Article

El Niño, La Niña, and the Forecastability of the Realized Variance of Heating Oil Price Movements

Mehmet Balcilar [1], Elie Bouri [2], Rangan Gupta [3,*] and Christian Pierdzioch [4]

1 Department of Economics, Eastern Mediterranean University, Famagusta 99268, Turkey; mehmet@mbalcilar.net
2 Adnan Kassar School of Business, Lebanese American University, Beirut P.O. Box 13-5053, Lebanon; elie.elbouri@lau.edu.lb
3 Department of Economics, University of Pretoria, Pretoria 0002, South Africa
4 Department of Economics, Helmut Schmidt University, 22008 Hamburg, Germany; macroeconomics@hsu-hh.de
* Correspondence: rangan.gupta@up.ac.za

Abstract: We use the heterogenous autoregressive (HAR) model to compute out-of-sample forecasts of the monthly realized variance (RV) of movements of the spot and futures price of heating oil. We extend the HAR–RV model to include the role of El Niño and La Niña episodes, as captured by the Equatorial Southern Oscillation Index (EQSOI). Using data from June 1986 to April 2021, we show evidence for several model configurations that both El Niño and La Niña phases contain information useful for forecasting subsequent to the realized variance of price movements beyond the predictive value already captured by the HAR–RV model. The predictive value of La Niña phases, however, seems to be somewhat stronger than the predictive value of El Niño phases. Our results have important implications for investors, as well as from the perspective of sustainable decisions involving the environment.

Keywords: El Niño-Southern Oscillation; realized heating oil price volatility; heterogenous autoregression; forecasting

JEL Classification: C22; Q41; Q54

1. Introduction

As pointed out by Trenberth et al. [1], the El Niño Southern Oscillation (ENSO) is an irregularly periodic variation in winds and sea surface temperatures over the tropical eastern Pacific Ocean, which, in turn, impacts the climate of much of the tropical and subtropical zones. In this regard, the warming phase of the sea temperature is known as El Niño, while the cooling phase is called La Niña, with the two phases lasting several months each and typically occurring every few years with varying intensities. Because El Niño phases tend to be warmer and drier overall, people will require less heating and more cooling, while the opposite is true for La Niña phases. As a result, El Niño and La Niña phases are likely to cause fluctuations in the price of heating oil, given its crucial role as a source of heat in the winter months. Not surprisingly, while analyzing the role of the ENSO on the prices of oil, gasoline, heating oil, diesel fuel, jet fuel and propane, [2] detect the strongest impact for heating oil.

Against the backdrop, we build on this line of research by verifying whether this evidence of the impact of ENSO on the first moment of the price of heating oil can be extended to out-of-sample forecasting of its second moment (that is, the variance of price movements), which is likely to be affected via variable demand for heating oil during El Niño and La Niña phases, associated with extreme warm and cold temperatures during a year. In this regard, we can further hypothesize that with the La Niña events associated

with extreme cold weather, they are likely to have a relatively stronger influence on the volatility of heating oil than the corresponding El Niño phases. Moreover, the ENSO has been held (partially) responsible for some of the world's greatest humanitarian disasters, with [3,4] documenting the global El Niño droughts of 1876–1878, 1888–1891, and 1896–1902, which contributed to the deaths of between 30 and 60 million people in India, China and Brazil over the 26 year period. Moreover, these extreme weather conditions can constrain the supply of rain-driven agricultural commodities and induce food price and generalized inflation [5–8], which may trigger social unrest as shown empirically by [9] and both theoretically and empirically in [10]. In this regard, [11–14] show that El Niño and La Niña shocks may even have played a role in a substantial number of civil conflicts. In this sense, the ENSO cycle can serve as a proxy for rare disaster risks [15], which in turn contribute to jump risks [16,17] that account for a large part of the volatility process of asset and commodity prices [18–20]. Naturally, one can expect an impact on the variance of price movements of heating oil via the volatility jumps channel as well, over and above the heating oil demand fluctuations route due to El Niño and La Niña events.

Yet, from a statistical perspective, it is well-established that in-sample predictability does not necessarily guarantee out-of-sample predictability gains emerging from a specific predictor, and it is in fact the latter approach that often provides a more robust test of the appropriateness of an econometric model [21]. Hence, rather than focusing on in-sample predictions, we use an out-of-sample forecasting experiment to test the predictive value of El Niño and La Niña for the variance of movements of the price of heating oil. At the same time, heating oil is second only to gasoline in terms of products derived from crude oil, and given that commodity market participants need real-time volatility forecasts for effective investment, hedging risk and arbitrage strategies, the economic importance of the question we ask in this paper cannot be overemphasized. This is more so given the close linkage between energy-based commodities and the financial markets in the wake of the process of financialization of the former, and with investors looking for alternative safe investment assets following the Global Financial Crisis (see, for example, [22,23] for detailed discussions in this regard).

Given this, we forecast the monthly realized variance (RV) of movements of the spot and future heating oil price using an extended version of the widely-studied heterogeneous autoregressive (HAR)–RV model [24]. We extend the model to incorporate the role of El Niño and La Niña phases that occurred during the period from June 1986 to April 2021. In this regard, it is important to note that measuring variance in terms of RV, defined as the sum of squared daily returns of prices observed over a given month (see [25]), provides an observable and unconditional metric of variance, which is otherwise a latent process. Accordingly, we differ from the existing literature (see for example, [26–32]) on modeling and forecasting heating oil volatility based on various types of univariate GARCH models, under which the conditional variance is a deterministic function of model parameters and historical data, and, hence, is not model-free as in the case of RV. Moreover, as discussed in [33,34], the benchmark HAR–RV model captures the generally observed long-memory and multi-scaling properties of volatility, despite having a simplistic structure. Notably, one important feature of the HAR–RV model is that it employs volatilities from different time resolutions to forecast the realized heating oil price volatility. Accordingly, the model captures the main idea motivating the heterogeneous market hypothesis (EMH, [35]) in a simple and straightforward way. The EMH indicates that different groups of market participants populate financial markets (that is, in our case, the market for heating oil), and that these groups differ in their sensitivity to information flows at various time horizons. For example, in the energy market, speculators and traders are concerned with short-term investment horizons, whereas institutional investors are concerned with long-term investment horizons.

At this stage, to emphasize our contribution, it is important that we provide some further details on the literature cited above dealing with heating oil volatility forecasting, with these studies also simultaneously often dealing with crude oil, natural gas and gasoline. While in what follows we present a brief outline of the existing literature on forecasting heating oil price volatility, the reader is referred to [36] for a more detailed discussion of the literature on volatility forecasting of multiple energy commodities.

In this regard, Sadorsky cite26 showed that within the GARCH family models, the threshold GARCH (TGARCH) framework proved to be the most appropriate for the volatility of heating oil futures returns over the daily periods from 1988 to 2003. Nomikos and Pouliasis [27] compared the forecasting ability of a Markov-Switching (MS)–GARCH model to the basic GARCH on a dataset containing information about daily futures prices from 1991 to 2008 for heating oil, to depict the dominance of the former due to its ability to account for the special volatility characteristics. Arouri [28], while analyzing both daily spot and futures prices from 1986 to 2011, highlighted the role of a Fractionally Integrated GARCH (FIGARCH) model in its ability to better describe the time-varying volatility process, as it can capture long-memory (which decreases after incorporating structural breaks), a phenomenon we also account for via the HAR–RV model. Given this, the FIGARCH model (suitably modified to account for structural breaks) is shown to perform better than any other related model in forecasting heating oil volatility. Along similar lines, ref. [29] showed that the Autoregressive Fractionally Integrated Moving Average (ARFIMA)-FIGARCH model dominated various other GARCH-class models in forecasting volatility for heating oil derived from daily futures data covering 1995–2012. Different to the existing studies, ref. [31] pointed out that the unbiased extreme value volatility estimator, namely the Add RS, when incorporated into an ARFIMA model of heating oil futures volatility, outperforms the alternative models from the GARCH family over the daily period from 1996 to 2013. In a similar vein, in terms of methodological innovation, ref. [32] compared support vector regression (SVR) with the GARCH models, and showed the superiority of the former approach in forecasting the volatility of heating oil futures based on daily data from 2015 to 2019. However, this study also indicated that, within the GARCH models, asymmetric versions tend to perform better—an observation drawn earlier by [26]. Interestingly, while all the above studies relied on univariate models, ref. [30] stressed the need to look at multivariate (the Full BEKK [37], the Diagonal BEKK and Scalar BEKK) models while forecasting weekly spot prices of heating oil, as these frameworks capture the time-varying correlation between the returns processes of the energy-based commodities.

