Commodity markets, just like asset prices, are known to be functions of the state of the economy (Bekiros et al., 2015). In this regard, a recently growing literature emphasizes the role of economic policy uncertainty on business cycles (see, e.g., Bloom (2009), Colombo (2013), Jones and Olson (2013), Mumtaz and Zanetti (2013), Karnizova and Li (2014) and Jurado et al. (2015) for detailed reviews) which in turn affects oil-price movements (see, e.g., Kang and Ratti (2013b,a), Antonakakis et al. (2014) and Aloui et al. (2016)). On the demand side, uncertainties can also drive economic concerns on the part of consumers, thus affecting the level of consumption growth in the economy. Therefore, considering the suggestion by Bernanke (2016) that both oil and the stock markets tend to move together, as they both react to a common factor reflecting global aggregate demand, one obvious channel that links uncertainty to oil market movements is its potential effect on growth expectations for both output and consumption. Equity-market uncertainty also feeds into oil price movements because, as Bloom (2009)’s firm-based theoretical framework notes, equity-market uncertainty affects hiring and investment and, hence, the production decisions of firms. In this regard, the empirical evidence relating to oil price movements and stock market volatility can be found in Kang and Ratti (2013b,a). While these channels are likely to cause economic uncertainties to affect oil market movements at lower frequencies, high frequency (for example, daily) impacts can originate from other possible channels.

For instance, uncertainty can affect oil return dynamics via its contribution to jump risk in oil prices. There is growing evidence suggesting that jumps account for a large part of the variation in crude oil prices and that a substantial part of the risk premium in oil derivatives prices is due to jumps (see, e.g., Larsson and Nossman (2011), Christoffersen et al. (2016) and Baum and Zerilli (2016)). Therefore, it could be argued that economic uncertainties contribute to the presence of jumps in oil prices, which in turn drive return and volatility dynamics in the oil market. Hence, even though the oil market is one of the most deep and liquid markets, complemented by a set of oil-related derivative instruments, if the jump risks emanating from economic policy and equity market uncertainties cannot be effectively hedged through the variety of available instruments, then these uncertainties are likely to influence the oil market. Alternatively, given that these uncertainties have been shown to affect equity markets (see, e.g., Chuliá et al. (2017) for a detailed literature review in this regard), investors might in fact move funds to diversify their portfolios by investing in the commodity market, including in oil, in the hope of hedging portfolio risks (Andreasson et al., 2016). These sudden movements of investments into the oil market could also impact on returns and the volatility of crude oil. As indicated by Ji and Guo (2015b,a), uncertainty tends to move the oil market through a behavioural channel as well, i.e.,

by affecting market participants' psychological expectations.

Against this backdrop, the objective of this paper is to analyze the role of recently developed news-based measures of economic policy uncertainty (EPU) and equity market uncertainty (EMU) by Baker et al. (2013) in causing in the Granger (1969) sense and forecasting oil returns and their risk. Given the possibility that the oil market is also likely to drive these uncertainty measures—see, e.g., Kang and Ratti (2013b,a) and Antonakakis et al. (2014)—we employed a modified bivariate quantile causality-based model, as developed by Balcilar et al. (2016a,b,c); notably, it combines the causality in quantile test of Jeong et al. (2012) with the k -th order nonparametric Granger causality test of Nishiyama et al. (2011). By resorting to this quantile-based analysis of causality, we evaluated the impact of news-based measures of uncertainty on both the returns and the risk of oil across a collection of quantiles. Testing for causality in risk allows us to shed some light on the volatility spillover phenomenon, since at times, the simple causality in returns series may not exist, but there may be significant relationships at higher moments. Notably, using quantile-based methods allows us to analyze the causality structure depending on the volatility state (high versus low).

Balcilar et al. (2016a) developed the framework we employed in this study to analyze the causality relationships running from EPU and EMU to oil returns and risk. They concluded that, for oil returns, EPU and EMU have strong predictive power over the entire distribution, barring the regions around the median, but for risk the predictability virtually covers the entire distribution, with some exceptions in the tails.¹ We extend the paper by Balcilar et al. (2016a) in various important ways. First, we made use of a rolling window procedure, by which we provided a time-varying approach to the in-sample quantile causality for both oil returns and risk. This is important, given that we detected structural breaks in the estimated conditional distributions over time; therefore, the full-sample quantile causality, as in Balcilar et al. (2016a), could possibly be misleading. Indeed, in contrast to Balcilar et al. (2016a), the empirical findings arising from the rolling window analysis suggest that EPU and EMU are not always relevant drivers in causing and forecasting the conditional quantiles of oil returns and risk, but are only so during particular periods, especially ones of market distress. Furthermore, our results show evidence of stronger relationships between the two uncertainty indexes and oil risk, with respect to oil returns. This important finding is consistent with the fact that EPU and EMU are uncertainty indexes and, therefore, are directly connected to oil risk, quantified by its volatility, which itself is a measure of dispersion, or uncertainty.

Second, starting with Balcilar et al. (2016a,b,c), in which causality in risk is implemented by using squared returns, we went further by directly considering the realized volatility of oil. We found that the two approaches provide similar implications at central quantiles levels, but differ for the extreme quantiles, thus challenging the use of squared oil returns for the analyses of volatility causation between

¹Note that Balcilar et al. (2016a) focus just on the in-sample analysis. In contrast, we provide here an extensive analysis both in-sample and out-of-sample.

oil and uncertainty measures.

Finally, we accompanied the causality exercise with a forecasting analysis. In contrast to Bekiros et al. (2015), where the authors focus only on a point forecast of oil returns, we were able to analyze the density forecast for both oil returns and risk. In particular, we made use of the causality detected in designing the quantile regression models (Koenker and Bassett, 1978), and in doing that we adjusted the original estimated quantiles to guarantee their coherence; that is, their increasing monotonicity in $\tau \in (0, 1)$. Indeed, the approach introduced by Koenker and Bassett (1978) allows the estimation of the single quantiles individually. As a consequence, when the analysis focuses on many quantiles, they may cross each other at specific quantile levels. Then, from a large collection of corrected quantiles, we built the conditional density of oil returns and risk through a non-parametric kernel-based method. This again is more informative than the point forecasts, since we were able to understand the role of EPU and EMU in forecasting oil movements in different phases (bearish, normal and bullish) of the market. Moreover, extending the analysis to the entire conditional distribution is of relevant importance in evaluating the uncertainty associated with the single point estimates and forecasts. Additionally, we did not restrict our attention to the in-sample analysis, but also evaluated the out-of-sample performance of the method we proposed by using various suitable tests. Both in- and out-of-sample results show evidence of the heterogeneous effects of EPU and EMU over the different regions of the oil movements distribution.

We also compared our approach with a competitive method; that is, an exponential generalized autoregressive conditional heteroskedasticity model (EGARCH), which gave stronger evidence of the causality impact of EPU and EMU when we focused on oil returns, but a weaker evidence when we focused on oil variance. Furthermore, we directly compared our approach with the EGARCH model in terms of predictive accuracy using various testing methods. Notably, our approach overperformed the EGARCH model in the majority of cases. Therefore, the methods we propose could provide useful insights for many of the decision makers in several areas of economics and finance, such as risk management, pricing and trading strategies, when the focus is placed on instruments depending on oil returns and risk.

The rest of the paper is organized as follows: Section 2 presents the details of the methodologies pursued, while we describe the empirical set-up in Section 3. Section 4 presents the results and Section 5 concludes with an economic discussion of the results obtained.

2 Data and methodology

In the empirical analysis, we made use of four series: oil prices, the realized volatility of oil future prices and the two uncertainty indexes—EPU and EMU. Note that, instead of using the VIX², a

²Often referred to as the fear index or the fear gauge, it represents one measure of the market's expectation of stock-market volatility over the next 30 day period.

popular measure of the implied volatility of S&P500 index options, we use the news-based measure of EMU index to ensure that both our measures of uncertainty (i.e., EPU and EMU) are derived in a similar method (i.e., news articles-based) and hence, the results, in terms of their relationship with oil are comparable. Also, since we are extending the work of Balcilar et al. (2016a), we too use the same measures of uncertainties as these authors, though ideally we should be using some measure of global uncertainty, as oil is a global market. But then again, we are analyzing the WTI oil price movements, which are likely to be affected more due to US-based uncertainties, besides the fact that, there is no global measure of uncertainty available at daily frequency. Moreover, as indicated by Ajmi et al. (2015), US uncertainty leads uncertainties of other economies around the world, thus making it an important source and a relevant proxy of global uncertainty. For oil prices and both EMU and EPU, we considered a sample period starting on January 2, 1986 and ending on April 23, 2015, for a total of 7646 days. In contrast, we estimated the daily realized volatility of oil future prices from five-minutes data from January 2, 2007 to April 23, 2015.³ Appendix B provides a detailed description of the series and of their properties, as well as a discussion concerning the estimator we used to recover the daily realized volatility.

Let $\{y_t\}_{t \in T}$ be the time series of oil returns computed as $y_t = \log(oil_t) - \log(oil_{t-1})$, where oil_t is the oil spot price at time t . We also denote by $\{x_{1,t}\}_{t \in T}$ and $\{x_{2,t}\}_{t \in T}$ the logarithm of the Economic Policy Uncertainty (EPU) and the logarithm of the Equity Market Uncertainty (EMU), respectively. Finally, let bpv_t be the estimated daily realized volatility of oil future prices.

Next, we move to a brief description of the methodological tools we used in our empirical analysis.

2.1 Causality in quantiles

We first consider the methods we used to shed light on the causality relations between oil return dynamics and the two uncertainty indexes in a bivariate framework. For simplicity of notation, we used x_t in place of $x_{1,t}$ or $x_{2,t}$ when we studied the causality implications of EPU or EMU on y_t .

Let $F_{y_t|Z_{t-1}}(y_t|Z_{t-1})$ and $F_{y_t|Y_{t-1}}(y_t|Y_{t-1})$ be the distributions of y_t conditional on Z_{t-1} and Y_{t-1} respectively, where $Y_{t-1} \equiv (y_{t-1}, \dots, y_{t-p})$ and $Z_{t-1} \equiv (y_{t-1}, \dots, y_{t-p}, x_{t-1}, \dots, x_{t-q})$, for $(p, q) > 1$. Besides, $Q_\tau(Z_{t-1}) \equiv Q_\tau(y_t|Z_{t-1})$ and $Q_\tau(Y_{t-1}) \equiv Q_\tau(y_t|Y_{t-1})$ are the τ -th quantiles of y_t conditional on Z_{t-1} and Y_{t-1} , respectively, for $\tau \in (0, 1)$.

In studying the Granger causality in quantiles, we followed Jeong et al. (2012): x_t does not cause the τ -th quantile of y_t with respect to Z_{t-1} if $Q_\tau(Z_{t-1}) = Q_\tau(Y_{t-1})$. In contrast, x_t is a *prima facie* cause in the τ -th quantile of y_t with respect to Z_{t-1} if $Q_\tau(Z_{t-1}) \neq Q_\tau(Y_{t-1})$. We evaluated the null hypothesis of no causality (H_0) using the test statistic proposed by Jeong et al. (2012), which we denote as $\hat{J}_T^*$ throughout the paper. Notably, $\hat{J}_T^*$ is asymptotically distributed as $\mathcal{N}(0, 1)$.⁴

³The series of oil prices was recovered from Thomson Reuters Datastream. The data and the details about EPU and EMU are available at <http://www.policyuncertainty.com/>. High-frequency data were obtained from TickData.com.

⁴Other contributions that use $\hat{J}_T^*$ are, for instance, Balcilar et al. (2016a,b,c).

More recently, Balcilar et al. (2016a,b,c) extended the approach introduced by Jeong et al. (2012) to the second moment of y_t . This novel approach aims at testing for the presence of quantile causality when considering the density of the risk or of the dispersion characterizing the variable y_t . In particular, Balcilar et al. (2016a,b,c) studied the causality impact on the τ -th quantile of oil's variance by replacing y_t by y_t^2 , preserving the asymptotic distribution of $\hat{J}_T^*$. Balcilar et al. (2016a,b,c) proved the validity of this approach, starting with the findings in Nishiyama et al. (2011).

Nevertheless, using squared values of the dependent variable, that is, y_t^2 does not provide a direct focus on causation from the uncertainty measures to the quantiles of oil risk. In fact, if there is some quantile causality in the first-order moment, this might distort the outcomes of quantile causality in the second-order moment, as pointed out by Nishiyama et al. (2011) and Balcilar et al. (2016a,b,c). More clearly, the test involving squares is used for causality up to the second-order moment, and can be interpreted as a test for causality in risk only if there is no quantile causality for y_t . We then extended the existing literature by directly implementing the Jeong et al. (2012) test to realized measures of volatility. In this way, we could evaluate the fitness of the method used in Balcilar et al. (2016a,b,c), where the squared returns of the variables of interest were taken as a proxy of their second moment, for analyzing the causality relationships in quantiles. By replacing y_t or y_t^2 with bpv_t , that is, oil's bipower variation (Barndorff-Nielsen and Shephard, 2004), in $\hat{J}_T^*$, we tested the causality impact of EPU and EMU on the τ -th quantile of the y_t 's realized variance.

2.2 Quantiles and density forecasting

In this section we summarize the methods we used for studying the forecasting implications of EMU and EPU on y_t , y_t^2 and bpv_t . As mentioned in the introduction, one of the research questions deals with evaluating the potentially different impact of the two uncertainty indexes. A forecasting exercise allows for a direct comparison of EPU and EMU, which can be jointly introduced in a single model for testing their statistical and forecasting impact. In particular, we aimed to forecast both the conditional quantiles and distributions of y_t , y_t^2 and bpv_t taking into account the information associated with EPU and EMU.

The testing approach discussed in Section 2.1 allows for the detection of causation, but it does not provide guidance as to the specification of any particular functional form. We decided to adopt the simplest specification for the quantiles, that is, a linear one, resorting to the quantile regression introduced by Koenker and Bassett (1978). We stress that the model is linear only at the quantiles, and we did not impose a distributional assumption on the residuals. We then used a parametric model for the quantiles and recovered the density forecasts in a non-parametric way.