Given the brief description of the academic literature, it appears that our current paper is the first paper that evaluates the out-of-sample forecasting power of the El Niño and La Niña phases for the realized variance of heating oil price movements, which are likely to be affected via demand and volatility jumps channels, based on an extended HAR–RV model. Our focus is not necessarily to contribute to the technicality of volatility forecasting models, but rather to consider the usage of innovative predictors, namely the El Niño and La Niña phases.

Section 2 provides the dataset. Section 3 describes the HAR–RV model. Section 4 presents the results from our forecasting experiment. Section 5 concludes the paper.

2. Data

We use daily spot and futures prices of heating oil (U.S. dollars and cents per gallon) as represented by the New York Harbor No. 2 Heating Oil Spot Price FOB and the continuous NYMEX New York HARBOR ULSD futures contract, respectively, with the price data extracted from DataStream. After having computed the continuously compounded (i.e., log) daily returns of the spot and futures prices, we compute the monthly realized variance by taking the sum of daily squared returns over a month, where the overall sample period runs from June 1986 to April 2021. Figure 1 shows the daily price data that we used to construct the monthly realized variances. Table 1 depicts the key summary statistics of

daily returns. Both returns series are slightly negatively skewed, are leptokurtic and they are slightly positively (negatively) skewed in the case of spot (futures) returns. The two top panels of Figure 2 show the monthly realized variances of the spot and futures returns of the heating oil price. While the overall pattern of RV is similar across the spot and futures data, there are also noticeable differences visible in the figure, for example, in the early 1990s, in 2000 and in 2020 at the end of the sample period.

Table 1. Summary Statistics of Returns.

Statistic	Spot	Futures
Mean	0.0002	0.0002
Standard deviation	0.0246	0.0229
Skewness	-1.4094	-1.1954
Kurtosis	37.1312	21.3331
AR(1)	0.0002	-0.0236

Note: The number of observations on returns is 9109 for both spot and futures returns.
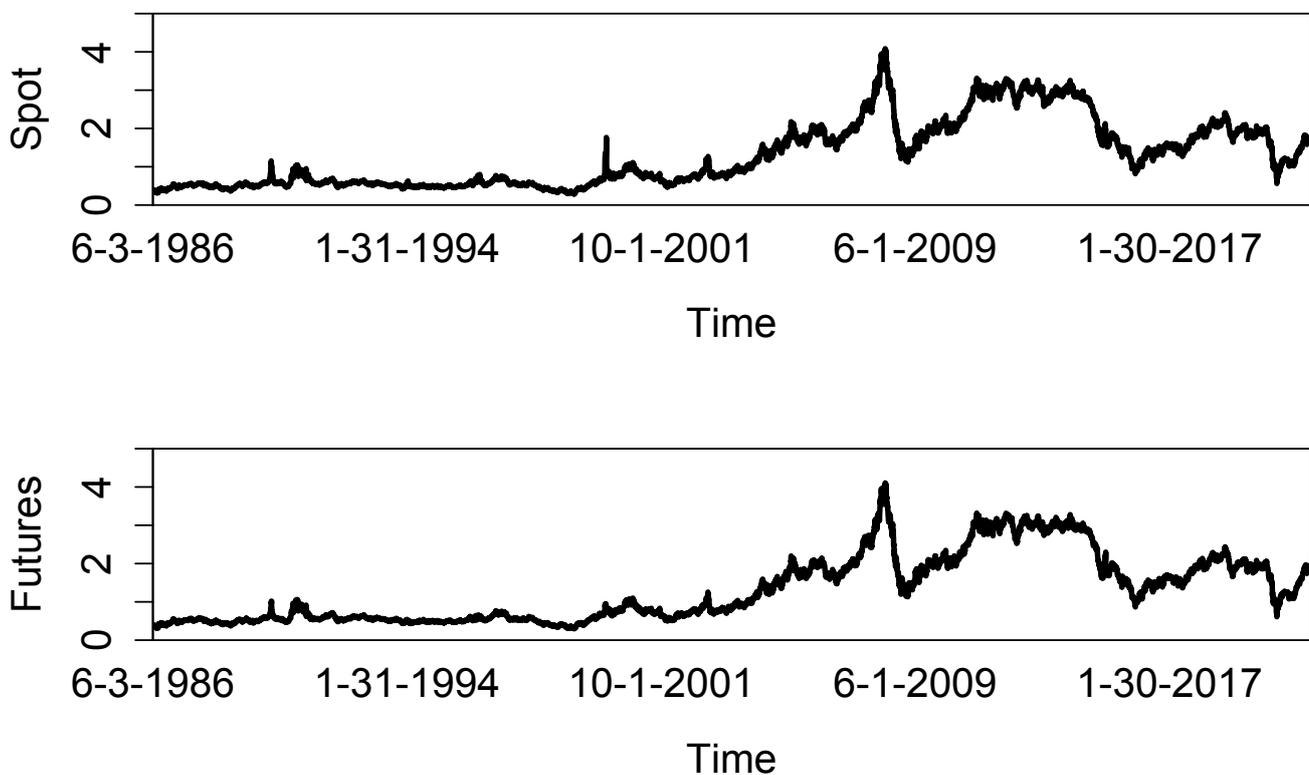
Figure 1. The Spot and Futures Data. Note: The figure shows the daily spot und futures price data (in US dollars per gallon) used to construct the monthly realized variance of movements of the price of heating oil.
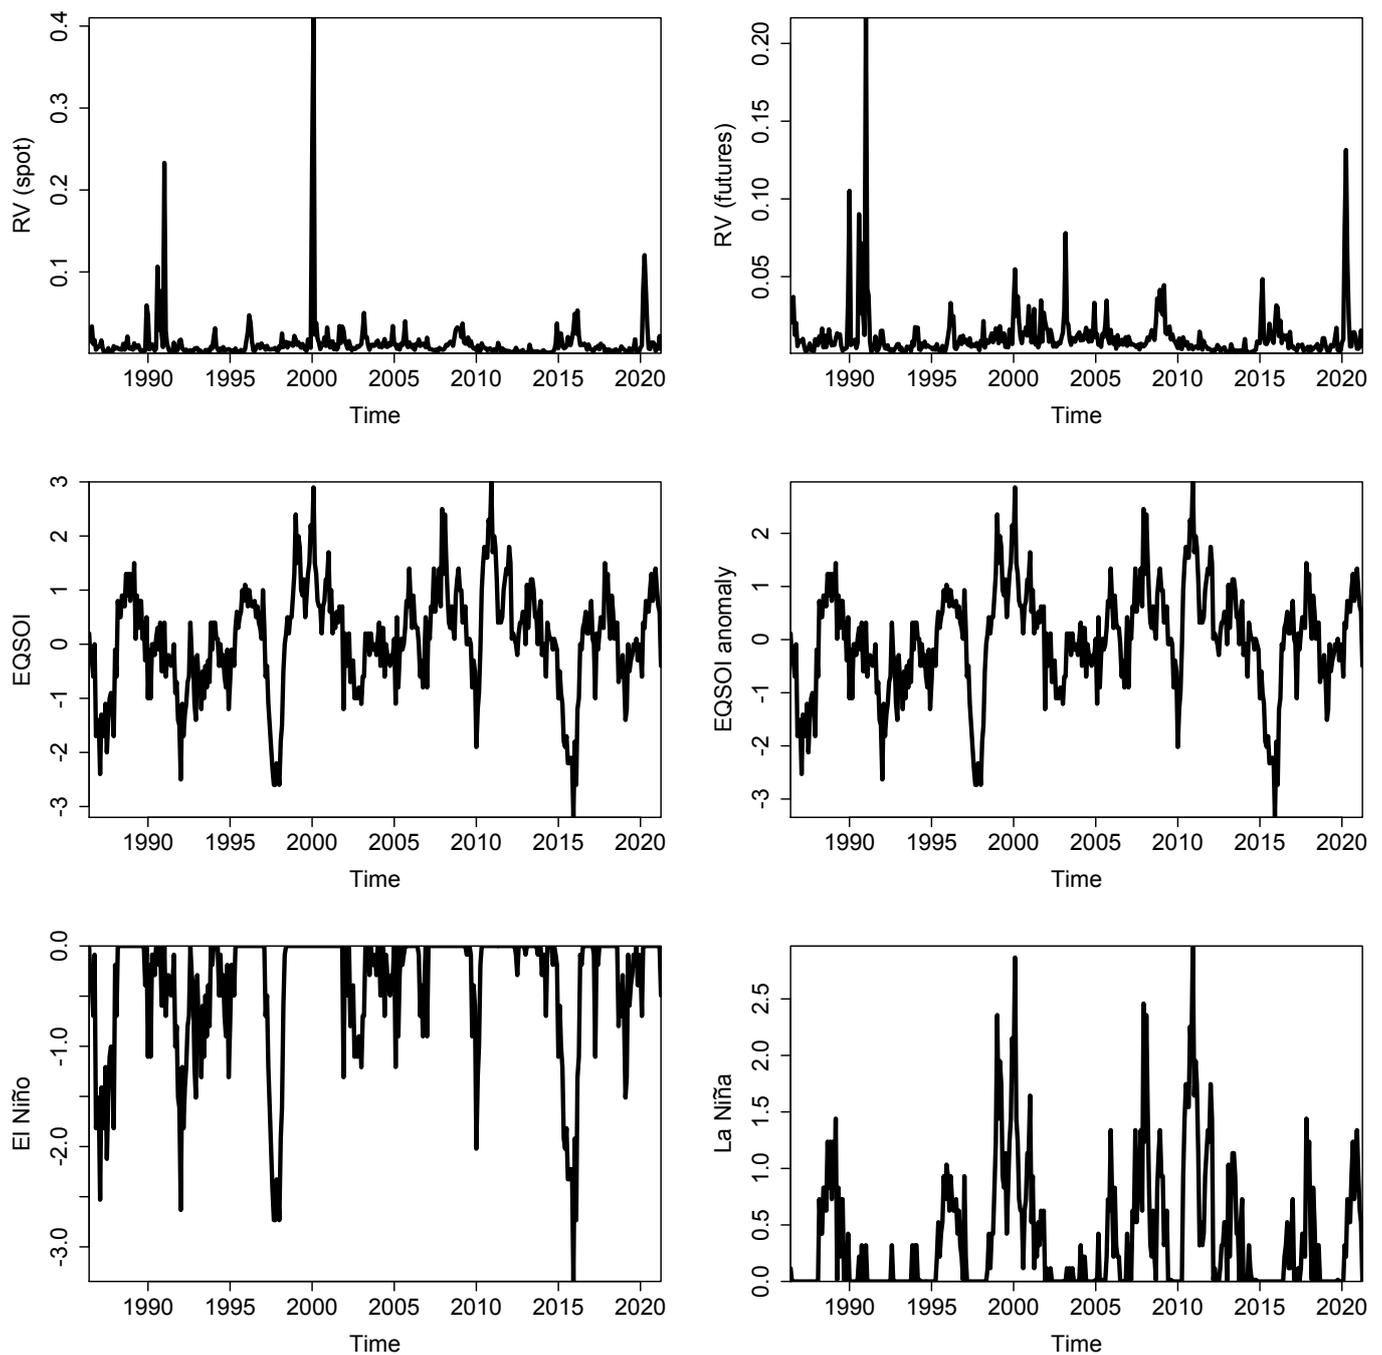
Figure 2. The *RVad* and EQSOI Data.

In order to measure the ENSO cycle, we rely on the Southern Oscillation Index (SOI) as compiled and published by the Bureau of Meteorology, Government of Australia (http://www.bom.gov.au/climate/current/soihtm1.shtml, accessed on 21 May 2021). The SOI summarizes the development and intensity of El Niño or La Niña events in the Pacific Ocean. From a methodological point of view, the SOI is computed by taking the pressure differences between Tahiti and Darwin. Such pressure differences are important because low atmospheric pressure tends to occur over warm water, while high pressure tends to occur over cold water, in part because of deep convection over the warm water. El Niño phases are, by definition, periods of sustained warming of the central and eastern tropical Pacific Ocean. La Niña phases are periods of sustained cooling of the central and eastern tropical Pacific Ocean. As such, the result is a decrease and an increase in the strength of

the Pacific trade winds, respectively. Sustained negative (positive) realizations of the SOI below (above) -7 ($+7$) usually indicate El Niño (La Niña) phases.

It should be mentioned, however, that a drawback of the SOI data is because both Darwin and Tahiti are located well south of the Equator, implying that the surface air pressure at both locations is less directly linked to ENSO. In order to address this drawback, a new index called the Equatorial Southern Oscillation Index (EQSOI) was constructed. For further details, see https://www.climate.gov/news-features/blogs/enso/why-are-there-so-many-enso-indexes-instead-just-one (accessed on 21 May 2021). The data for this index are collected at two new regions, both centered on the Equator, where the western one is located over Indonesia while the eastern one is located over the equatorial Pacific, close to the South American coast. We collect data on the EQSOI from the Climate Prediction Center of the National Oceanic and Atmospheric Administration (US Department of Commerce). See https://www.cpc.ncep.noaa.gov/data/indices/ (accessed on 21 May 2021).

Like [38], we first compute the SOI and EQSOI "anomalies", defined as the deviation of the SOI and EQSOI in any given month from its historical average, normalized (divided) by its historical standard deviation. The two middle panels of Figure 2 show the EQSOI data and the EQSOI anomaly. Equipped with the anomalies, we identify sustained negative SOI and EQSOI anomaly values below -1 (above $+1$) as indicators of El Niño (La Niña) phases. In order to capture the two phases of the ENSO cycle, we create a dummy variable which assumes the value of one whenever the SOI and EQSOI anomalies are negative (positive), and zero otherwise. Finally, we multiply the dummy variable with the SOI and EQSOI anomalies to obtain a metric for the intensity of the El Niño (La Niña) phases. In our empirical analysis, we use the El Niño and La Niña phases, depicted in the two lower panels of Figure 2, associated with EQSOI as a measure of the ENSO, given its stronger reliability. We use the El Niño and La Niña phases based on the SOI data for a robustness check.

3. The HAR–RV Model

The HAR–RV model proposed by [24] has become the workhorse model in empirical finance for the analysis of the dynamics of the realized variance (RV) of asset-price returns. While the HAR–RV model is typically applied to studying high-frequency data, we use a variant of the HAR–RV model to study the monthly realized variance (RV) of spot and futures movements of the price of heating oil. Our benchmark HAR–RV model is given by:

$$RV_{t+h} = \beta_0 + \beta_1 RV_t + \beta_2 RV_{q,t} + \beta_3 RV_{y,t} + \epsilon_{t+h}, \tag{1}$$

where h = the forecast horizon, RV_{t+h} = the average (spot or futures) realized variance observed during the forecast horizon, $RV_{q,t}$ = the average quarterly realized variance from month $t-3$ to $t-1$, and $RV_{y,t}$ = the average yearly realized variance from month $t-12$ to $t-1$). We estimate the following coefficients β_j, $j = 0, 1, 2, 3$, and ϵ_t = a disturbance term.

The predictive value of the benchmark HAR–RV model is compared to that of each of the following extended HAR–RV models:

$$RV_{t+h} = \beta_0 + \beta_1 RV_t + \beta_2 RV_{q,t} + \beta_3 RV_{y,t} + \beta_4 \text{El Niño}_t + \epsilon_{t+h}, \tag{2}$$

$$RV_{t+h} = \beta_0 + \beta_1 RV_t + \beta_2 RV_{q,t} + \beta_3 RV_{y,t} + \beta_5 \text{La Niña}_t + \epsilon_{t+h}, \tag{3}$$

$$RV_{t+h} = \beta_0 + \beta_1 RV_t + \beta_2 RV_{q,t} + \beta_3 RV_{y,t} + \beta_4 \text{El Niño}_t + \beta_5 \text{La Niña}_t + \epsilon_{t+h}. \tag{4}$$

In our empirical analysis, we focus on the out-of-sample predictive performance of our various HAR–RV models. We compute out-of-sample forecasts by estimating the models on rolling-estimation windows of different lengths. Equipped with the out-of-sample forecasts, we then evaluate the forecasts by means of the Clark and West [39] test. The null hypothesis of this test is that the two models that a researcher compares have an equal out-of-sample mean-squared prediction error (MSPE). The Clark–West test is easy to implement because it requires regressing the quantity $\hat{f}_{t+h} = (RV_{t+h} - \hat{RV}_{A,t+h})^2 - [(RV_{t+h} - \hat{RV}_{B,t+h})^2 - (\hat{RV}_{A,t+h} - \hat{RV}_{B,t+h})^2]$ on a constant. A hat over a variable marks

a forecast of *RV* and the subindices *A* and *B* denote the two models being compared (*B* denotes the larger model). We reject the null hypothesis if the t-statistic of the constant in this regression model is significantly positive (one-sided test; we use Newey–West robust standard errors to compute the t-statistic). In addition, we use the root mean squared forecasting error (RMSFE) and the mean absolute forecasting error (MAFE) statistics to compare the forecasting performance of the various HAR–RV models. We carry out our forecasting experiments using the R language for statistical computing [40].