Given $\mathcal{W}_{t-1} \equiv (y_{t-1}, \dots, y_{t-p}, x_{1,t-1}, \dots, x_{1,t-q}, x_{2,t-1}, \dots, x_{2,t-r})$, for $(p, q, r) > 1$, the first step consisted in forecasting the conditional quantiles of y_t , by estimating a quantile regression model (Koenker

and Bassett, 1978), which led to the following conditional quantile function:⁵

$$\begin{aligned} Q_\tau(y_t|\mathcal{W}_{t-1}) &= \alpha_0(\tau) + \beta_1(\tau)y_{t-1} + \dots + \beta_p(\tau)y_{t-p} + \delta_1(\tau)x_{1,t-1} \\ &+ \dots + \delta_q(\tau)x_{1,t-q} + \lambda_1(\tau)x_{2,t-1} + \dots + \lambda_r(\tau)x_{2,t-r}. \end{aligned} \quad (1)$$

We then obtained from (1) the forecast with a horizon of one period ahead as follows:

$$\begin{aligned} \widehat{Q}_\tau(y_{t+1}|\mathcal{W}_t) &= \widehat{\alpha}_0(\tau) + \widehat{\beta}_1(\tau)y_t + \dots + \widehat{\beta}_p(\tau)y_{t-p+1} + \\ &+ \widehat{\delta}_1(\tau)x_{1,t} + \dots + \widehat{\delta}_q(\tau)x_{1,t-q+1} + \widehat{\lambda}_1(\tau)x_{2,t} + \dots + \widehat{\lambda}_r(\tau)x_{2,t-r+1}. \end{aligned} \quad (2)$$

The standard quantile regression approach allows estimating individual quantiles, but it does not guarantee their coherence, i.e. their increasing monotonicity in $\tau \in (0, 1)$. For instance, it might have occurred that the predicted 95-th percentile of the response variable was lower than the 90-th percentile. If quantiles cross, corrections must be applied to obtain a valid conditional distribution of volatility. For instance, to cope with the crossing problem, Koenker (1984) applied parallel quantile planes, whereas Bondell et al. (2010) estimated the quantile regression coefficients with a constrained optimization method. We followed a different approach, that is, the one proposed by Zhao (2011). Given a collection of ϑ predicted conditional quantiles $(\widehat{Q}_{\tau_1}(y_{t+1}|\mathcal{W}_t), \dots, \widehat{Q}_{\tau_\vartheta}(y_{t+1}|\mathcal{W}_t))$, for $0 < \tau_j < \tau_{j+1} < 1$, $j = 1, \dots, \vartheta - 1$, we first rearranged them into ascending order, by making use of the quantile bootstrap method proposed by Chernozhukov et al. (2010). We denote the rearranged quantiles as $Q_{\tau_1}^*(y_{t+1}|\mathcal{W}_t), \dots, Q_{\tau_\vartheta}^*(y_{t+1}|\mathcal{W}_t)$.

We then estimated the entire conditional distribution with a nonparametric kernel method. The predicted density is obtained as follows:

$$\widehat{f}_{y_{t+1}|\mathcal{W}_t}(y^*|\mathcal{W}_t) = \frac{1}{\vartheta h_\vartheta} \sum_{i=1}^{\vartheta} K_e \left(\frac{y^* - Q_{\tau_i}^*(y_{t+1}|\mathcal{W}_t)}{h_\vartheta} \right), \quad (3)$$

where y^* are evenly interpolated points that generate the support of the estimated distribution, h_ϑ is the bandwidth and $K_e(\cdot)$ is the kernel function. Following Gaglianone and Lima (2012), we used the Epanechnikov kernel as $K_e(\cdot)$.

By implementing the same method described above, we also forecasted the quantiles and densities of y_t^2 , which we respectively denoted as $\widehat{Q}_\tau(y_{t+1}^2|\mathcal{W}_{2,t})$ and $\widehat{f}_{y_{t+1}^2|\mathcal{W}_{2,t}}(y^{2*}|\mathcal{W}_{2,t})$, by replacing y_t in (1)–(3) with y_t^2 , given $\mathcal{W}_{2,t-1} \equiv (y_{t-1}^2, \dots, y_{t-p}^2, x_{1,t-1}, \dots, x_{1,t-q}, x_{2,t-1}, \dots, x_{2,t-r})$. Likewise, we forecasted the quantiles and densities of bpv_t ($\widehat{Q}_\tau(bpv_{t+1}|\mathcal{W}_{3,t})$ and $\widehat{f}_{bpv_{t+1}|\mathcal{W}_{3,t}}(bpv^*|\mathcal{W}_{3,t})$) by replacing y_t in (1)–(3) with bpv_t , given $\mathcal{W}_{3,t-1} \equiv (bpv_{t-1}, \dots, bpv_{t-p}, x_{1,t-1}, \dots, x_{1,t-q}, x_{2,t-1}, \dots, x_{2,t-r})$.

We computed the coefficients' standard errors using the bootstrap method (Efron, 1979), the ad-

⁵We followed a standard practice, reporting the expected conditional quantiles where the intercept includes a parameter and the quantile of the error term; we refer the reader to Koenker (2005) for details.

vantages of which are well-known: it assumes no particular distribution of the errors, it is not based on asymptotic model properties and it is available regardless of the complexity of the statistic of interest. Among all the available bootstrapping methods, we made use of the xy -pair method (Kocherginsky, 2003), which provides various advantages for quantile regression problems (Davino et al., 2014).

We evaluated the predictive accuracy of the methods described above by implementing the tests proposed by Berkowitz (2001), Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011), where the variance of the differences of the losses/scores provided by the competitive models was computed using the heteroskedasticity and autocorrelation consistent (HAC) estimator. We give the main details of the tests in Appendix A.

3 Empirical set-up: dynamic analysis and rolling window procedure

As noticed by Balcilar et al. (2016a), the relationships between y_t or y_t^2 and the uncertainty indexes are not stable over time. The authors implemented the Bai and Perron (2003)'s test, detecting the presence of multiple structural breaks in oil return series for the EPU- and EMU-based VARs.⁶ Here, we implemented the DQ and the SQ tests introduced by Qu (2008), which best fit our analysis, as they reveal structural changes in regression quantiles with unknown timing. Following Tillmann and Wolters (2015), whose study focuses on US inflation persistence, we proceeded in two stages. First, we used the DQ test to capture possible changes in the entire conditional distribution of the response variable. Given that we did not have any prior information about the part of the conditional distribution affected by breaks, we took into account a broad number of quantiles levels, that is, $\tau = \{0.05, 0.1, 0.15, \dots, 0.95\}$. Second, we implemented the SQ test at the dates where the null hypothesis of the DQ test was rejected at the level of 0.01. For simplicity, we implemented the SQ test at three quantile levels, that is, $\tau = \{0.1, 0.5, 0.9\}$. Notably, the DQ and the SQ tests highlighted the presence of several structural breaks for different regions of the conditional distributions of y_t , y_t^2 and bpv_t . As a result, the conclusions drawn from the full sample analysis might be misleading also in the quantile regression framework.⁷

In capturing the dynamics in the relations between the variables of interest, we differed from Balcilar et al. (2016a) by implementing a rolling window procedure for causality testing, model estimating and forecast computing. The window used to estimate the model has a width of 500 observations, roughly two years. We believe an estimation window of this size represents a good balance between the precision of quantile regression estimates and the possible presence of a time-change in the parameters of the quantile regression. Moreover, to create a balance between flexibility, efficiency and computational burden, we re-estimated the model with steps of five days. Hence, the first window we considered here

⁶We were also able to detect four (18/01/1991, 26/03/2003, 02/12/2008, and 05/11/2011) and five (18/02/1999, 24/03/2003, 31/05/2007, 11/12/2008 and 05/11/2011) breaks with EPU and EMU being the independent variables respectively, in relation to oil returns.

⁷The results obtained by implementing the DQ and the SQ tests are more thoroughly analyzed in Appendix C.

includes the observations recorded between the first and the 500-th day of the sample. It is important to highlight that when we studied the conditional distribution and quantiles of y_t (and then y_t^2), $t = 1$, that is, the first day of our analysis was January 7, 1986 (here the first two lags of y_t and y_t^2 are available). In the case of bpv_t , in contrast, $t = 1$ coincides with January 9, 2007 due to the availability of data (the dataset is described in Appendix B). In this way, the first window we have in the case of bpv_t coincides with the 1097-th window determined in the case of y_t and y_t^2 .

At $t = 500$ (December 7, 1987, for y_t and y_t^2 , December 8, 2008 for bpv_t), we computed for the first time the test statistic $\hat{J}_T^*$ discussed in Section 2.1 at different quantiles levels, with τ ranging from 0.05 to 0.95 and a step of 0.05, for a total of 19 $\hat{J}_T^*$ values. Given the step of five days in the rolling window procedure, $\hat{J}_T^*$ is computed also in $t = 505, 510, \dots$, until the available time series were exploited.

Figure 1: The dynamic correlations, $\rho(\cdot)$, of the variables y_t , y_t^2 , bpv_t , $x_{1,t}$, $x_{2,t}$. The linear correlations coefficients were computed using a rolling window procedure with a window size of 500 observations and a step of 5 days.

As a preliminary experiment, after determining the rolling windows through the previously mentioned procedure, we could compute the correlations between the variables of interest for each subsample. The series are affected by structural breaks, as mentioned above, and consistent with that evidence, the correlations are not constant over time, as can be seen in Figure 1. We could also condition the dynamic correlations displayed in Figure 1, according to the fact that y_t , y_t^2 and bpv_t are, respectively, lower (greater) than their τ -th in-sample quantiles, for $\tau = \{0.1, 0.5, 0.9\}$. For instance, for each window, we computed the correlation between y_t and $x_{1,t}$ from the values of the pair $(y_t, x_{1,t})$, taking into account only the days where y_t was lower (greater) than its τ -th in-sample quantile. By repeating this method for all the subsamples, we computed the mean of the conditional correlations, denoted as $\bar{\rho}(y_t, x_{1,t})$. The results are given in Table 1. From this table we can see that, for each pair of variables and for each τ , the results differ considerably depending on whether we conditioned the correlation coefficients on the values of y_t , y_t^2 or bpv_t being greater or lower than their respective

in-sample τ quantiles.

Table 1: Conditional average correlations

	$\tau = 0.1$		$\tau = 0.5$		$\tau = 0.9$	
	lower	greater	lower	greater	lower	greater
$\bar{\rho}(y_t, x_{1,t})$	-0.1415	0.0286	-0.0837	0.0785	-0.0474	0.1428
$\bar{\rho}(y_t, x_{2,t})$	-0.1750	0.0272	-0.1023	0.0941	-0.0716	0.1535
$\bar{\rho}(y_t^2, x_{1,t})$	-0.0600	0.0610	0.0064	0.0853	0.0179	0.1277
$\bar{\rho}(y_t^2, x_{2,t})$	-0.0649	0.0810	-0.0176	0.1140	0.0190	0.1993
$\bar{\rho}(bpv_t, x_{1,t})$	-0.0374	0.1845	0.0326	0.1676	0.1612	0.1072
$\bar{\rho}(bpv_t, x_{2,t})$	-0.0584	0.2508	-0.0129	0.2564	0.1664	0.1148

The table reports the average correlations between the variables y_t , y_t^2 , bpv_t , $x_{1,t}$ and $x_{2,t}$. The correlation coefficients were computed by conditioning the pairs of the variables on the values of y_t , y_t^2 and bpv_t , respectively, such that y_t , y_t^2 or bpv_t are lower or greater than their τ -th quantiles, for $\tau = \{0.1, 0.5, 0.9\}$.

To summarize, the results discussed above highlight the importance of using the quantile regression method, given the asymmetric relations among the variables at different τ levels. In particular, at $t = 500$, we estimated for the first time the parameters of the quantile regression model (1) when we focused on $Q_\tau(y_t|\mathcal{W}_{t-1})$, by setting τ from 0.01 to 0.99, with step of 0.01, to obtain quantile vectors of length 99. We proceeded similarly for $Q_\tau(y_t^2|\mathcal{W}_{2,t-1})$ and $Q_\tau(bpv_t|\mathcal{W}_{3,t-1})$. In contrast with the causality analysis, the finer grid of quantiles used in the forecasting exercise was due to the need to estimate the conditional distributions of y_t , y_t^2 and bpv_t with adequate precision. Given the parameter estimates obtained at time $t = 500$, we computed the forecasts of the conditional quantiles and distributions of y_t , y_t^2 and bpv_t for $t = 501, \dots, 505$. Note, we were not making a five-step-ahead forecast, but were simply fixing the model parameters for five days, computing five one-step-ahead forecasts. For instance, to recover the quantile forecasts, we multiplied the values of the predictors observed in $t = 500, \dots, 504$ by the coefficients estimated in $t = 500$.

The second window includes the observations between the 6-th and the 505-th day. Hence, at $t = 505$, we computed for the second time, updating the previous output obtained in $t = 500$, the estimated parameters with which we forecasted, for $t = 506, \dots, 510$, the conditional quantiles and distributions of y_t , y_t^2 and bpv_t . The procedure went on until the entire dataset was completely exploited. We stress that we made such a choice only to reduce computation time. In fact, for each estimation sample, we were, in practice, estimating 99 quantile regressions for each model specification. Re-estimating the model every five observations allowed the computation time to be sensibly reduced.

As for the implementation of the tests proposed by Berkowitz (2001), Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011) (see the details of these tests in Appendix A), starting from $t + 1 = 501$, we compared the forecasts formulated in t with the respective out-of-sample observations y_{t+1} , y_{t+1}^2 and bpv_{t+1} . Therefore, on the basis of those comparisons, we computed, for each point in the forecast sample, the scores characterizing the respective tests. In our analysis, the forecasting evaluation was not carried out on the full sample, but on rolling intervals consisting of

$M = 500$ periods. We motivate this choice by the possibility of controlling, in that way, for changes in the impact of uncertainty measures on the improvement of density forecasts. Therefore, in $t = 1000$ (November 6, 1989 for y_t and y_t^2 , November 8, 2010 for bpv_t), we computed for the first time the four test statistics mentioned above. Similarly, by updating the scores of each test by one period ahead, we computed the statistics for the second time in $t = 1001$, and the procedure was continued until the entire dataset was completely exploited. Note that we made use of two windows: the first refers to the model estimation, while the latter defines the range over which we evaluated the density forecast performances of the restricted and unrestricted models.

In applying the test introduced by Jeong et al. (2012), as in Balcilar et al. (2016a), we might have determined the lag order q using the Schwarz Information Criterion computed on the VAR comprising oil returns and EPU or EMU, and estimated over the full sample. However, our analysis adopted the rolling window procedure described above. Consequently, on the one hand, a large q would imply huge computational costs and, on the other, q most likely would change from one window to another. For this reason, as a rule of thumb, and to create a balance between the precision of the analyses and computational burden, we set $q = 2$ in applying the causality test in quantiles. Likewise, for the forecasting exercise, we set $p = q = r = 2$.

4 Empirical results

4.1 Time variation of quantile causality

We first analyzed the causality in quantiles as detected by the test proposed by Jeong et al. (2012), discussed in Section 2.1. We report in Figure 2 the values of the test statistic $\hat{J}_T^*$ when we studied the causality implications of $x_{1,t}$ (the logarithm of EPU) on the quantiles of y_t (oil return). In contrast, Figure 3 displays the output of the test implemented using $x_{2,t}$ (the logarithm of EMU). The results in Figures 2—3 are very similar: periods in which $\hat{J}_T^*$ takes low values (pointing out the low or inexistent power of the two uncertainty indexes in causing the y_t 's quantiles) are followed by periods of relevant peaks. Nevertheless, the periods in which EPU and EMU are not significant in causing the y_t 's quantiles are more persistent. Moreover, we can see that the causality relations are stronger at the central τ 's levels.

Next, we focus on the impact of EPU and EMU on oil's variance. We observe in Figures 4—5 the results obtained for the quantiles of oil's realized variance (bpv_t), whereas Figures 6—7 display the results we obtained for the quantiles of the squared oil returns (y_t^2). For both bpv_t and y_t^2 we detected a stronger impact of EPU and EMU with respect to the case of y_t , as the periods in which $x_{1,t}$ and $x_{2,t}$ are significant in the causality relationships become more persistent. In contrast with the other cases, the values of $\hat{J}_T^*$ in Figures 6—7 exhibit evident peaks on the extreme quantile levels in some periods. Hence, for causality in risk assessment purposes, we could think that using y_t^2 in place of bpv_t would

Figure 2: The causality of the Economic Policy Uncertainty (in logarithm) on the quantiles of oil's returns. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

Figure 3: The causality of the Market Equity Uncertainty (in logarithm) on the quantiles of oil's returns. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

be a good approximation just for the central quantiles levels; on the other hand, when we consider the tails, this approximation could lead to inaccurate results. We stress that the difference between the results for y_t^2 and bpv_t depend on both the use of oil future prices volatility as a proxy for oil spot prices volatility and on y_t^2 being a noisy proxy for oil spot price volatility.

The aforesaid results are confirmed by those given in Table 2, where we report the p-values of the χ^2

Figure 4: The causality of the Economic Policy Uncertainty (in logarithm) on the quantiles of oil's realized volatility. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

Figure 5: The causality of the Market Equity Uncertainty (in logarithm) on the quantiles of oil's realized volatility. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

test of independence, at a significance level of 0.05, implemented on the variables described below. In particular, we transformed $\hat{J}_T^*$ into a dichotomous variable taking the value of 1 if the null hypothesis of non-causality is rejected at the 0.05 level for a given value of τ , and the value of 0 otherwise. Therefore, the χ^2 test is computed using contingency tables built for each of the pairs in $\{y_t, y_t^2, bpv_t\}$, depending on whether the causal variable is $x_{1,t}$ (in the left panel of Table 2) or $x_{2,t}$ (in the right panel of Table 2). For instance, in the left panel of Table 2, $\chi^2(y_t, y_t^2)$ denotes the p-value of the χ^2 test applied to the dichotomized values displayed in Figures 2—6, for a given value of τ . Given that the time series

Figure 6: The causality of the Economic Policy Uncertainty (in logarithm) on the quantiles of oil's squared returns. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

Figure 7: The causality of the Market Equity Uncertainty (in logarithm) on the quantiles of oil's squared returns. The figure reports the values of $\hat{J}_T^*$ for different quantile levels, computed using a rolling procedure with a window size of 500 observations and step of 5 periods ahead. The blue cutting plane represents the 5% confidence bound.

of y_t and y_t^2 cover a longer time period with respect to bpv_t , for consistency we applied the χ^2 test of independence by using the same time interval for all the variables, which starts on December 8, 2008 (i.e. the closing point of the first window for bpv_t , where we computed for the first time the values in Figures 4—5). We point out that we never rejected the null hypothesis of independence (at the 0.05 level) in the case of $\chi^2(y_t, y_t^2)$, whereas we rarely rejected it for $\chi^2(y_t, bpv_t)$. In contrast, the null hypothesis of independence was almost always rejected, with very low p-values, for the pair $\{y_t^2, bpv_t\}$; the exceptions occurred at extreme τ levels. Therefore, using y_t^2 in the Jeong et al. (2012)'s test for

analysing causality in risk, in place of bpv_t , leads to satisfactory approximations for the central values of τ ; that evidence does not hold at extreme quantiles levels. We might also read this result from a different viewpoint: the use of oil future prices to recover a proxy for oil spot price volatility seems appropriate at central quantiles only.

Table 2: χ^2 test of independence

τ	causal variable $x_{1,t}$			causal variable $x_{2,t}$		
	$\chi^2(y_t, y_t^2)$	$\chi^2(y_t, bpv_t)$	$\chi^2(y_t^2, bpv_t)$	$\chi^2(y_t, y_t^2)$	$\chi^2(y_t, bpv_t)$	$\chi^2(y_t^2, bpv_t)$
0.05	0.61	1.00	0.99	1.00	0.54	0.77
0.15	0.20	1.64E-03	8.33E-06	0.82	0.05	0.03
0.25	0.58	0.02	7.49E-05	0.93	0.12	0.46
0.35	0.84	7.33E-03	8.24E-04	1.00	0.55	2.11E-05
0.45	1.00	0.03	2.48E-03	1.00	1.00	1.32E-08
0.55	1.00	2.46E-03	5.61E-03	0.89	0.05	9.97E-03
0.65	0.96	0.08	0.03	0.80	0.51	4.97E-03
0.75	1.00	0.61	2.39E-03	0.78	0.14	0.02
0.85	1.00	0.94	0.27	0.46	0.38	0.02
0.95	1.00	0.11	1.00	0.95	0.54	0.02

The table shows the p-values of the χ^2 test of independence, at a significance level of 0.05. The test was applied by transforming $\widehat{J}_T^*$ into a dichotomous variable taking a value of 1 if the null hypothesis of non-causality in quantile, for a given value of τ , was rejected at the 0.05 level, and a value of 0 otherwise. Therefore, the χ^2 test was computed from contingency tables built for each of the pairs in $\{y_t, y_t^2, bpv_t\}$, depending on whether the causal variable was $x_{1,t}$ (in the left panel of Table 2) or $x_{2,t}$ (in the right panel of Table 2).

4.2 Model assessment

We first estimated the parameters of the quantile regression model in (1), that is, $Q_\tau(y_t|\mathcal{W}_{t-1}) = \alpha_0(\tau) + \beta_1(\tau)y_{t-1} + \beta_2(\tau)y_{t-2} + \delta_1(\tau)x_{1,t-1} + \delta_2(\tau)x_{1,t-2} + \lambda_1(\tau)x_{2,t-1} + \lambda_2(\tau)x_{2,t-2}$. By replacing y_t by y_t^2 (bpv_t) and $\mathcal{W}_{t-1}$ by $\mathcal{W}_{2,t-1}$ ($\mathcal{W}_{3,t-1}$) we also estimated the parameters of $Q_\tau(y_t^2|\mathcal{W}_{2,t-1})$ and $Q_\tau(bpv_t|\mathcal{W}_{3,t-1})$. After estimating the parameters of these quantile regression models for each of the subsamples resulting from the rolling window procedure, we computed their respective average values and standard deviations.

We checked that, on average, the coefficients' p-values were greater than 0.05, pointing out that the explanatory variables are not always statistically significant in explaining the conditional quantiles of y_t , y_t^2 and bpv_t over time. Hence, we report in Table 3 the mean (columns 2—7) and the standard deviation (columns 8—13) of the coefficients, conditional on their respective p-values being less than or equal to 0.05. These average values are denoted as $\bar{\beta}_j(\tau)$, $\bar{\delta}_j(\tau)$ and $\bar{\lambda}_j(\tau)$, whereas the standard deviations are denoted as $\sigma_{\beta_j(\tau)}$, $\sigma_{\delta_j(\tau)}$, $\sigma_{\lambda_j(\tau)}$, for $j = \{1, 2\}$. For simplicity, we display the results obtained at $\tau = \{0.1, 0.5, 0.9\}$.

Starting with $Q_\tau(y_t|\mathcal{W}_{t-1})$'s estimated parameters, on average, the impact of the explanatory variables changes according to the τ levels, counter-evidence for the so-called location-shift hypothesis, which assumes the homogeneous effects of the covariates across the conditional quantiles of the response variable. It is possible to observe a precise trend of the coefficients' values over τ : negative for $\bar{\beta}_j(\tau)$

Table 3: Quantile regression output

τ	$\bar{\beta}_1(\tau)$	$\bar{\beta}_2(\tau)$	$\bar{\delta}_1(\tau)$	$\bar{\delta}_2(\tau)$	$\bar{\lambda}_1(\tau)$	$\bar{\lambda}_2(\tau)$	$\sigma_{\beta_1(\tau)}$	$\sigma_{\beta_2(\tau)}$	$\sigma_{\delta_1(\tau)}$	$\sigma_{\delta_2(\tau)}$	$\sigma_{\lambda_1(\tau)}$	$\sigma_{\lambda_2(\tau)}$
Estimates in $Q_\tau(y_t \mathcal{W}_{t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)y_{t-j} + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$												
0.1	0.024	17.893	-0.835	-0.040	-0.529	-0.020	19.122	5.997	0.399	0.694	0.207	0.683
0.5	-11.531	-7.864	-0.464	0.291	0.011	0.130	1.613	0.543	0.139	0.332	0.257	0.260
0.9	-16.625	-18.386	0.559	0.829	0.555	0.765	15.213	6.581	0.146	0.414	0.401	0.119
Estimates in $Q_\tau(y_t^2 \mathcal{W}_{2,t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)y_{t-j}^2 + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$												
0.1	0.467	1.027	0.001	-4e-04	-2e-04	-1e-05	0.444	0.903	3e-05	1e-03	2e-05	3e-04
0.5	6.711	7.977	0.009	0.011	0.003	0.005	5.032	2.895	0.004	0.007	0.005	0.003
0.9	67.069	85.650	0.038	0.020	0.050	0.004	27.312	34.422	0.016	0.065	0.031	0.046
Estimates in $Q_\tau(bpv_t \mathcal{W}_{3,t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)bpv_{t-j} + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$												
0.1	29.679	13.277	0.002	0.005	0.001	0.001	4.512	2.831	0.002	0.002	0.000	0.000
0.5	44.738	25.433	0.004	0.004	0.001	0.001	4.187	6.557	0.001	0.001	0.000	0.001
0.9	78.199	59.270	0.023	0.012	0.006	0.002	21.408	14.328	0.003	0.003	0.003	0.002

The table reports the average values (%), in columns 2—7, and the standard deviations (%), in columns 8—13, computed for the subsamples determined by the rolling window procedure, of the estimated parameters, conditional on their being statistically significant at the level of 0.05. The rolling window procedure was applied by using a window size of 500 observations and steps of 5 days ahead.

and positive for $\bar{\delta}_j(\tau)$, $\bar{\lambda}_j(\tau)$, $j = \{1, 2\}$. On average, the lags of y_t have a positive impact on the left tail of the response variable's conditional distribution; on the other hand, their effects become negative at medium-high τ levels. This was expected as past negative returns lead to an increase in the series dispersion, and thus moved the 0.1 (0.9) quantile further to the left (right), with an additional effect on the median. On the contrary, positive returns shrink the density toward the median, which then also moves to the right. We interpret this evidences as a form of asymmetry, where the sign of the shocks lead to opposite effects on the quantiles, and thus on the distribution, of the target variable.

The opposite phenomenon was observed for $x_{1,t-j}$ and $x_{2,t-j}$, $j = \{1, 2\}$; for the uncertainty indexes, we were expecting these signs. In fact, an increase in uncertainty moves the lower quantiles to the left and the upper quantiles to the right, with the impact on the median being smaller than that on other quantiles for $j = 1$. With the exception of $x_{2,t-j}$, $j = \{1, 2\}$, the coefficients of the other explanatory variables are less volatile at the central levels of τ . In Table 4 we report the number of subsamples in which each coefficient turns out to be statistically significant at the level of 0.05. It is possible to see that at τ equal to 0.1, 0.5 and 0.9, $x_{2,t-2}$, $x_{2,t-1}$ and y_{t-1} record, respectively, the highest number of periods in which their coefficients are statistically significant.

Moving to the estimation of $Q_\tau(bpv_t|\mathcal{W}_{3,t-1})$'s parameters, all the predictors have a positive impact on the estimated quantiles. Also, the magnitude of the coefficients is almost always a positive function of τ , mainly in the case of $\bar{\beta}_1(\tau)$ and $\bar{\beta}_2(\tau)$, which reach considerable levels at medium-high τ values.