As an extension and a robustness check (and given the results of the causality-in-quantiles test reported in Table A1 at the end of the paper), we consider a quantile-regression version of the HAR–RV model (see [41,42], among others). The quantile-regression variant of the HAR–RV model is tailored to study whether the predictive value of the El Niño and La Niña phases differs across the quantiles of the conditional distribution of the realized variance. The quantile-regression HAQR–RV model is estimated by minimizing

$$\hat{\mathbf{b}}_\alpha = \arg\min \sum_t \rho_\alpha (RV_{t+h} - \mathbf{X_t}\mathbf{b}_\alpha), \tag{5}$$

where $\rho_\alpha = \epsilon_{t+h}(\alpha - \mathbf{1}(\epsilon_{t+h} < 0))$ denotes the standard check function, α denotes a quantile, **1** denotes the indicator function, and $\mathbf{b}_\alpha$ denotes a vector of quantile-specific parameters to be estimated (a hat denotes an estimated parameter). In Equation (5), we simplify the notation by letting the matrix $\mathbf{X_t}$ denote the predictors of the HAQR–RV model.

In order to evaluate forecasts for the various quantiles of the conditional distribution of *RV*, we use an out-of-sample relative loss criterion [43,44]. The out-of-sample relative-loss criterion, $\mathcal{R}_o$, is defined as:

$$\mathcal{R}_o = 1 - \frac{\sum_{t=\tau}^{T}[\alpha + (1-2\alpha)\mathbf{1}(RV_t - \hat{RV}_{t,\alpha} < 0)]|RV_t - \hat{RV}_{t,\alpha}|^p}{\sum_{t=\tau}^{T}[\alpha + (1-2\alpha)\mathbf{1}(RV_t - \bar{RV}_{t,\alpha} < 0)]|RV_t - \bar{RV}_{t,\alpha}|^p}, \tag{6}$$

where $\bar{RV}_t$ denotes the forecast derived from the benchmark model given in Equation (1), τ is the first period for which a forecast is available. It should be noted that the benchmark forecast is also indexed by the quantile parameter, α, showing that the benchmark model is also estimated as an HARQ–RV model.

The parameter, p, governs the functional form used to compute the out-of-sample relative loss criterion. Specifically, in case we set $p = 1$, a forecaster evaluates forecast errors under a lin–lin loss function (L1; this case corresponds to the check function of the quantile regression framework), while in the case of $p = 2$, a forecaster uses a quad–quad loss function (L2) to evaluate forecast errors [45].

It should be noted that the quantile parameter, $\alpha \in (0,1)$, governs the asymmetry of the loss function. The loss function is symmetric for $\alpha = 0.5$ and is symmetric and quadratic if, in addition, $p = 2$. Setting $\alpha = 0.5$ and $p = 1$, in turn, implies that the the loss from a forecast error increases in the absolute magnitude of the forecast error. If $\alpha > 0.5$ ($\alpha < 0.5$), then the loss from underpredicting (overpredicting) *RV* exceeds the loss from an overprediction (underprediction) of the same size, that is, for $\alpha \neq 0.5$ the loss function is asymmetric.

For a given quantile parameter and a given functional form of the loss function, the HARQ–RV model, extended to include the El Niño and/or La Niña phases, outperforms the benchmark HARQ–RV model when the out-of-sample relative loss criterion is positive, $\mathcal{R}_o > 0$. When the out-of-sample relative loss criterion assumes a negative value, $\mathcal{R}_o < 0$, the benchmark model performs better than the extended model.

4. Empirical Results

Table 2 summarizes our main forecasting results. We report, for the EQSOI data, results for the CW test (p-values) for different forecast horizons ($h = 1, 3, 6, 12, 24, 36, 48$) and various lengths of the rolling estimation window (48, 60, 72, 96 and 120 months). In Panels A to C, we document the results for spot *RV*, while Panels D to F contain the results for futures *RV*. In the case of spot *RV* (Panel A), we observe that El Niño phases add to the predictive performance of the HAR–RV model when we use an intermediate window length and forecast horizons of up to two years. The test results become largely insignificant for the short and the long window length and for the longer forecast horizons. When we consider La Niña phases in Panel B, in contrast, most test results are significant. Only for a combination of a long estimation window and a short to intermediate forecast horizon are the test results not statistically significant. A similar result emerges when we combine El Niño and La Niña phases in a unified forecasting model (Panel C).

Table 2. Clark–West Test Results.

Panel A: Spot RV (El Niño)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1766	0.1141	0.1336	0.0135	0.4691	0.1135	0.1155
Window length = 60	0.0993	0.0354	0.0606	0.0319	0.0219	0.1080	0.2243
Window length = 72	0.0797	0.0462	0.0585	0.0376	0.0772	0.4138	0.7867
Window length = 96	0.1459	0.0943	0.0916	0.0875	0.1299	0.7561	0.8137
Window length = 120	0.1027	0.0960	0.1349	0.1423	0.1578	0.6694	0.7623
Panel B: Spot RV (La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0355	0.0078	0.0374	0.0180	0.3356	0.6720	0.0502
Window length = 60	0.0405	0.0581	0.0284	0.0096	0.0008	0.0107	0.0013
Window length = 72	0.0813	0.0904	0.0606	0.0290	0.0054	0.0020	0.0020
Window length = 96	0.1410	0.1247	0.1048	0.0536	0.0394	0.0299	0.0172
Window length = 120	0.1553	0.1632	0.1706	0.1034	0.0822	0.1033	0.0734
Panel C: Spot RV (El Niño and La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0191	0.0017	0.0099	0.0074	0.4377	0.0695	0.0368
Window length = 60	0.0288	0.0457	0.0102	0.0046	0.0009	0.0397	0.0186
Window length = 72	0.0650	0.0729	0.0335	0.0206	0.0050	0.0177	0.0065
Window length = 96	0.1312	0.1145	0.0837	0.0487	0.0392	0.0052	0.0040
Window length = 120	0.1369	0.1375	0.1342	0.0931	0.0740	0.0298	0.0274
Panel D: Futures RV (El Niño)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0221	0.0069	0.0159	0.0277	0.0232	0.0448	0.0685
Window length = 60	0.0459	0.0123	0.0222	0.0545	0.0683	0.0584	0.0658
Window length = 72	0.2770	0.0685	0.1112	0.1005	0.2046	0.5192	0.7192
Window length = 96	0.4088	0.0964	0.1847	0.3417	0.6243	0.9867	0.9144
Window length = 120	0.3303	0.0517	0.0616	0.0122	0.0024	0.3246	0.4821

Table 2. *Cont.*

Panel E: Futures RV (La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1697	0.0652	0.1175	0.3629	0.0340	0.0737	0.0333
Window length = 60	0.0167	0.0266	0.0168	0.0321	0.0024	0.0239	0.0072
Window length = 72	0.2276	0.0947	0.0545	0.0398	0.0039	0.0022	0.0015
Window length = 96	0.1098	0.0800	0.0501	0.0717	0.1550	0.2296	0.0582
Window length = 120	0.1188	0.1280	0.0809	0.0655	0.0411	0.0700	0.0647
Panel F: Futures RV (El Niño and La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
tabularnewline Window length = 48	0.0096	0.0008	0.0024	0.0346	0.0091	0.0245	0.0228
Window length = 60	0.0119	0.0060	0.0114	0.0419	0.0155	0.0106	0.0032
Window length = 72	0.1960	0.0479	0.0613	0.0299	0.0022	0.0016	0.0021
Window length = 96	0.1563	0.0472	0.0487	0.0365	0.1722	0.2836	0.0678
Window length = 120	0.1839	0.0555	0.0558	0.0175	0.0200	0.0558	0.0383

Note: This table reports results (*p*-values) of the Clark–West test for an equal mean squared prediction error (MSPE) for alternative forecast horizons and alternative lengths of the rolling estimation window used to compute forecasts. The baseline HAR–RV model is the benchmark model, and the HAR–RV that features El Niño and/or La Niña phases as predictors is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter *h* denotes the forecast horizon (in months). The *p*-values are based on robust standard errors.