As for $Q_\tau(y_t^2|\mathcal{W}_{2,t-1})$, just $\bar{\delta}_2(0.1)$, $\bar{\lambda}_1(0.1)$ and $\bar{\lambda}_2(0.1)$ are negative; nevertheless these coefficients take very low values. With the exception of $\bar{\lambda}_2(\tau)$, all the other coefficients exhibit, on average, an increasing trend over τ . Interestingly, the coefficients of $Q_\tau(y_t^2|\mathcal{W}_{2,t-1})$ and of $Q_\tau(bpv_t|\mathcal{W}_{3,t-1})$ follow a similar trend. The larger impact of the lagged realized volatility and of the squared lagged returns on the upper quantiles again was expected, signalling that large movements (either positive or negative)

lead to a huge increase in the risk.

We report in Table 4 the number of subsamples in which each coefficient is statistically significant at the 0.05 level. By way of reminder, the total number of rolled windows is equal to 1430 in the case of y_t and y_t^2 , and is equal to 334 for bpv_t , due to the limited availability of data. The most relevant evidence in Table 4 is that bpv_{t-1} almost always turns out to be significant over time, for all the levels of τ ; secondly, bpv_{t-2} also reaches high levels of $n_{\beta_2(\tau)}$. Thus, the realized volatility of oil exhibits relevant persistence in its lags, consistent with the findings in Corsi (2009). Another important variable is $x_{2,t-1}$ at $\tau = 0.9$, given that it is statistically significant in 186 out of 334 subsamples. As for y_t and y_t^2 , at $\tau = \{0.1, 0.9\}$, y_{t-1} , y_{t-2} and $x_{1,t-1}$ record the highest number of periods in which their coefficients are statistically significant at the 5% level.

The p-values in Table 4 indicate that EPU and EMU are significant in forecasting oil's returns and risk in a few cases. This evidence is consistent with the results shown in Figures 2—3, where EPU and EMU seldom cause the y_t 's quantiles. In contrast, EPU and EMU almost always cause the variance's quantiles (Figures 4—7). This difference might be due to the fact that we obtained the p-values in Table 4 by including both $x_{1,t-j}$ and $x_{2,t-j}$, with $j = \{1, 2\}$, in the same quantile regression model, whereas we implemented the Jeong et al. (2012)'s test by using either $x_{1,t-j}$ or $x_{2,t-j}$. Hence, we also estimated the quantile regression models discussed above, including either $x_{1,t-j}$ or $x_{2,t-j}$. The results are qualitatively similar to those with both indexes.⁸ Consequently, from a forecasting perspective, EMU and EPU might be considered substitutes. A further explanation is that we estimated a linear quantile regression model, whereas the Jeong et al. (2012)'s test does not have any particular functional form. Therefore, the different specification of the two approaches can lead to different conclusions when analyzing the causality implications of EPU and EMU on oil's risk.

Table 4: Persistence of significance over the rolled windows

τ	$n_{\beta_1(\tau)}$	$n_{\beta_2(\tau)}$	$n_{\delta_1(\tau)}$	$n_{\delta_2(\tau)}$	$n_{\lambda_1(\tau)}$	$n_{\lambda_2(\tau)}$
Estimates in $Q_\tau(y_t \mathcal{W}_{t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)y_{t-j} + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$						
0.1	112	37	165	114	54	190
0.5	102	9	30	64	113	35
0.9	322	156	93	47	105	53
Estimates in $Q_\tau(y_t^2 \mathcal{W}_{2,t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)y_{t-j}^2 + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$						
0.1	43	63	50	13	35	24
0.5	142	50	150	57	145	103
0.9	174	27	126	117	69	90
Estimates in $Q_\tau(bpv_t \mathcal{W}_{3,t-1}) = \alpha_0(\tau) + \sum_{j=1}^2 \beta_j(\tau)bpv_{t-j} + \sum_{j=1}^2 \delta_j(\tau)x_{1,t-j} + \sum_{j=1}^2 \lambda_j(\tau)x_{2,t-j}$						
0.1	330	81	33	20	13	10
0.5	310	298	32	26	33	109
0.9	333	176	0	3	186	29

The table reports the number of subsamples, determined by the rolling window procedure, in which each coefficient turns out to be statistically significant at a level of 0.05. The rolling window procedure was applied by using a window size of 500 observations and steps of 5 days ahead.

⁸Results are available upon request.

Next, we show that the method we used turns out to be very effective in obtaining valid estimations of distributions. For instance, Subfigure 8(a) shows the conditional distribution of y_t , estimated from the first of the rolled windows. We emphasize that the problem of crossing in quantiles vanishes by applying the quantile bootstrap method proposed by Chernozhukov et al. (2010). Moreover, using the Epanechnikov kernel method results in obtaining smoother distributions. Subfigure 8(b) shows the conditional density of y_t , respectively, estimated from the first and the last windows of the rolling window procedure, when applying the Epanechnikov kernel method. Notably, the shape of each density changes over time, thus supporting the need for a rolling evaluation.⁹

Figure 8: Conditional distribution and density of y_t . Subfigure (a) displays the conditional distribution of y_t , estimated from the first subsample determined through the rolling window procedure. ‘Original’, ‘adjusted’ and ‘kernel’ stand for the distributions arising directly from Model (1), the one obtained by adjusting the original estimates through the quantile bootstrap method proposed by Chernozhukov et al. (2010), and the one built by means of the Epanechnikov kernel, respectively. Subfigure (b) shows the conditional density of y_t , estimated from the first and the last windows determined by the rolling window procedure, by applying the Epanechnikov kernel method.

4.3 The asymmetric impact of the uncertainty measures

In this section, we present an evaluation of the possible asymmetric effects of the uncertainty indexes on oil movements. In doing so, we first centred to zero both $x_{1,t-j}$ and $x_{2,t-j}$, $j = \{1, 2\}$, by subtracting their respective means; these new variables are denoted by $x_{1,t-j}^*$ and $x_{2,t-j}^*$, $j = \{1, 2\}$, respectively. Next, we defined the indicator functions $\mathbf{1}_{\{x_{1,t-j}^* < 0\}}$ and $\mathbf{1}_{\{x_{2,t-j}^* < 0\}}$, $j = \{1, 2\}$, which take a value of 1 if the condition in $\{\cdot\}$ is true and a value of 0 otherwise, and obtained the following model:

$$\begin{aligned}
Q_\tau^d(y_t | \mathcal{W}_{t-1}) &= \alpha_0(\tau) + \beta_1(\tau)y_{t-1} + \beta_2(\tau)y_{t-2} + \delta_1^d(\tau)x_{1,t-1}^* + \delta_2^d(\tau)x_{1,t-2}^* + \lambda_1^d(\tau)x_{2,t-1}^* \\
&+ \lambda_2^d(\tau)x_{2,t-2}^* + \delta_1^*(\tau)\mathbf{1}_{\{x_{1,t-1}^* < 0\}}x_{1,t-1}^* + \delta_2^*(\tau)\mathbf{1}_{\{x_{1,t-2}^* < 0\}}x_{1,t-2}^* \\
&+ \lambda_1^*(\tau)\mathbf{1}_{\{x_{2,t-1}^* < 0\}}x_{2,t-1}^* + \lambda_2^*(\tau)\mathbf{1}_{\{x_{2,t-2}^* < 0\}}x_{2,t-2}^* + \varepsilon_t^d(\tau),
\end{aligned} \tag{4}$$

Similarly, we also estimated the parameters of $Q_\tau^d(y_t^2 | \mathcal{W}_{2,t-1})$ ($Q_\tau^d(bpv_t | \mathcal{W}_{3,t-1})$) by replacing y_t in (4) with y_t^2 (bpv_t) and $\mathcal{W}_{t-1}$ with $\mathcal{W}_{2,t-1}$ ($\mathcal{W}_{3,t-1}$). In evaluating the asymmetric effects of EPU

⁹We observed the same phenomenon for y_t^2 and bpv_t . The results are available on request.

(EMU) on oil movements, it is important to notice that the impact of $x_{1,t-j}^*$ ($x_{2,t-j}^*$), $j = \{1, 2\}$, was quantified by $\delta_j^d(\tau)$ ($\lambda_j^d(\tau)$) if $x_{1,t-j}^* \geq 0$ ($x_{2,t-j}^* \geq 0$); in contrast, its impact is equal to $\delta_j^d(\tau) + \delta_j^*(\tau)$ ($\lambda_j^d(\tau) + \lambda_j^*(\tau)$) if $x_{1,t-j}^* < 0$ ($x_{2,t-j}^* < 0$). We report the estimates of these new models accounting for the asymmetric impact of EPU and EMU in Table 5. Here, we display the average values of the coefficients over the rolled subsamples (window size of 500 observations and steps of 5 days ahead), conditional on them being statistically significant at a level of 5%; we also report their standard deviations. More precisely, for instance, in order to evaluate correctly the asymmetric effects of $x_{1,t-1}^*$, for each window, we considered the cases where all the coefficients $[\widehat{\delta}_1^d(\tau), \widehat{\delta}_1^{**}(\tau) = \widehat{\delta}_1^d(\tau) + \widehat{\delta}_1^*(\tau)]$ are simultaneously significant, and then we computed their average values.

Table 5: The asymmetric impact of uncertainty on oil movements

τ	$\bar{\delta}_1^d(\tau)$	$\bar{\delta}_1^{**}(\tau)$	$\bar{\delta}_2^d(\tau)$	$\bar{\delta}_2^{**}(\tau)$	$\bar{\lambda}_1^d(\tau)$	$\bar{\lambda}_1^{**}(\tau)$	$\bar{\lambda}_2^d(\tau)$	$\bar{\lambda}_2^{**}(\tau)$
Estimates of the $Q_\tau^d(y_t \mathcal{W}_{t-1})$'s parameters								
0.1	-2.44 (1.9)	-0.12 (0.4)	2.42 (1.2)	-1.16 (0.6)	-1.07 (1.0)	0.21 (0.2)	-1.74 (0.5)	0.33 (1.1)
0.5	0.95 (0.4)	-0.33 (0.2)	0.33 (1.1)	-0.38 (0.2)	-0.92 (0.8)	0.32 (0.1)	-1.53 (0.1)	0.29 (1.1)
0.9	2.20 (1.9)	-0.92 (0.8)	1.45 (1.1)	-0.50 (0.7)	-0.65 (1.9)	0.54 (0.4)	-1.00 (0.1)	0.33 (0.1)
Estimates of the $Q_\tau^d(y_t^2 \mathcal{W}_{2,t-1})$'s parameters								
0.1	1.00 (8.2)	-0.10 (2.2)	-2.92 (10.2)	0.19 (3.1)	2.23 (2.9)	-0.01 (3.2)	0.06 (0.5)	-0.02 (0.6)
0.5	2.68 (0.8)	-0.30 (1.4)	-2.07 (2.6)	0.17 (0.3)	3.13 (1.0)	-0.12 (0.2)	1.90 (1.2)	-0.09 (1.3)
0.9	14.14 (12.0)	-0.52 (3.1)	29.51 (15.2)	-2.90 (3.9)	5.56 (3.7)	-1.28 (1.8)	4.34 (4.8)	-1.89 (2.5)
Estimates of the $Q_\tau^d(bpv_t \mathcal{W}_{3,t-1})$'s parameters								
0.1	0.45 (2.3)	-0.11 (3.6)	0.09 (2.2)	-0.05 (1.9)	0.02 (4.8)	-0.05 (3.9)	0.08 (1.2)	-0.01 (3.5)
0.5	1.32 (1.3)	-0.38 (2.1)	0.02 (1.4)	-0.06 (2.9)	0.05 (3.7)	-0.94 (2.5)	1.31 (5.7)	-0.03 (2.1)
0.9	3.26 (9.4)	-0.78 (8.1)	3.23 (9.4)	-1.23 (9.3)	2.2 (8.7)	-1.02 (13.6)	2.28 (9.2)	-0.73 (8.9)

The table reports the average values (%), computed over the subsamples derived from the rolling window procedure, of the uncertainty indexes' coefficients (in brackets we report their standard deviations) estimated for the following models: $Q_\tau^d(y_t|\mathcal{W}_{t-1}) = \alpha_0(\tau) + \beta_1(\tau)y_{t-1} + \beta_2(\tau)y_{t-2} + \delta_1^d(\tau)x_{1,t-1}^* + \delta_2^d(\tau)x_{1,t-2}^* + \lambda_1^d(\tau)x_{2,t-1}^* + \lambda_2^d(\tau)x_{2,t-2}^* + \delta_1^*(\tau)\mathbf{1}_{\{x_{1,t-1}^* < 0\}}x_{1,t-1}^* + \delta_2^*(\tau)\mathbf{1}_{\{x_{1,t-2}^* < 0\}}x_{1,t-2}^* + \lambda_1^*(\tau)\mathbf{1}_{\{x_{2,t-1}^* < 0\}}x_{2,t-1}^* + \lambda_2^*(\tau)\mathbf{1}_{\{x_{2,t-2}^* < 0\}}x_{2,t-2}^*$, $Q_\tau^d(y_t^2|\mathcal{W}_{2,t-1}) = \alpha_0(\tau) + \beta_1(\tau)y_{t-1}^2 + \beta_2(\tau)y_{t-2}^2 + \delta_1^d(\tau)x_{1,t-1}^* + \delta_2^d(\tau)x_{1,t-2}^* + \lambda_1^d(\tau)x_{2,t-1}^* + \lambda_2^d(\tau)x_{2,t-2}^* + \delta_1^*(\tau)\mathbf{1}_{\{x_{1,t-1}^* < 0\}}x_{1,t-1}^* + \delta_2^*(\tau)\mathbf{1}_{\{x_{1,t-2}^* < 0\}}x_{1,t-2}^* + \lambda_1^*(\tau)\mathbf{1}_{\{x_{2,t-1}^* < 0\}}x_{2,t-1}^* + \lambda_2^*(\tau)\mathbf{1}_{\{x_{2,t-2}^* < 0\}}x_{2,t-2}^*$ and $Q_\tau^d(bpv_t|\mathcal{W}_{3,t-1}) = \alpha_0(\tau) + \beta_1(\tau)bpv_{t-1} + \beta_2(\tau)bpv_{t-2} + \delta_1^d(\tau)x_{1,t-1}^* + \delta_2^d(\tau)x_{1,t-2}^* + \lambda_1^d(\tau)x_{2,t-1}^* + \lambda_2^d(\tau)x_{2,t-2}^* + \delta_1^*(\tau)\mathbf{1}_{\{x_{1,t-1}^* < 0\}}x_{1,t-1}^* + \delta_2^*(\tau)\mathbf{1}_{\{x_{1,t-2}^* < 0\}}x_{1,t-2}^* + \lambda_1^*(\tau)\mathbf{1}_{\{x_{2,t-1}^* < 0\}}x_{2,t-1}^* + \lambda_2^*(\tau)\mathbf{1}_{\{x_{2,t-2}^* < 0\}}x_{2,t-2}^*$. We conditioned the estimates on them being statistically significant at a level of 0.05. The rolling window procedure was applied by using a window size of 500 observations and step of 5 days ahead. $\bar{\delta}_j^{**}(\tau)$ is the conditional average value of the sums $\delta_j^d(\tau) + \delta_j^*(\tau)$, computed from the windows where the two coefficients are simultaneously significant; similarly, $\bar{\lambda}_j^{**}(\tau)$ is the conditional average value $\lambda_j^d(\tau) + \lambda_j^*(\tau)$, $j = \{1, 2\}$.

We can see from Table 5 that, on average, both the uncertainty indexes have asymmetric effects on oil movements, and that the impact is stronger in the states where they take high values, i.e. when $x_{1,t-j}^*$ and $x_{2,t-j}^*$ take positive values. Indeed, the means of $\widehat{\delta}_j^d(\tau)$ and $\widehat{\lambda}_j^d(\tau)$ are almost always greater, in absolute value, than the means of $(\widehat{\delta}_j^d(\tau) + \widehat{\delta}_j^*(\tau))$ and $(\widehat{\lambda}_j^d(\tau) + \widehat{\lambda}_j^*(\tau))$, $j = \{1, 2\}$, respectively. This is a somewhat expected result suggesting that increases in uncertainty do have a larger impact on oil movements compared with decreases in uncertainty.

4.4 Out-of-sample forecast evaluation of quantile causality

We evaluated the predictive accuracy of the methods discussed in Section 2.2, focusing on the contribution of the two uncertainty indexes ($x_{1,t}$ and $x_{2,t}$) in forecasting the quantiles and the distributions of y_t , y_t^2 and bpu_t . The case of y_t is displayed in Figure 9. Starting from Berkowitz (2001), we tested the null hypothesis of correct specification and report in Figure 9(a) the p-values obtained from four different models: Model 1 includes all the predictors, that is, y_{t-j} , $x_{1,t-j}$ and $x_{2,t-j}$, with $j = \{1, 2\}$. Model 2 has just y_{t-j} , whereas Model 3 comprises y_{t-j} and $x_{1,t-j}$. Finally, Model 4 includes y_{t-j} and $x_{2,t-j}$, with $j = \{1, 2\}$. For all the four models, we can see that there are periods in which the null hypothesis of correct specification is not rejected at the 0.05 significance level, and others where the null hypothesis is rejected. Notably, during the years 1999—2002 and 2008, the null hypothesis is not rejected just for Model 1, highlighting the importance of exploiting the joint predictive power of EPU and EMU.

Figure 9: Predictive accuracy for the y_t conditional distribution and quantiles over the rolled windows. The tests were applied by placing different weights on the different regions of the conditional distribution of the response variable. In Subfigure 9(a), Model 1 includes all the selected predictors, that is, y_{t-j} , $x_{1,t-j}$ and $x_{2,t-j}$, with $j = \{1, 2\}$. Model 2 has just y_{t-j} , whereas Model 3 comprises y_{t-j} and $x_{1,t-j}$. Finally, Model 4 includes y_{t-j} and $x_{2,t-j}$, with $j = \{1, 2\}$. In Subfigures 9(b)—9(d) the black horizontal lines represent the 5% confidence bounds, whereas in Subfigure 9(a) the black horizontal line represents the value threshold of 0.05.

The direct comparisons between the restricted (the one including just the lags of y_t) and the unrestricted (which includes also the lags of $x_{1,t}$ and $x_{2,t}$) models, based on the tests proposed by Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011), were evaluated and are reported in Figures 9(b)—9(d). Here, the tests statistics change their signs over time, highlighting

periods in which the unrestricted model works better, followed by others where the best performance is recorded by the restricted model.¹⁰ Nevertheless, the null hypothesis of equal performance is not always rejected at the 5% level and the periods in which the unrestricted model records the best performance, statistically significant, are less frequent than the ones in which it is outperformed by the restricted model. In general, all the tests give evidence of the best performance of the unrestricted model in the second half of the 2000s and at the end of the 2000s. Besides, the Diks et al. (2011) test detects further periods, namely in the middle of the 1990s and at the beginning of the 2000s, where the unrestricted model outperforms the restricted one, mainly when we focus on the center of the distribution.

Here we focus on risk, and for that purpose, we report in the following the case of oil's volatility, estimated through the realized bipower variation (subplots on the right side), and the case of squared oil returns (on the left side). In particular, the tests proposed by Berkowitz (2001), Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011) are reported in Figures 10—13, respectively. We notice that, during the years 2010—2015, i.e. the period in common between the y_t^2 and the bpv_t series, there is no evidence in favour of the (statistically significant) better performance of the unrestricted model, highlighting the poor contribution of EPU and EMU in forecasting oil risk's conditional distribution and quantiles.

As for the years before 2010 where we focus just on y_t^2 , the Berkowitz (2001) test shows periods in which the null hypothesis is not rejected just for Model 4, such as the end of the 1980s, and between the years 2003-2005 and 2009-2010, highlighting the relevant contribution of EMU in the forecasting exercise. Among the other tests, the Diks et al. (2011) test records the highest number of windows where the null hypothesis is rejected in favour of the better performance of the unrestricted model, mainly in the left tail and at the center of the y_t^2 conditional distribution. The tests developed by Diebold and Mariano (2002) and Gneiting and Ranjan (2011) show positive evidence at high y_t^2 quantiles, in the years 2009-2010.

We now assemble the information from the four tests discussed above to highlight the periods where EPU and EMU turn out to be crucial in forecasting the y_t conditional distribution and quantiles. Hereafter, we focus just on y_t and y_t^2 , given the poor evidence observed in the case of bpv_t during the years 2010-2015. For this purpose, we computed, for each of the applied tests—those developed by Berkowitz (2001), Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011)—a dummy variable denoted D_t^{pred} taking the value of 1 if the unrestricted model records a (statistically) better performance at the level of 0.05 than the restricted one at t , and the value of 0 otherwise. In the case of the Berkowitz (2001) test, D_t^{pred} takes a value of 1 if the null hypothesis is not rejected for the unrestricted model (which includes all the available covariates), and for the same window, is rejected for the restricted one (which includes just the lagged values of y_t or y_t^2) at t . In order to clean

¹⁰We emphasize that, for the Diks et al. (2011) test, the unrestricted model outperforms the restricted one if the test statistic is positive. The opposite holds for the Diebold and Mariano (2002) and Gneiting and Ranjan (2011) tests.

the series from the periods where the better performance of one of the two models lasts for only a few day, being negligible, we computed, for each test, the following moving average:

$$SD_t^{pred} = \frac{1}{2M_s + 1} \sum_{t-M_s}^{t+M_s} D_t^{pred}. \quad (5)$$

In our work, we set $M_s = 10$, hence the moving averages span by 21 days, and identified with more smoothness the periods where the unrestricted model performed better than those where $SD_t^{pred} \geq 0.5$.

Figure 10: The p-values of the Berkowitz (2001) test over the rolled windows. Subfigure 10(a) and 10(b) refer, respectively, to the forecasts of the distributions of y_t^2 and bpv_t . Model 1 includes all the selected predictors, that is, y_{t-j}^2 (in Subfigure 10(a)) or of bpv_{t-j} (in Subfigure 10(b)), $x_{1,t-j}$ and $x_{2,t-j}$, with $j = \{1, 2\}$. Model 2 has just the lags of y_{t-j}^2 (in Subfigure 10(a)) or of bpv_{t-j} (in Subfigure 10(b)). Model 3 comprises y_{t-j}^2 (in Subfigure 10(a)) or of bpv_{t-j} (in Subfigure 10(b)) and $x_{1,t-j}$, whereas Model 4 includes y_{t-j}^2 (in Subfigure 10(a)) or of bpv_{t-j} (in Subfigure 10(b)) and $x_{2,t-j}$, with $j = \{1, 2\}$. The black horizontal line represents the threshold value of 0.05.

Figure 11: The Diebold and Mariano (2002) test statistic values over the rolled windows. Subfigures 11(a) and 11(b) refer, respectively, to the forecasts of the y_t^2 and the bpv_t distributions, where we compared the restricted (which contain just the lags of y_t^2 and bpv_t) and the unrestricted models (which include also the lags of $x_{1,t}$ and $x_{2,t}$). The test was applied for three different τ values: 0.05 (green lines), 0.50 (yellow lines) and 0.95 (blue lines). The black horizontal lines represent the 5% confidence bounds.

We display in Figures 14–15 the periods where EPU and EMU are crucial in the forecasting exercise, with a different colour for each test. From Figure 14, we can see that EPU and EMU played an important role in forecasting the y_t 's conditional quantiles and distributions during the years

Figure 12: The Diks et al. (2011) test statistic values over the rolled windows. Subfigures 12(a) and 12(b) refer, respectively, to the forecasts of the y_t^2 and the bpv_t distributions, where we compared the restricted (which contain just the lags of y_t^2 and bpv_t) and the unrestricted models (which include also the lags of $x_{1,t}$ and $x_{2,t}$). The test was applied by placing greater emphasis on the center (green lines), on the right tail (blue lines) and on the left tail (yellow lines) of the conditional distributions. The black horizontal lines represent the 5% confidence bounds.

Figure 13: The Gneiting and Ranjan (2011) test statistic values over the rolled windows. Subfigures 13(a) and 13(b) refer, respectively, to the forecasts of the y_t^2 and the bpv_t distributions, where we compared the restricted (which contain just the lags of y_t^2 and bpv_t) and the unrestricted models (which include also the lags of $x_{1,t}$ and $x_{2,t}$). The test was applied by placing greater emphasis on the center (green lines), on the right tail (blue lines) and on the left tail (yellow lines) of the conditional distributions. The black horizontal lines represent the 5% confidence bounds.

2005—2007 and 2008—2010. These periods are close to two special events: the ‘2008 oil price bubble’, which spans the years 2007—2008, and the US subprime crisis, marked by the Lehman Brothers’ default in September 2008. As for y_t^2 (Figure 15), similar to y_t , we have evidence of the crucial role of EPU and EMU during the years 2005—2007 and 2008—2010.

Therefore, EPU and EMU turn out to be crucial variables in forecasting the conditional quantiles and distributions of oil movements only in some periods. Such evidence supports the use of the rolling window procedure to capture these dynamics, than carrying out a full sample analysis. As for bpv_t and y_t^2 , with which we focus on risk, the results arising from the forecasting exercise provide similar suggestions in terms of the average behaviour of the coefficients of both $Q_\tau(y_t^2|\mathcal{W}_{2,t-1})$ and $Q_\tau(bpv_t|\mathcal{W}_{3,t-1})$. Some similarities are observed also in terms of a predictive accuracy assessment, in the sense that, during the period in common between the bpv_t and the y_t^2 series, all the implemented

Figure 14: Evidence of the crucial impact of EPU and EMU in forecasting the y_t conditional distribution and quantiles over time, detected by several tests and linked to the y_t (a), $x_{1,t}$ (b), and $x_{2,t}$ (c) series. The subfigures display the evidences resulting from the tests proposed by Diks et al. (2011), with focuses on the right (A) and the center (B) parts of the distribution; Gneiting and Ranjan (2011), with focus on the right tail of the distribution (C); Diebold and Mariano (2002) with focuses on the 95-th (D) and the 5-th percentile (E); and Berkowitz (2001) (F).

Figure 15: Evidence of the crucial impact of EPU and EMU in forecasting the y_t^2 conditional distribution and quantiles over time, detected by several tests and linked to the y_t (a), $x_{1,t}$ (b), and $x_{2,t}$ (c) series. The subfigures display the evidence resulting from the tests proposed by Diks et al. (2011), with focuses on the center (A) and the left (B) parts of the distribution; Gneiting and Ranjan (2011), with focus on the right tail of the distribution (C); Diebold and Mariano (2002), with focus on the 95-th percentile (D); and Berkowitz (2001) (E).

tests suggest poor evidence in favour of significant contributions of EPU and EMU in improving forecasting precision. As for the causality exercise, the periods in which the two uncertainty indexes

significantly improve forecasting accuracy are more persistent with respect to the ones recorded in the case of y_t ; the reason for this might be the following: with the squared returns of oil, we focused on the y_t volatility, a measure of dispersion (and thus uncertainty) that better fits the nature of EPU and EMU, which are themselves uncertainty indicators.