For the futures RV, the results are slightly different. The forecasting model that features El Niño phases as an additional predictor (Panel D) only adds predictive value to the HAR–RV model for the two short rolling estimation windows. The test results for the other model configurations are statistically insignificant. For La Niña phases, in contrast, the test results are stronger (Panel E). The CW test yields significant results for all combinations of window length and forecast horizon with only a few exceptions, which are not systematically linked either to the window length or the forecast horizon. It is clear, thus, that La Niña phases are an important source of the superior forecasting performance that we find in Panel F, where we consider the forecasting model that features both El Niño and La Niña phases as additional predictors.

It is interesting to examine whether the test results are robust to a change in the definition of realized variance. To this end, we report in Table 3 the test results for the square root of realized variance, $\sqrt{RV}$, which researchers in empirical finance also often consider as a metric of realized "volatility". The results that we document in Panel A witness that the contribution of El Niño phases to the predictive performance of the HAR–RV model, in the case of spot $\sqrt{RV}$, is concentrated at the shorter forecast horizons. The overall picture that emerges, however, is that La Niña phases (Panel B) have a more systematic statistically significant effect on the forecasting performance of the HAR–RV model than on El Niño phases. Only the test results for the long rolling estimation window are not significant. We conclude that La Niña phases are a major driver of the results, especially as far as the longer forecast horizons are concerned, for the combined model in Panel C. For the futures-based $\sqrt{RV}$ (Panels D to F), the results are not too different. We observe that the contribution of El Niño phases to the forecasting performance of the HAR–RV model, in terms of statistical significance, is concentrated at the two short rolling estimation windows and at the intermediate forecast horizons. The overall picture that emerges from a comparison with the results for La Niña phases in Panel D, however, is that, with the exception of the long rolling estimation window, the test results for La Niña phases are statistically significant for a broader array of combinations of window length and forecast horizons. As a result, it is not surprising that the test results in Panel F, for a model that features both El Niño and La Niña phases as predictors, are largely significant.

Table 3. Test Results for $RV^{0.5}$.

Panel A: Spot RV (El Niño)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48

Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0304	0.0238	0.0682	0.1023	0.0544	0.1206	0.1089
Window length = 60	0.1012	0.0267	0.0837	0.1018	0.0648	0.0804	0.1197
Window length = 72	0.0803	0.0304	0.0780	0.0976	0.1459	0.3367	0.5523
Window length = 96	0.0703	0.0373	0.0717	0.1230	0.1610	0.6337	0.6159
Window length = 120	0.0606	0.0432	0.0673	0.0956	0.1230	0.8169	0.8414

Panel B: Spot RV (La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0320	0.0281	0.0073	0.0089	0.0095	0.0692	0.0320
Window length = 60	0.0131	0.0300	0.0126	0.0045	0.0002	0.0160	0.0050
Window length = 72	0.0568	0.0740	0.0551	0.0103	0.0001	0.0020	0.0011
Window length = 96	0.1485	0.1622	0.1374	0.0354	0.0188	0.0199	0.0107
Window length = 120	0.1913	0.1931	0.2123	0.1376	0.1327	0.1325	0.0926

Panel C: Spot RV (El Niño and La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0112	0.0079	0.0061	0.0038	0.0299	0.0443	0.0160
Window length = 60	0.0071	0.0139	0.0070	0.0128	0.0119	0.0146	0.0177
Window length = 72	0.0327	0.0403	0.0235	0.0057	0.0011	0.0113	0.0119
Window length = 96	0.1179	0.1284	0.0959	0.0249	0.0188	0.0223	0.0129
Window length = 120	0.1702	0.1438	0.1346	0.1113	0.1212	0.0938	0.0600

Panel D: Futures RV (El Niño)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0980	0.0106	0.0207	0.0503	0.0400	0.0463	0.0825
Window length = 60	0.1202	0.0033	0.0256	0.0819	0.0926	0.0887	0.1176
Window length = 72	0.1844	0.0219	0.0589	0.0471	0.1017	0.3484	0.5596
Window length = 96	0.2597	0.0470	0.1493	0.2718	0.3910	0.8508	0.6995
Window length = 120	0.3717	0.0831	0.0858	0.0601	0.0327	0.5982	0.6202

Panel E: Futures RV (La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0241	0.0187	0.0279	0.0936	0.0190	0.0655	0.0203
Window length = 60	0.0086	0.0208	0.0210	0.0144	0.0027	0.0264	0.0079
Window length = 72	0.0913	0.0670	0.0728	0.0269	0.0048	0.0200	0.0191
Window length = 96	0.0665	0.0748	0.0774	0.0743	0.0769	0.0958	0.0561
Window length = 120	0.1264	0.1566	0.1785	0.0965	0.0917	0.1095	0.0891

Panel F: Futures RV (El Niño and La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0295	0.0021	0.0043	0.0224	0.0181	0.0194	0.0095
Window length = 60	0.0161	0.0039	0.0069	0.0206	0.0233	0.0131	0.0080
Window length = 72	0.0787	0.0252	0.0284	0.0067	0.0063	0.0310	0.0379
Window length = 96	0.0944	0.0399	0.0590	0.0478	0.0857	0.1089	0.0552
Window length = 120	0.1770	0.0821	0.1001	0.0670	0.0874	0.0932	0.0593

Note: This table reports results (*p*-values) of the Clark–West test for an equal mean squared prediction error (MSPE) for alternative forecast horizons and alternative lengths of the rolling estimation window used to compute forecasts. The baseline HAR–RV model is the benchmark model, and the HAR–RV that features El Niño and/or La Niña phases as predictors is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter *h* denotes the forecast horizon (in months). The *p*-values are based on robust standard errors.

Table 4 summarizes the results of an additional robustness check, where we replace our EQSOI-based predictors with corresponding predictors based on the SOI data. The results corroborate our EQSOI-based results. Most test results are statistically significant and show that accounting for El Niño and La Niña phases helps to improve forecasts of *RV* of heating oil price movements at various forecast horizons.

Table 4. Test Results for SOI.

Panel A: Spot RV (El Niño)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1606	0.0285	0.0333	0.0122	0.5127	0.0322	0.0797
Window length = 60	0.0204	0.0010	0.0008	0.0021	0.0056	0.1035	0.3536
Window length = 72	0.0543	0.0156	0.0051	0.0096	0.0049	0.2092	0.1960
Window length = 96	0.1470	0.0958	0.0437	0.0520	0.0849	0.9016	0.2992
Window length = 120	0.1672	0.1758	0.1509	0.1140	0.2000	0.9645	0.7523
Panel B: Spot RV (La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0675	0.0549	0.1364	0.0119	0.1180	0.3431	0.0071
Window length = 60	0.0105	0.0216	0.0041	0.0032	0.0022	0.0001	0.0004
Window length = 72	0.0638	0.0436	0.0343	0.0316	0.0098	0.0002	0.0051
Window length = 96	0.1169	0.0981	0.0692	0.0588	0.0450	0.0299	0.0276
Window length = 120	0.0656	0.1245	0.1047	0.0400	0.0344	0.0409	0.0232
Panel C: Spot RV (El Niño and La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.2468	0.0192	0.0743	0.0182	0.1948	0.0011	0.0038
Window length = 60	0.0122	0.0106	0.0009	0.0032	0.0011	0.0004	0.0003
Window length = 72	0.0545	0.0303	0.0074	0.0156	0.0052	0.0001	0.0016
Window length = 96	0.1256	0.0951	0.0608	0.0538	0.0426	0.0194	0.0227
Window length = 120	0.1051	0.1241	0.0890	0.0628	0.0356	0.0147	0.0193
Panel D: Futures RV (El Niño)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1421	0.0776	0.0139	0.0011	0.0238	0.0411	0.0397
Window length = 60	0.0677	0.0075	0.0009	0.0013	0.0115	0.0147	0.0162
Window length = 72	0.1734	0.0596	0.0078	0.0053	0.0188	0.0244	0.0272
Window length = 96	0.2947	0.1487	0.0414	0.1266	0.6690	0.9802	0.7344
Window length = 120	0.1104	0.0412	0.0062	0.0052	0.0003	0.1895	0.3012
Panel E: Futures RV (La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.4347	0.0727	0.2163	0.3099	0.0625	0.1152	0.0574
Window length = 60	0.0861	0.0056	0.0095	0.0028	0.0070	0.0143	0.0035
Window length = 72	0.2762	0.0896	0.0600	0.0279	0.0089	0.0046	0.0013
Window length = 96	0.0469	0.0359	0.0665	0.0652	0.1159	0.0520	0.0087
Window length = 120	0.0569	0.0648	0.0742	0.0467	0.0120	0.0208	0.0164

Table 4. *Cont.*

Panel F: Futures RV (El Niño and La Niña)							
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.2481	0.1042	0.0253	0.0275	0.0129	0.0205	0.0102
Window length = 60	0.0289	0.0012	0.0005	0.0006	0.0028	0.0036	0.0017
Window length = 72	0.1368	0.0348	0.0034	0.0040	0.0037	0.0015	0.0005
Window length = 96	0.0787	0.0351	0.0243	0.0627	0.2061	0.1375	0.0198
Window length = 120	0.0986	0.0453	0.0106	0.0113	0.0083	0.0251	0.0204

Note: This table reports results (*p*-values) of the Clark–West test for an equal mean squared prediction error (MSPE) for alternative forecast horizons and alternative lengths of the rolling estimation window used to compute forecasts. The baseline HAR–RV model is the benchmark model, and the HAR–RV that features El Niño and/or La Niña phases as predictors is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter *h* denotes the forecast horizon (in months). The *p*-values are based on robust standard errors.