At this stage, it is important to point out that, although robust predictive inference is derived based on the causality-in-quantiles test, it would also be interesting to estimate the magnitude and direction of the effects of uncertainties on oil market movements at various quantiles. However, in a nonparametric framework, this is not straightforward. We will need to employ the first-order partial derivative. Estimation of the partial derivatives for nonparametric models can experience complications because nonparametric methods exhibit slow convergence rates, which can depend on the dimensionality and smoothness of the underlying conditional expectation function. What one could however do is to look at a statistic that summarizes the overall effect or the global curvature (i.e., the global sign and magnitude), but not the entire derivative curve. In this regard, a natural measure of the global curvature is the average derivative (AD). One could use the conditional pivotal quantile, based on approximation or the coupling approach of Belloni et al. (2017), to estimate the partial ADs. The pivotal coupling approach additionally can approximate the distribution of AD using Monte Carlo simulation. Given that in our case, the focus is on predictability of the oil market movements, and not necessarily on the sign (direction) of the effect at this stage, we leave this for future research. However, even though we cannot draw one-to-one correspondence between standard quantile regressions and our nonparametric causality-in-quantiles test, preliminary evidence tends to suggest that uncertainties tends to reduce oil returns and increase its volatility.¹¹

Next, we summarize our results and discuss the economic implications of these for various agents in the economy. First turning to oil returns, we observe that uncertainty indexes tend to have stronger effects when the economy and the oil market are in turmoil, or where either is in turmoil. This is because during these periods uncertainty is likely to be higher relative to its average. This line of thinking is vindicated by the asymmetry analysis we conducted, whereby we showed that higher uncertainty has a stronger effect than lower values of it. This result is in line with the vast literature that tends to suggest markets move more relative to bad news (i.e., higher uncertainty) than to good news. When predictability does exist for oil returns in calmer periods of the sub-sample, it is concentrated around the median of the conditional distribution, i.e., when the oil market is in its normal mode. Lack of predictability at the lower end of the distribution, could be an indication of possible herding in the market, while at the upper end—since the market is doing well in any case—investors probably do not need any information from possible predictors (in our case uncertainty) to make their investment decisions involving oil, and the market tends to function as a random walk. From an academic’s perspective, these results tend to suggest that the oil market is weakly efficient relative to measures of

¹¹Complete details of these results are available upon request from the authors.

uncertainty, when the oil market and the economy, in general, are performing well. During episodes of turmoil, investors need information on uncertainty to make profitable investments, by predicting the future path of oil. In other words, the weak-form of efficiency in the oil market fails to hold during tumultuous periods characterized by recessions. When we focus on oil price volatility, we observe that in general, uncertainty tends to predict volatility consistently, irrespective of the phase of the oil market, though stronger effects again are observed during episodes of economic stress. In this regard, our results on volatility, in particular, highlight the importance of focusing not only on in-sample, but also on out of sample predictability. As we show, in-sample results do not necessarily translate into out-of-sample results—an important observation for the academic. Note that volatility estimates, when interpreted as uncertainty, are required by investors as a key input for investment decisions and portfolio choices. Also, volatility is the most important variable in the pricing of derivative securities, i.e., to price an option, one needs reliable estimates of its volatility. Hence, the information content on uncertainty indexes, which is shown to affect oil market volatility is of immense value to investors. From a hedging perspective however, the oil market does not seem to be a good hedge against uncertainty, and hence is not necessarily a strong instrument for diversifying portfolio risks, especially during episodes of economic distress. As the literature has shown, not only oil returns, but also oil volatility tends to have tremendous repercussions on the real economy (Elder and Serletis, 2010). Given this, policy makers should build uncertainty proxies into their forecasting models, since we show that uncertainty indexes both move oil returns and, in particular, increase oil volatility. And in the process, they should be ready to undertake appropriate monetary policy action (expansionary policy) in the wake of increased oil market volatility originating from movements in uncertainty, especially when the uncertainty is higher than normal. Our results also tend to suggest that economies can move into deeper recessions through the adverse effect of the oil market, in the event where uncertainty increases during such periods, implying that stronger policy stances would be desired from the policy makers.

4.5 A robustness check against a GARCH model

As a last exercise, we compared the methods described in Section 2 with a competitive approach, namely a model belonging to the class of the generalized autoregressive conditional heteroskedasticity (GARCH) models (Bollerslev, 1986). We adopted the model to evaluate the impact of EPU and EMU on both the mean and the variance of the oil's returns. To avoid the introduction of positivity constraints on the model parameters, we used the exponential generalized autoregressive conditional heteroskedasticity (EGARCH) model introduced by Nelson (1991). Notably, the model also allows for asymmetry.

In particular, we estimated an ARX(2)-EGARCHX(1,1) model for y_t that includes, besides the lags, also $x_{1,t-j}$ and $x_{2,t-j}$ ($j = \{1, 2\}$) as external covariates in both the mean and the variance equations. We estimated the parameters of the model and forecasted the mean ($\hat{\mu}_{t+1}^{egarch}$) and the

standard deviation ($\hat{\sigma}_{t+1}^{egarch}$) of oil's returns, by means of the same rolling window scheme we used for the quantile regression (see Section 3). First, we studied the Granger causality of EPU and EMU, by analyzing the statistical significance of the coefficients related to $x_{1,t-1}$, $x_{1,t-2}$, $x_{2,t-1}$ and $x_{2,t-2}$ in the mean and variance equations. In dealing with possible model misspecification, we first estimated the coefficients' standard errors in a robust way using the method proposed by White (1982); additionally, we forecast the quantiles of y_{t+1} as $\hat{q}_\tau(y_{t+1}^{egarch}) = \hat{\mu}_{t+1}^{egarch} + \hat{\sigma}_{t+1}^{egarch} \cdot \hat{q}_\tau(z_t^{egarch})$, where $\hat{q}_\tau(z_t^{egarch})$ is the in-sample τ -th quantile of the standardized residuals, with $\tau \in (0, 1)$. Note that such a choice allowed us to avoid making a distributional assumption, thus reducing the risk of model misspecification. We then obtained the density forecast resulting from the EGARCH model, by mimicking the method we used to forecast the density of oil's returns from the quantile regression output (see Section 2.2).

Table 10 in Appendix D reports the number and the percentage of subsamples determined by the rolling window procedure where the coefficients of $x_{1,t-j}$ and of $x_{2,t-j}$ ($j = \{1, 2\}$) are statistically significant at the 0.05 level. Starting from the mean equation, we can see that $x_{1,t-1}$ and $x_{1,t-2}$ are significant in about 70% of the cases, whereas $x_{2,t-1}$ and $x_{2,t-2}$ record a percentage of about 60%. Therefore, we have a stronger evidence of the causality impact of EPU and EMU on the y_t 's mean when using the EGARCH model with respect to both the Jeong et al. (2012) test (see the values of $\hat{J}_T^*$ in Figures 2—3) and the quantile regression method (see the p-values of the $x_{i,t-j}$'s coefficients in Table 4). The percentages in Table 10 halve when taking into consideration the variance equation. Hence, the evidence of causality is stronger with the Jeong et al. (2012) test when we focus on oil's variance (see the values of $\hat{J}_T^*$ in Figures 4—7). Results are thus mixed, but still point to the presence of a significant impact of EMU and EPU on dynamic oil returns.

When moving to the density forecast analysis, we directly compared the EGARCH and the quantile regression methods in terms of out-of-sample accuracy, using the tests of Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011) focusing on different quantile levels or on different regions of the y_t 's conditional distribution. Positive (negative) values of the test statistic point to the best performance of the quantile regression (EGARCH) in Diks et al. (2011). The opposite holds in Diebold and Mariano (2002) and in Gneiting and Ranjan (2011). We report the results in Figures 16—18 (Appendix D). The density forecasts obtained from a collection of quantile regressions overperform the EGARCH-based density forecasts in the majority of cases. The differences between the two methods are often different in a statistically significant way.

5 Concluding remarks

In this work, we checked that the relationships between oil movements and the uncertainty indexes (EPU and EMU) are affected by structural breaks. The conclusions drawn from a full sample analysis, as in Balcilar et al. (2016a), would be misleading and, therefore, we implemented a rolling window

procedure to capture the dynamics among the variables involved.

First, we showed that the impact of EPU and EMU in causing—in the Granger (1969) sense—the quantiles of oil returns changes over time. Indeed, periods characterized by low or inexistent power of the two uncertainty indexes in causing oil returns are followed by periods of relevant causality evidence. Nevertheless, despite the changing in regimes over time, the periods in which the uncertainty indexes are significant in causing oil returns are less persistent than the ones characterized by no causality. In contrast, when we focused on risk, then considering the oil volatility estimated through the realized bipower variation (bpv_t) and the squared oil returns (y_t^2), we observed stronger causality impacts, as the periods in which $x_{1,t}$ and $x_{2,t}$ are significant in causing the bpv_t and the y_t^2 quantiles are more persistent. The reason might be the following: with the squared returns of oil, we focus on the y_t volatility, a measure of dispersion (and thus uncertainty) that better fits the nature of EPU and EMU, which themselves are uncertainty indicators. Balcilar et al. (2016a,b,c) used squared returns in their causality exercise to study causality in the second moment. We went further, by considering also the causality for the realized bipower variation of oil, comparing the two approaches. We observed that the causality results for bpv_t and y_t^2 are quite similar for central quantiles; in contrast, at extreme probability levels, the two approaches lead to different conclusions.

Similar results are observed for the forecasting exercise. First, as for the in-sample estimates, the coefficients of the two uncertainty indexes are not always statistically significant, at the 5% level. Therefore, EPU and EMU are useful predictors only in some periods. Moving to the coefficients' values, on average, EPU and EMU have a negative impact on the lower quantiles of oil returns, whereas their impact becomes positive at the upper quantiles. At the median level, the magnitude of the coefficients, in absolute value, is of moderate size, highlighting that EPU and EMU are critical indicators during particular periods, such as these characterized by large movements in the oil price, with potentially relevant effects in terms of inflation or deflation. In contrast, EPU and EMU positively affect the realized volatility and the squared returns series in all the regions of their conditional distribution, with the magnitude of their impact being a positive function of the quantile level. Again, this was an expected result, as the greater the uncertainty, the greater the risk in the oil market. Consistent with the in-sample evidence, the out-of-sample performance of our model was significantly improved by EPU and EMU only in particular periods, as suggested by suitable tests, namely those proposed by Berkowitz (2001), Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011). Notably, the different testing procedures provided relevant evidence of the significant improvements due to the two indexes during the years 2005—2007 and 2008—2010. These periods are close to two special events: the '2008 oil price bubble', which spans the years 2007-2008, and the US subprime crisis marked by the Lehman Brothers' default in September 2008. Interestingly, we also observed that the out-of-sample improvements given by EPU and EMU change over time according to the different regions of oil returns' and oil risk's conditional distributions.

In this study, we built the predictive conditional densities of oil’s returns and variance from the quantile regression output (Koenker and Bassett, 1978), focusing on the impact of EPU and EMU, as evaluated with an extensive empirical analysis. We also compared our approach with an EGARCH model. We checked that the causality relationships between the two uncertainty indexes (EPU and EMU) and oil’s returns are stronger when using the EGARCH model. In contrast, the causality relationships become more evident using the quantile regression method when focusing on oil’s variance. Moreover, different testing methods point out that our approach outperforms the EGARCH model in terms of predictive accuracy. The literature includes a wide set of additional alternative methods used for building conditional densities—see, e.g., Hyndman and Yao (2002)—and a proper comparison might reveal interesting findings, such as identifying outperforming approaches. Furthermore, we could also extend the bivariate framework we use here to study Granger causality in quantiles to a larger dimension. For instance, we could estimate a VAR including autoregressive components as done by White et al. (2015). Also, our approach, just like that of Nishiyama et al. (2011) and Jeong et al. (2012) upon which we build our econometric model on, is bivariate. But given the fact that, oil market movements are likely to be driven by other predictors, besides economic and financial uncertainties, we would ideally need a multivariate model. This, we believe, is also an important area of future research, which in turn, would require us to validate our results obtained when we control for additional predictors affecting the oil market. Note that, nonlinear causality tests are in general bivariate (see for example, Hiemstra and Jones (1994), Diks and Panchenko (2005, 2006)), however, in recent papers, Bai et al. (2010, 2011) have developed a multivariate version of the nonlinear causality test, but the framework is restricted to a conditional mean-based model and analyses only causality in the first-moment. Finally, in an interesting paper, Barrero et al. (2017) highlights that the impact of short- and long-run uncertainty is likely to be different on the macroeconomy, hence, the same can be tested in the oil market by decomposing movements in uncertainty into various frequencies using wavelet. We include these directions and ideas in our agenda for future research.

To conclude, first, we have contributed to the literature by showing that only in particular periods the Economic Policy and the Market Equity Uncertainty indexes are relevant drivers of oil movements. This evidence holds for both the in- and the out-of-sample analyses. Second, their impact is different according to the fact that the focus is placed on oil returns or on oil risk, therefore it is important to take into account both the effects in order to get a complete view of the market behaviour. Finally, both the causality and the forecasting exercises support the need to go beyond the point estimates, by analyzing the entire conditional distribution of oil movements, especially the tails. Indeed, the results depend deeply on the state of the oil market, differing as it does between bearish, normal and bullish conditions. Moreover, extending the analysis to the entire conditional distribution is of relevant importance in evaluating the uncertainty of the point estimates and forecasts. All of these ingredients might be of interest to many of decisions makers in several areas of economics and finance, such as risk

management, pricing and trading, when the instruments of interest depend on the oil price or risk.