Our empirical findings show that both El Niño and La Niña have predictive value in an out-of-sample forecasting experiment for the realized variance (and volatility) of heating oil price movements, as measured by means of spot and futures price data. On balance, however, the predictive value of La Niña appears to be somewhat more robust in our out-of-sample forecasting experiment across the various model configurations that we have studied than the predictive value of El Niño. The results that we report in Table 5 illustrate this finding. Specifically, we report in Table 5 (for the EQSOI data in Panel A and the SOI data in Panel B) the average RMSFE statistics and the average ratios of the MAFE statistics as computed, for every model, across the five different lengths of the rolling-estimation window and the six different forecast horizons. The ratios are defined by dividing the RMSFE (MAFE) obtained for the baseline HAR–RV model by the corresponding RMSFE (MAFE) obtained for the HAR–RV model extended to include the El Niño and/or La Niña phases. Hence, a ratio larger than one indicates that the extended model produces, on average, more accurate forecasts than the baseline HAR–RV model. The punchline of the results summarized in Table 5 is that extending the the HAR–RV model to include El Niño phases tends to improve forecast accuracy, but extending the HAR–RV model to incorporate La Niña phases leads to a larger improvement of forecast accuracy. In fact, including both El Niño and La Niña phases in the forecasting model yields RMSFE and MAFE ratios that are close to the ratios reported for the forecasting model that features the La Niña phases as the only additional predictor.

Table 5. RMSFE and MAFE Ratios.

Panel A: EQSOI			
results.table	El Niño	La Niña	Both
RMSFE / Spot	0.9954	1.0137	1.0124
RMSFE / Futures	1.0012	1.0364	1.0284
MAFE / Spot	1.0057	1.0524	1.0680
MAFE / Futures	1.0070	1.0392	1.0411
Panel B: SOI			
results.table	El Niño	La Niña	Both
RMSFE / Spot	1.0026	1.0474	1.0420
RMSFE / Futures	1.0130	1.0396	1.0333
MAFE / Spot	1.0080	1.0621	1.0638
MAFE / Futures	1.0178	1.0388	1.0384

Note: This table reports the ratio of the average root mean squared forecasting errors (RMSFE) and the ratios of the mean absolute forecasting errors (MAFE). The ratios are defined by dividing the RMSFE (MAFE) obtained for the baseline HAR–RV model by the RMSFE (MAFE) obtained for the HAR–RV model extended to include El Niño and/or La Niña phases. The average RMSFE (MAFE) is computed by computing the mean RMSFE (MAFE) across the five different lengths of the rolling estimation window and the six different forecast horizons.