References

- Ajmi, A., Gupta, R., and Kanda, P. (2015). Causality between economic policy uncertainty across countries: Evidence from linear and nonlinear tests. *Frontiers in Finance and Economics*, 11:73–102.
- Aloui, R., Gupta, R., and Miller, S. M. (2016). Uncertainty and crude oil returns. *Energy Economics*, forthcoming.
- Andreasson, P., Bekiros, S., Nguyen, D., and Uddin, G. (2016). Impact of speculation and economic uncertainty on commodity markets. *International Review of Financial Analysis*, 43:115–127.
- Antonakakis, N., Chatziantoniou, I., and Filis, G. (2014). Dynamic spillovers of oil price shocks and economic policy uncertainty. *Energy Economics*, 44:433–447.
- Bai, J. and Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 6:72–78.
- Bai, Z., Li, H., Wong, W.-K., and Zhang, B. (2011). Multivariate causality tests with simulation and application. *Statistics & Probability Letters*, 81:1063–1071.
- Bai, Z., Wong, W.-K., and Zhang, B. (2010). Multivariate linear and nonlinear causality tests. *Mathematics and Computers in Simulation*, 81:5–17.
- Baker, S., Bloom, N., and Davis, S. (2013). Measuring economic policy uncertainty. *Chicago Booth Research Paper No.13-02*.
- Balcilar, M., Bekiros, S., and Gupta, R. (2016a). The role of news-based uncertainty indices in predicting oil markets: a hybrid nonparametric quantile causality method. *Empirical Economics*, forthcoming.
- Balcilar, M., Gupta, R., Kyei, C., and Wohar, M. (2016b). Does economic policy uncertainty predict exchange rate returns and volatility? evidence from a nonparametric causality-in-quantiles test. *Open Economies Review*, 27:229–250.
- Balcilar, M., Gupta, R., and Pierdzioch, C. (2016c). Does uncertainty move the gold price? new evidence from a nonparametric causality-in-quantiles test. *Resources Policy*, 49:74–80.
- Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 4:1–30.
- Barrero, J., Bloom, N., and Wright, I. (2017). Short and long run uncertainty. *NBER Working Paper No. 23676*.

- Baum, C. and Zerilli, P. (2016). Jumps and stochastic volatility in crude oil futures prices using conditional moments of integrated volatility. *Energy Economics*, 53:175–181.
- Bekiros, S., Gupta, R., and Paccagnini, A. (2015). Oil price forecastability and economic uncertainty. *Economics Letters*, 132:125–128.
- Belloni, A., Chernozhukov, V., Chetverikov, D., and Fernandez-Val, I. (2017). Conditional quantile processes based on series or many regressors. *ArXiv e-prints*: <https://arxiv.org/abs/1105.6154>.
- Berkowitz, J. (2001). Testing density forecasts, with applications to risk management. *Journal of Business and Economic Statistics*, 19(4):465–474.
- Bernanke, B. (2016). The relationship between stocks and oil prices. *Brookings*. Retrieved from <https://www.brookings.edu/blog/ben-bernanke/2016/02/19/the-relationship-between-stocks-and-oil-prices/>.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77:623–685.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.
- Bondell, H., Reich, B., and Wang, H. (2010). Non-crossing quantile regression curve estimation. *Biometrika*, 97(4):825–838.
- Chang, C., Caporin, M., and McAleer, M. (2016). Are the S&P 500 index and crude oil, natural gas and ethanol futures related for intra-day data? *Working paper*.
- Chernozhukov, V., Fernandez-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78(3):1093–1125.
- Christoffersen, P., Jacobs, K., and Li, B. (2016). Dynamic jump intensities and risk premiums in crude oil futures and options markets. *The Journal of Derivatives*, 24:8–30.
- Chuliá, H., Gupta, R., Uribe, J., and Wohar, M. (2017). Impact of us uncertainties on emerging and mature markets: Evidence from a quantile-vector autoregressive approach. *Journal of International Financial Markets, Institutions and Money*, 48:178–191.
- Colombo, V. (2013). Economic policy uncertainty in the us: Does it matter for the euro area? *Economics Letters*, 121:39–42.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.
- Davino, C., Furno, M., and Vistocco, D. (2014). *Quantile regression: theory and applications*. Wiley.

- Dickey, D. and Fuller, W. A. (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, 49(4):1057–1072.
- Diebold, F. and Mariano, R. (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 20(1):134–144.
- Diks, C. and Panchenko, V. (2005). A note on the hiemstra-jones test for granger noncausality. *Studies in Nonlinear Dynamics and Econometrics*, 9:1–7.
- Diks, C. and Panchenko, V. (2006). A new statistic and practical guidelines for nonparametric granger causality testing. *Journal of Economic Dynamics and Control*, 30:1647–1669.
- Diks, C., Panchenko, V., and van Dijk, D. (2011). Likelihood-based scoring rules for comparing density forecasts in tails. *Journal of Econometrics*, 163:215–230.
- Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26.
- Elder, J. and Serletis, A. (2010). Oil price uncertainty. *Journal of Money Credit and Banking*, 42:1137–1159.
- Gaglianone, W. and Lima, L. (2012). Constructing density forecasts from quantile regressions. *Journal of Money, Credit and Banking*, 44(8):1589–1607.
- Giglio, S., Kelly, B. T., and Qiao, X. (2012). Systemic risk and the macroeconomy: An empirical evaluation. *Working paper*.
- Gneiting, T. and Ranjan, R. (2011). Comparing density forecasts using threshold and quantile weighted proper scoring rules. *Journal of Business and Economic Statistics*, 29:411–422.
- Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37:424–438.
- Hamilton, J. D. (1983). Oil and the macroeconomy since world war ii. *Journal of Political Economy*, 91:228–248.
- Hamilton, J. D. (2008). Oil and the macroeconomy. in *New Palgrave Dictionary of Economics*, 2nd edition, edited by Steven Durlauf and Lawrence Blume, Palgrave MacMillan Ltd.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity*, Spring, pages 215–259.
- Hiemstra, C. and Jones, J. (1994). Testing for linear and nonlinear granger causality in the stock price-volume relation. *Journal of Finance*, 49:1639–1664.

- Hyndman, R. and Yao, Q. (2002). Nonparametric estimation and symmetry tests for conditional density functions. *Journal of nonparametric statistics*, 14:259–278.
- Jarque, C. and Bera, A. (1987). A test for normality of observations and regression residuals. *International Statistical Review*, 55:163–172.
- Jeong, K., Hardle, W. K., and Song, S. (2012). A consistent nonparametric test for causality in quantile. *Econometric Theory*, 28:861–887.
- Ji, Q. and Guo, J. (2015a). Market interdependence among commodity prices based on information transmission on the internet. *Physica A*, 426:35–44.
- Ji, Q. and Guo, J. (2015b). Oil price volatility and oil-related events: an internet concern study perspective. *Applied Energy*, 137:256–264.
- Jones, P. M. and Olson, E. (2013). The time-varying correlation between uncertainty, output, and inflation: Evidence from a dcc-garch model. *Economics Letters*, 118:33–37.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *The American Economic Review*, 105:1177–1216.
- Kang, W. and Ratti, R. A. (2013a). Oil shocks, policy uncertainty and stock market return. *International Financial Markets, Institutions and Money*, 26:305–318.
- Kang, W. and Ratti, R. A. (2013b). Structural oil price shocks and policy uncertainty. *Economic Modelling*, 35:314–319.
- Karnizova, L. and Li, J. (2014). Economic policy uncertainty, financial markets and probability of us recessions. *Economics Letters*, 125:261–265.
- Kocherginsky, M. (2003). *Extensions of the Markov chain marginal bootstrap*. PhD Thesis, University of Illinois Urbana-Champaign.
- Koenker, R. (1984). A note on l-estimators for linear models. *Statistics and Probability Letters*, 2(6):323–325.
- Koenker, R. (2005). *Quantile regression*. Number 38. Cambridge university press.
- Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.
- Larsson, K. and Nossman, M. (2011). Jumps and stochastic volatility in oil prices: Time series evidence. *Energy Economics*, 33:504–514.
- Mumtaz, H. and Zanetti, F. (2013). The impact of the volatility of monetary policy shocks. *Journal of Money, Credit and Banking*, 45:535–558.

- Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59:347–370.
- Nishiyama, Y., Hitomi, K., Kawasaki, Y., and Jeong, K. (2011). A consistent nonparametric test for nonlinear causality - specification in time series regression. *Journal of Econometrics*, 165:112–127.
- Phillips, P. C. B. and Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2):335–346.
- Qu, Z. (2008). Testing for structural change in regression quantiles. *Journal of Econometrics*, 146:170–184.
- Tillmann, P. and Wolters, M. H. (2015). The changing dynamics of US inflation persistence: a quantile regression approach. *Studies in Nonlinear Dynamics & Econometrics*, 19(2):161–182.
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50:1–25.
- White, H., Kim, T., and Manganelli, S. (2015). VAR for VaR: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics*, 187:169–188.
- Zhao, Y. (2011). Real time density forecasts of output and inflation via quantile regression. *Working paper*.

Appendix

A Evaluation of the predictive accuracy

In what follows, we give the main details on the tests we used to evaluate the predictive accuracy of the methods described in Section 2.2. Here we focus on the conditional quantiles and distribution of y_t . The same methodology applies to y_t^2 and bpv_t .

We first used the Berkowitz test for an absolute assessment of the density forecasts recovered from (3). Besides, we also implemented the test on a restricted model, which excluded the lags of $x_{1,t}$ and $x_{2,t}$ to assess the joint contribution of EPU and EMU. We also evaluated the contribution of each uncertainty index separately, by adding to the restricted model only the lagged values of $x_{1,t}$ when we focused on EPU, or the lagged values of $x_{2,t}$, when we considered EMU.

The approach proposed by Berkowitz (2001) evaluates the fitness of a specific sequence of density forecasts, relative to the unknown data-generating process. However, given a certain model, the Berkowitz test has power only for misspecifications of the first two moments, but in practice, this model could be misspecified at higher-order moments. In that case, a valid solution consists in comparing density forecasts, i.e. performing a relative comparison given a specific measure of accuracy. Hence,

in addition to the approach proposed by Berkowitz (2001), we also considered the tests introduced by Diebold and Mariano (2002), Diks et al. (2011) and Gneiting and Ranjan (2011).

We implemented the test developed by Diebold and Mariano (2002) on the basis of the losses generated by the unrestricted and the restricted models, denoted by $L_{\tau,t+1}(y_{t+1}|\mathcal{W}_t)$ and $L_{\tau,t+1}(y_{t+1}|\mathcal{Y}_t)$, respectively. We emphasize that the restricted model does not account for the possible impact of the uncertainty measures, whereas the unrestricted model includes the uncertainty measures in the information set $\mathcal{Y}_t$. Among the various loss functions adopted in the literature, following Giglio et al. (2012), we made use of those defined as follows:

$$L_{\tau,t+1}(y_{t+1}|\mathcal{W}_t) = (\tau - \mathbf{1}_{\{y_{t+1} - Q_{\tau}^*(y_{t+1}|\mathcal{W}_t) < 0\}}) [y_{t+1} - Q_{\tau}^*(y_{t+1}|\mathcal{W}_t)], \quad (6)$$

$$L_{\tau,t+1}(y_{t+1}|\mathcal{Y}_t) = (\tau - \mathbf{1}_{\{y_{t+1} - Q_{\tau}^*(y_{t+1}|\mathcal{Y}_t) < 0\}}) [y_{t+1} - Q_{\tau}^*(y_{t+1}|\mathcal{Y}_t)], \quad (7)$$

where $\mathbf{1}_{\{\cdot\}}$ is an indicator function taking a value of 1 if the condition in $\{\cdot\}$ is true and a value of 0 otherwise. In evaluating the Diebold and Mariano (2002) test statistic, we focused on the quantile levels $\tau = \{0.05, 0.5, 0.95\}$.

The following tests focus on the density forecast, thus allowing for a much broader evaluation of the relevance of the uncertainty indexes. As for the Diks et al. (2011) test we used the score function of the unrestricted model defined as:

$$S^{csl}(y_{t+1}|\mathcal{W}_t) = w_{csl,t}(y_{t+1}) \log \hat{f}(y_{t+1}|\mathcal{W}_t) + (1 - w_{csl,t}(y_{t+1})) \log \left[1 - \int w_{csl,t}(s) \hat{f}(s|\mathcal{W}_t) ds \right], \quad (8)$$

where $w_{csl,t}(\cdot)$ is a weighting function by which we focus on the density's region of interest, whereas the second addend in (8) avoids the mistake of attaching comparable scores to density forecasts that have similar tail shapes but may have completely different tail probabilities (Diks et al., 2011).

Let $\bar{y}_1$ and $\bar{y}_3$ be the in-sample first and third quartile of y_t , respectively, we set $w_{csl,t}(y_{t+1}) = \mathbf{1}_{\{y_{t+1} \leq \bar{y}_1\}}$ when we focus on the left tail, $w_{csl,t}(y_{t+1}) = \mathbf{1}_{\{\bar{y}_1 \leq y_{t+1} \leq \bar{y}_3\}}$ when we place the attention on the center of the distribution, $w_{csl,t}(y_{t+1}) = \mathbf{1}_{\{y_{t+1} \geq \bar{y}_3\}}$ when we consider the right tail. We implemented the Diks et al. (2011) test by comparing (8) with the score function of the restricted model (which excludes the lags of $x_{1,t}$ and $x_{2,t}$), that is, $S^{csl}(y_{t+1}|\mathcal{Y}_t)$. We then tested the null hypothesis of equal performance between the unrestricted and the restricted model.

Finally, again focusing on the unrestricted model, the score proposed by Gneiting and Ranjan (2011) was defined as follows:

$$S^{gr}(y_{t+1}|\mathcal{W}_t) = \frac{1}{I-1} \sum_{i=1}^I w(\tau_i) Q S_{\tau_i} \left[\hat{F}^{-1}(\tau_i|\mathcal{W}_t), y_{t+1} \right], \quad (9)$$

where $\tau_i = i/I$ and

$$QS_{\tau_i} \left[\hat{F}^{-1}(\tau_i | \mathcal{W}_t), y_{t+1} \right] = 2 \left[\mathbf{1}_{\{y_{t+1} < \hat{F}^{-1}(\tau_i | \mathcal{W}_t)\}} - \tau_i \right] (\hat{F}^{-1}(\tau_i | \mathcal{W}_t) - y_{t+1}). \quad (10)$$

It is interesting to observe that the quantity defined in (10) is similar to the one in (6); nevertheless, the loss given in (9) is more informative than $L_{\tau, t+1}(y_{t+1} | \mathcal{W}_t)$, as it is equal to the weighted average of several $QS_{\tau_i} \left[\hat{F}^{-1}(\tau_i | \mathcal{W}_t), y_{t+1} \right]$ values computed for a sufficiently large grid of probabilities levels.