We next turn to a quantiles-based analysis. As for in-sample predictability, we find, based on the multivariate version of the the nonparametric *k*-th order causality-in-quantiles test of [46], which allows us to test for causality in (spot and futures) returns and squared returns (i.e., variance) simultaneously over their respective conditional distributions (thus defining the states of the heating oil market) strong evidence of in-sample predictability from both El Niño and La Niña phases, when controlling for La Niña or El Niño in the model, respectively—see the results reported in Table A1 at the end of the paper. The results further show that the evidence of in-sample predictability is not concentrated at one or a few quantiles, justifying our decision to use the HARQ–RV model to compute quantile-specific forecasts of *RV*, and to use the out-of-sample relative-loss criterion to evaluate the forecasts relative to the forecasts we obtain from the benchmark HARQ–RV model. In order to save journal space, we compare the models that feature either El Niño or La Niña phases as an additional predictor. Figures 3 and 4 summarize the results for the spot *RV* under the lin–lin and quad–quad loss function. For the lin–lin loss function, the colored areas depicted in the figure show that, according to the out-of-sample relative-loss criterion, using El Niño phases contributes to the overall forecasting performance of the HARQ–RV model mainly at the short to intermediate forecast horizons, while the areas where the benchmark HARQ–RV model performs better than the extended model become larger as the forecast horizon increases. The results for the La Niña phases, in contrast, show that the contribution of this predictor to the forecasting performance of the model increases in terms of the out-of-sample relative-loss criterion when we study the intermediate to long forecast horizons. For the quad–quad loss function, in turn, the HARQ–RV model extended to include La Niña phases as an additional predictor outperforms the benchmark HARQ–RV model in terms of the out-of-sample relative-loss criterion for most combinations of the quantile parameter and the length of the rolling estimation window for all forecast horizons being studied. A key message conveyed by the results is that the overall contribution of La Niña phases to the forecasting performance of the model is stronger than that of El Niño phases. In this regard, it should also be noted that, as indicated by the legends of the figures, the maximum gains a forecaster can realize in terms of the out-of-sample relative-loss criterion are often larger for La Niña phases than for El Niño phases.
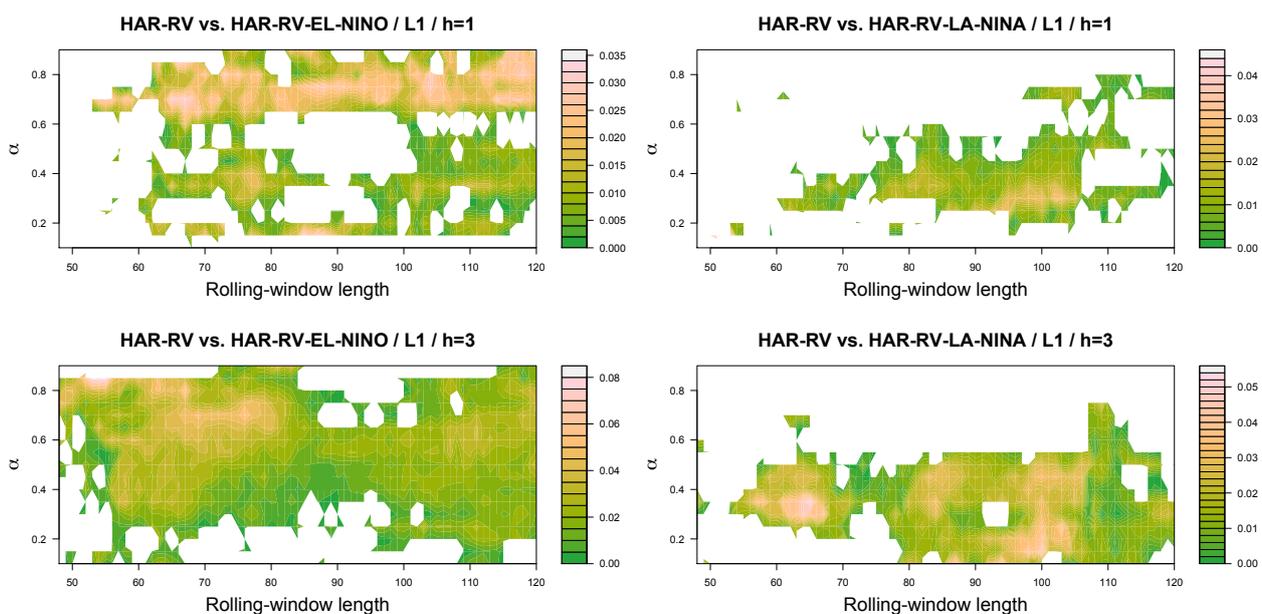
Figure 3. *Cont.*
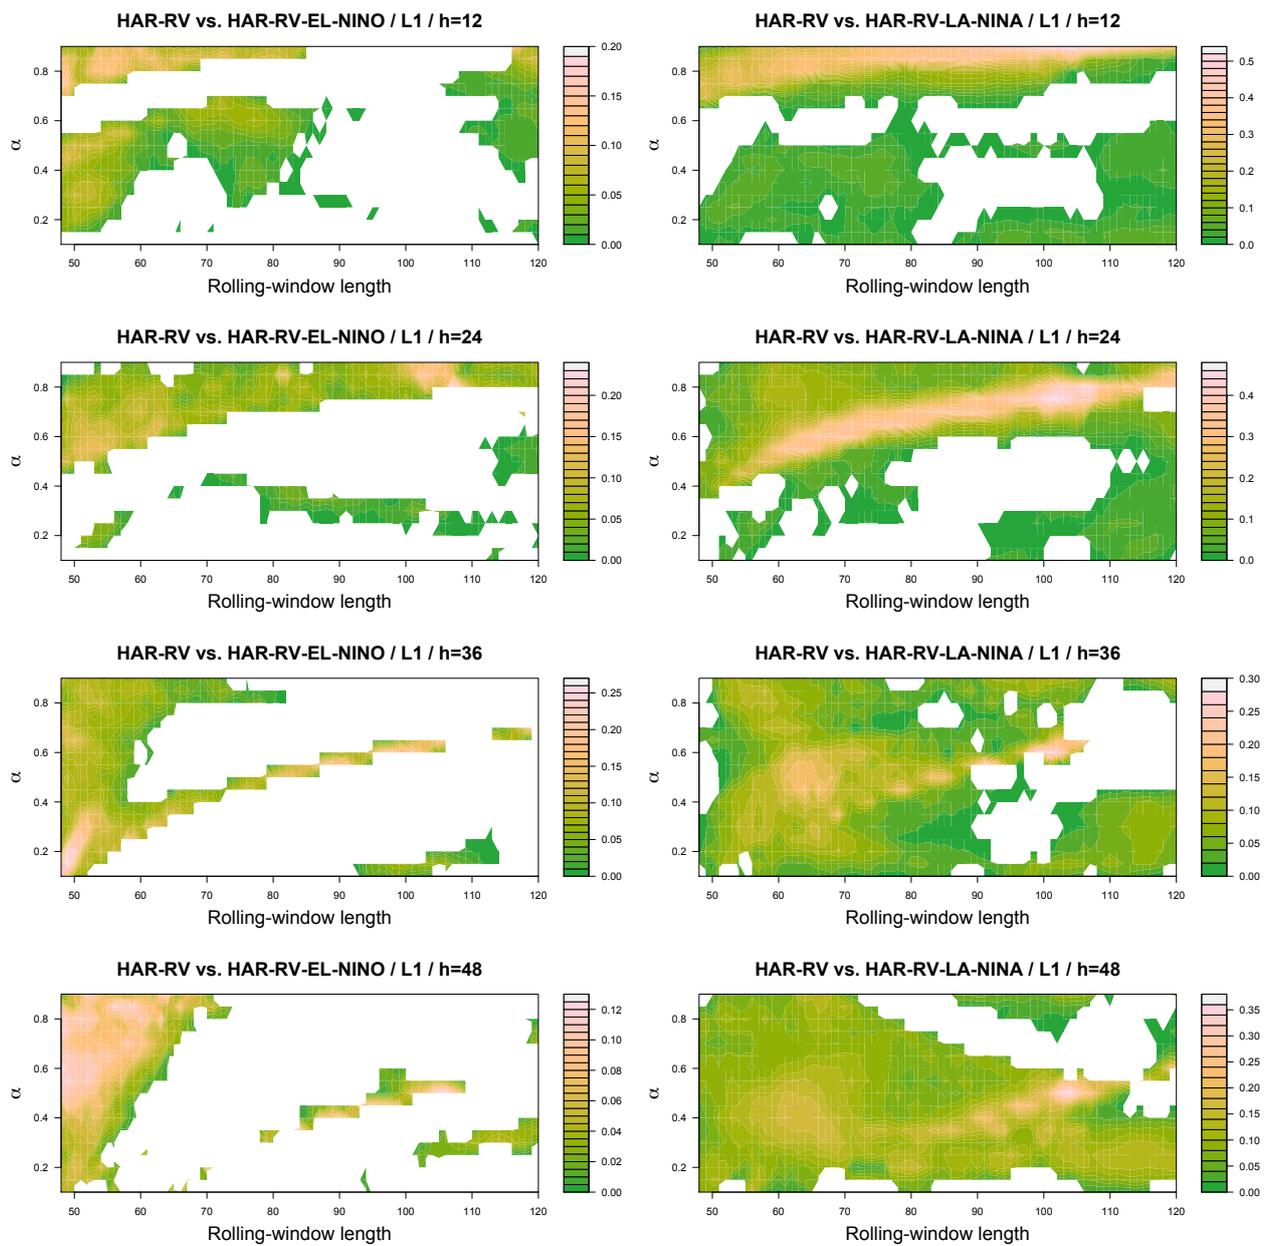
Figure 3. Results for the HARQ–RV Model (Spot; Lin–Lin Loss Function). Note: This figure reports the out-of-sample relative loss criterion for a lin–lin loss function (L1) and an extended HARQ–RV models that features either El Niño or La Niña phases as an additional predictor. The baseline model is the benchmark HARQ–RV model. White spaces indicate areas where the out-of-sample relative loss criterion takes on negative values (the baseline model performs better than the extended model). The rolling estimation window is varied from 48, 49, 50, ... to 120 observations. The results are based on the EQSOI data. The parameter *h* denotes the forecast horizon (in months).
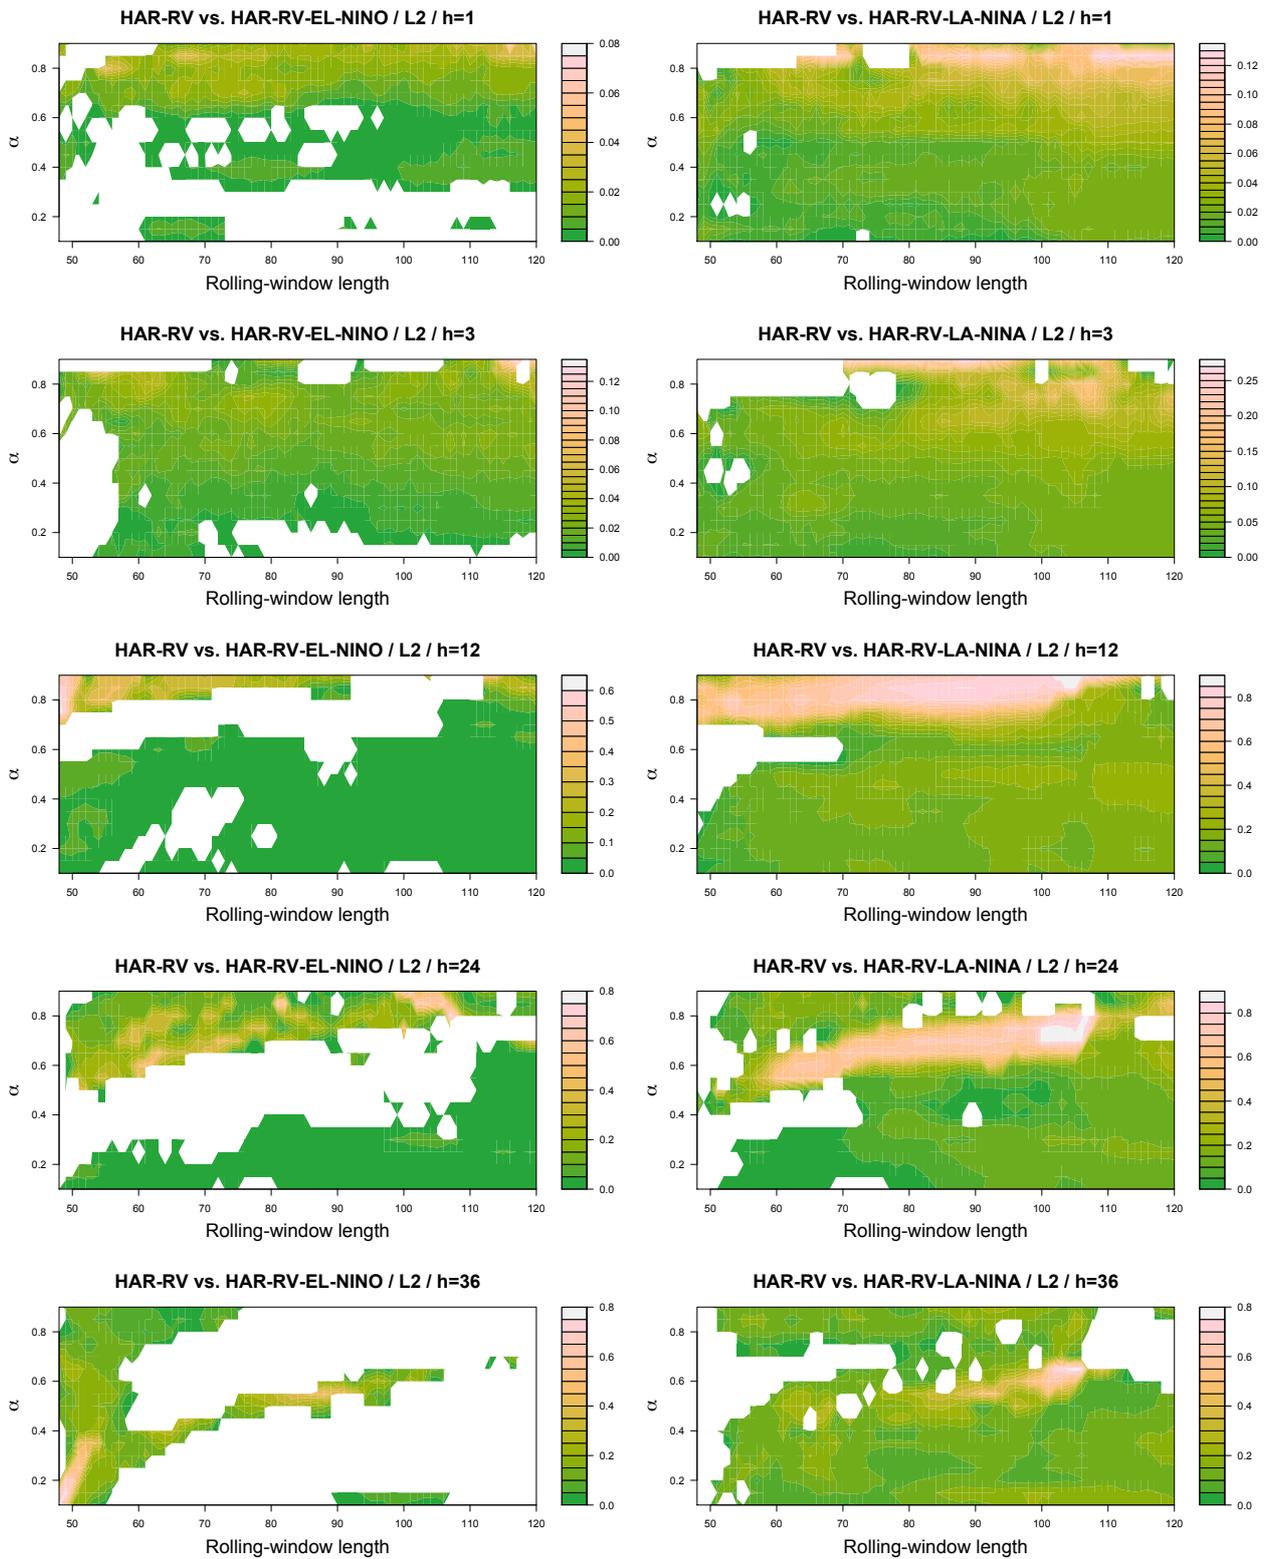
Figure 4. *Cont.*
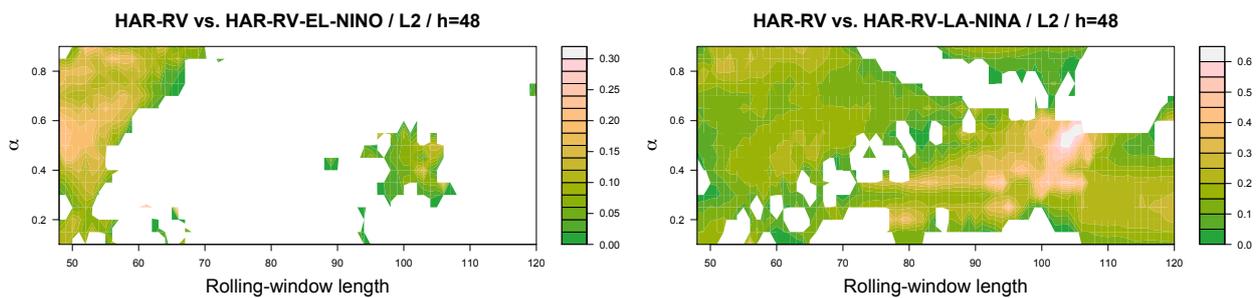
Figure 4. Results for the HARQ–RV Model (Spot; Quad–Quad Loss Function). Note: This figure reports the out-of-sample relative loss criterion for a quad-quad loss function (L2) and an extended HARQ–RV models that features either El Niño or La Niña phases as an additional predictor. The baseline model is the benchmark HARQ–RV model. White spaces indicate areas where the out-of-sample relative loss criterion takes on negative values (the baseline model performs better than the extended model). The rolling estimation window is varied from 48, 49, 50, ... to 120 observations. The results are based on the EQSOI data. The parameter *h* denotes the forecast horizon (in months).