As for the weight function, as suggested by Gneiting and Ranjan (2011), we set $w(\tau_i) = \tau_i(1 - \tau_i)$, $w(\tau_i) = \tau_i^2$, $w(\tau_i) = (1 - \tau_i)^2$ to assign greater importance to the center, the right tail and the left tail of the distribution, respectively. Similarly, we denote the score arising from the restricted model as $S^{gr}(y_{t+1} | \mathcal{Y}_t)$; we stress that we obtained the score by replacing $\mathcal{W}_t$ by $\mathcal{Y}_t$ in (10). Again, the test evaluates the null hypothesis of zero average score differentials.

B Data description

We denote by $\{y_t\}_{t \in T}$ the series of oil returns, that is, $y_t = \log(oil_t) - \log(oil_{t-1})$, where oil_t is the spot price of the West Texas Intermediate (WTI) crude oil at day t . $\{oil_t\}_{t \in T}$ is not stationary: both the augmented Dickey and Fuller (1981) and the Phillips and Perron (1988) tests do not reject the null hypothesis of unit root with p-values of 0.2623 and 0.2112, respectively; differently, the p-values of the two tests are less than 0.01 for both $\{y_t\}_{t \in T}$ and $\{y_t^2\}_{t \in T}$.

EPU and *EMU* are two indexes measuring the US economic policy and equity market uncertainty. *EPU* is built from the newspaper archives of the Access World New's NewsBank service, by restricting the focus on the United States and referencing the number of articles containing at least one of the terms belonging to three sets. The first set is 'economic/economy', the second is 'uncertain/uncertainty' and the third set is 'legislation/deficit/regulation/congress/federal reserve/white house'. Using the same news source, *EMU* is built from articles containing the terms previously mentioned and one or more of the following: "equity market/equity price/stock market". From *EPU* and *EMU* we computed their logarithms, which we denote as $\{x_{1,t}\}_{t \in T}$ and $\{x_{2,t}\}_{t \in T}$ respectively; these are not affected by unit root: in both the cases the p-values of the augmented Dickey and Fuller (1981) and the Phillips and Perron (1988) tests are less than 0.01.

As for the daily oil volatility, we focused on 5-minute data for the WTI crude oil future; we worked on the continuous series by rolling contracts on the basis of their daily volume. We considered here the 5 minutes' data recorded from 8:00 a.m. to 3:00 p.m. of each trading day from January 2, 2007 to April 23, 2015.¹², that is the time interval when most of the trading activities take place—see Chang et al. (2016).

¹²The time interval for which we estimate the oil volatility is due to the availability of data. The dataset was recovered from TickData.

We estimated the oil volatility by means of the realized bipower variation (Barndorff-Nielsen and Shephard, 2004), defined in (11):

$$bpv_t = \mu_1^{-2} \frac{M}{M-1} \sum_{i=2}^{M-1} |oil_{t,i}| |oil_{t,i+1}|, \quad (11)$$

where $oil_{t,i}$ is the i -th intra-day price of oil on day t , M is the number of intra-day intervals into which the trading day t is divided, $\mu_1 = \sqrt{2/\pi}$. The realized bipower variation in (11) is a robust-to-jumps estimator of the integrated volatility.

We are aware that the volatility of the spot oil might differ from the volatility of the oil future. However, given the availability of the oil future high frequency database, we decided to introduce this additional variable into the analyses to evaluate whether the oil volatility future has an informative content different from that of the squared oil returns, when pointing at the analysis of the causality in risk. This corresponds with the use of oil future realized volatility as a proxy for spot oil volatility.

Table 6: Descriptive statistics

Variable	Mean	St. Deviation	Min	Max	skewness	kurtosis	JB
y_t	1.0396e-04	0.0248	-0.4069	0.1924	-0.7639	18.3329	75632.00 (0.0000)
y_t^2	6.1412e-04	0.0026	0.0000	0.1655	37.9947	2294.8710	1.675e+09 (0.0000)
$x_{1,t}$	4.3665	0.6776	1.2185	6.5780	-0.2679	3.2726	115.15 (0.0000)
$x_{2,t}$	3.8459	1.0575	1.5688	7.8655	0.2718	2.7157	119.87 (0.0000)
bpv_t	2.9742e-04	4.1933e-04	0.0000	0.0044	3.9251	22.8180	40883.00 (0.0000)

The table reports some descriptive statistics computed for y_t , y_t^2 , $x_{1,t}$, $x_{2,t}$ and bpv_t . From left to right we report the mean, the standard deviation, the minimum and maximum values, the skewness and the kurtosis indexes and the Jarque and Bera (1987)'s test statistic JB (in brackets we report the corresponding p-values). bpv_t has a lower sample size with respect to the other variables due to the availability of the data.

We report in Table 6 some descriptive statistics computed for the variables described above. y_t and y_t^2 have average values close to zero, with standard deviations equal to 0.0248 and 0.0026, respectively; y_t ranges from -0.4069 to 0.1924 and its distribution is affected by negative skewness and leptokurtosis. Both y_t^2 and bpv_t have strong positive skewness and leptokurtosis, due to the presence, in their right tails, of relevant extreme values, mainly those detected in the time interval between the mid-2008 and the mid-2009. The uncertainty indexes, $x_{1,t}$ and $x_{2,t}$, are centered around 4.366 and 3.8459, with standard deviations equal to 0.6776 and 1.0575, respectively. Their distributions are slightly skewed, quite mesokurtic and affected by the presence of a few extreme values in the tails. Notably, in the period between January 2, 2007 and April 23, 2015, bpv_t and y_t^2 have a linear correlation coefficient equal to 0.48. That moderate value might be due to both the differences between the volatility of spot and future, as well as to the fact that the squared returns are a noisy proxy of the true (and unknown) variance. Finally, the null hypothesis of the Jarque and Bera (1987)'s test is rejected with low p-values. The statistics analyzed above point out heavy-tailed distributions of the variables of interest, with the presence of extreme values, mainly y_t and y_t^2 , suggesting the wisdom of using the

quantile regression (Koenker and Bassett, 1978) in the forecasting exercise, rather than the ordinary least squares approach, because the latter does not guarantee robust results in the presence of outliers.

C Structural breaks in quantiles and distributions

The tests proposed by Qu (2008) are subgradient and have good properties also in small samples. The tables containing the critical values of the DQ and the SQ tests are available in Qu (2008). The output of the DQ test, applied to the conditional quantiles and distribution of y_t is given in the left panel of Table 7. The number of breaks, detected at the level of 0.01, is equal to 7. The results of the SQ test are given in the right panel of Table 7; here, we can see that the breaks mainly affect the extreme conditional quantiles of y_t , rather than the central ones.

Table 7: Structural breaks in the conditional distribution and quantiles of y_t

Structural breaks in the conditional distribution		Structural breaks at specific quantiles		
Dates of breaks	DQ	$SQ (\tau = 0.1)$	$SQ (\tau = 0.5)$	$SQ (\tau = 0.9)$
20/03/1987	1.0733	2.7402 ***	1.2612	2.4531 ***
11/05/1989	1.0852	1.6406 **	1.5053	2.2918 ***
21/09/1990	1.0644	1.2656	1.5886 *	1.7786 **
05/11/1991	1.0826	2.3462 ***	1.2124	2.2366 ***
04/09/2000	1.0522	3.0317 ***	1.2306	2.6934 ***
16/08/2013	1.0598	2.1177 ***	1.4521	2.3642 ***
27/01/2015	1.1041	2.2104 ***	1.4634	1.3431

The table reports the output of the DQ and the SQ tests, introduced by Qu (2008). The former detects the presence of structural breaks in the conditional distribution of y_t at the level of 0.01, whereas the latter detects the presence of structural breaks at specific quantiles, namely at $\tau = \{0.1, 0.5, 0.9\}$; *, ** and *** refer, respectively, to the 10%, 5% and 1% significance level.

Similarly, we show the results of the two tests, arising from the estimation of the bpv_t conditional quantiles and distribution, in Table 8. The number of breaks is equal to 7 and are always present at $\tau = 0.5$. As regards the left tail ($\tau = 0.1$), two breaks are detected, respectively, in the second half of 2009 and at the beginning of 2015. Differently, four breaks are observed at $\tau = 0.9$, and one of them occurs in the period of the U.S. subprime crisis with high significance.

Table 8: Structural breaks in the conditional distribution and quantiles of bpv_t

Structural breaks in the conditional distribution		Structural breaks at specific quantiles		
Dates of breaks	DQ	$SQ (\tau = 0.1)$	$SQ (\tau = 0.5)$	$SQ (\tau = 0.9)$
16/10/2007	1.0619	1.5044	1.7023 **	1.4586
21/01/2009	1.0618	1.3262	1.5846 *	1.8923 ***
05/09/2011	1.0639	1.5819 *	1.5888 *	1.5158 *
10/04/2013	1.0602	1.2237	1.6981 **	1.6945 **
22/01/2014	1.0868	1.3855	1.9229 ***	1.5740 *
28/07/2014	1.0725	0.9388	1.6247 *	1.2472
19/01/2015	1.1091	2.1645 ***	1.7316 **	1.3013

The table reports the output of the DQ and the SQ tests, introduced by Qu (2008). The former detects the presence of structural breaks in the conditional distribution of bpv_t at the level of 0.01, whereas the latter detects the presence of structural breaks at specific quantiles, namely at $\tau = \{0.1, 0.5, 0.9\}$; *, ** and *** refer, respectively, to the 10%, 5% and 1% significance level.

Finally, Table 9 displays the case of y_t^2 . The number of breaks is equal to 13 and the structural changes occur mainly at medium-high levels of τ . In the years 2007—2015 (that is, the time period for which we also computed bpv_t) the null hypothesis of the DQ test is rejected three times at the level of 0.01; it is important to notice that two of those three rejections (the ones detected in the first half of 2009 and at the beginning of 2015) occur almost simultaneously to two of the breaks affecting the conditional distribution and quantiles of bpv_t . Therefore, starting from analyzing the presence of structural changes over time, we can find some similarities between bpv_t and y_t^2 .

Table 9: Structural breaks in the conditional distribution and quantiles of y_t^2

Structural breaks in the conditional distribution		Structural breaks at specific quantiles		
Dates of breaks	DQ	SQ ($\tau = 0.1$)	SQ ($\tau = 0.5$)	SQ ($\tau = 0.9$)
08/09/1986	1.0669	1.4449	1.7797 **	1.4096
08/12/1988	1.0630	1.3808	2.0024 ***	1.5430 **
09/10/1989	1.0599	1.2835	1.7168 **	1.5511 *
26/09/1990	1.1024	0.9704	1.3059	2.1648 ***
03/07/1991	1.0663	0.9263	2.0506 ***	1.9799 ***
27/01/1993	1.1142	0.9083	2.0441 ***	1.4157
16/03/1994	1.0746	1.4506	1.3358	1.7598 **
13/12/1994	1.0769	0.8698	1.8874 ***	0.8569
24/06/1996	1.0593	1.5555 *	1.7563 **	1.8971 ***
26/03/1999	1.0827	1.2611	1.0853	1.6581 *
18/05/2009	1.0926	1.3832	1.8667 ***	2.6657
04/12/2012	1.1001	0.9915	1.4124	1.4504
24/02/2015	1.0721	1.6582 **	1.4334	1.5175 *

The table reports the output of the DQ and the SQ tests, introduced by Qu (2008). The former detects the presence of structural breaks in the conditional distribution of y_t^2 at the level of 0.01, whereas the latter detects the presence of structural breaks at specific quantiles, namely at $\tau = \{0.1, 0.5, 0.9\}$; *, ** and *** refer, respectively, to the 10%, 5% and 1% significance level.

D Additional results

We provide additional results in Table 10 and in Figures 16—18.

Table 10: Statistical significance of EPU and EMU in the EGARCH model

	MEAN EQUATION		VARIANCE EQUATION	
variable	num.	perc.	num.	perc.
$x_{1,t-1}$	950	66.52	487	34.10
$x_{1,t-2}$	1022	71.57	514	35.99
$x_{2,t-1}$	874	61.20	459	32.14
$x_{2,t-2}$	849	59.45	437	30.60

The table reports the number (num.) and the percentages (perc.) of subsamples determined by the rolling window procedure in which the coefficients of the variables given in the first column are statistically significant at the 0.05 level. The estimates were obtained from an AR(2)-EGARCH(1,1) model on the oil's returns series which includes as external regressors $x_{1,t-j}$ and $x_{2,t-j}$ ($j=\{1,2\}$) in the mean and in the variance equations. The rolling window procedure was applied by using a window size of 500 observations and step of five days ahead.

Figure 16: The Diebold and Mariano (2002) test statistic values over the the rolled windows. The test compares the forecasting accuracy of 2 competitive methods, that is, the quantile regression and the exponential GARCH model. Positive (negative) values of the test statistic point out the best performance of the EGARCH (quantile regression) model. The test is applied at three different τ values: 0.05 (green lines), 0.50 (yellow lines) and 0.95 (blue lines). The black horizontal lines point out the 5% confidence bounds.

Figure 17: The Diks et al. (2011) test statistic values over the the rolled windows. The test compares the forecasting accuracy of two competitive methods, that is, the quantile regression and the exponential GARCH model. Positive (negative) values of the test statistic point out the best performance of the quantile regression (EGARCH) model. The test was applied by placing greater emphasis on the center (green lines), on the right tail (blue lines) and on the left tail (yellow lines) of the conditional distributions. The black horizontal lines point out the 5% confidence bounds.

Figure 18: The Gneiting and Ranjan (2011) test statistic values over the the rolled windows. The test compares the forecasting accuracy of two competitive methods, that is, the quantile regression and the exponential GARCH model. Positive (negative) values of the test statistic point out the best performance of the EGARCH (quantile regression) model. The test was applied by placing greater emphasis on the center (green lines), on the right tail (blue lines) and on the left tail (yellow lines) of the conditional distributions. The black horizontal lines point out the 5% confidence bounds.