In this regard, we note that the higher values of the statistic for the *k*-th order non-parametric causality-in-quantiles test, as reported in Table A1, confirm the finding of a stronger impact of La Niña phases relative to El Niño ones, especially at higher conditional quantiles, which corresponds to higher returns and volatility. In fact, based on the average derivative (AD) using the conditional pivotal quantile, based on the approximation or the coupling approach of [47], both El Niño and La Niña phases are found to have a positive relationship with the entire conditional distributions of returns and the volatility of heating oil price. These results are available from the authors upon request.

Figures A1 and A2 at the end of the paper (Appendix A) summarize the results for the futures *RV*. Again, we observe that the range of combinations of the quantile parameter and the length of the rolling estimation window that result in positive realizations of the out-of-sample relative-loss criterion tends to be wider for La Niña phases than for El Niño phases. Hence, to summarize, the message to take home from the results is that the out-of-sample relative loss criterion takes on positive values for many combinations of the quantile parameter, window length, and the the forecast horizon being studied, especially as far as the La Niña phases are concerned.

Finally, we present in Table 6 the results for upside ("good") and downside ("bad") *RV*, where the former is computed from positive daily returns, while the latter is computed from negative daily returns. This is an important issue, as [48] stresses that financial market participants care not only about the nature of volatility, but also about its level, with traders typically differentiating between good and bad volatilities, which, in turn, provides us the motivation to look at such disaggregation of *RV*. We present in Table 6 the results for the futures *RV*, but results for spot *RV* are qualitatively similar (and available from the authors upon request). The results show that El Niño phases add to the forecasting performance of the HAR–RV model estimated for upside and downside *RV* only for the two shortest rolling estimation windows. La Niña phases, in contrast, also carry information for several forecast horizons, which is useful for forecasting upside and downside *RV* when we study the longer rolling-estimation windows, in the case of downside *RV* mainly for the longer forecast horizons. As a consequence, it is not surprising that using both El Niño and La Niña phases as predictors of upside/downside *RV* results in significant test results for several combinations of the rolling estimation window and the forecast horizon.

Table 6. Test Results for Upside and Downside Variances.

Panel A: Futures Upside RV (El Niño)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0802	0.0083	0.0436	0.0651	0.0199	0.0351	0.0661
Window length = 60	0.2215	0.0261	0.0388	0.0385	0.0540	0.0393	0.0463
Window length = 72	0.2874	0.0776	0.0740	0.0650	0.0872	0.3049	0.5776
Window length = 96	0.6678	0.0958	0.2408	0.3105	0.3256	0.8476	0.7553
Window length = 120	0.4435	0.0650	0.1044	0.1062	0.0232	0.6889	0.6491

Panel B: Futures Upside RV (La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1360	0.0536	0.1111	0.1723	0.0265	0.0649	0.0304
Window length = 60	0.0655	0.0114	0.0059	0.0055	0.0026	0.0142	0.0020
Window length = 72	0.2895	0.0650	0.0326	0.0231	0.0081	0.0339	0.0170
Window length = 96	0.1860	0.0451	0.0337	0.0583	0.1715	0.1808	0.0711
Window length = 120	0.1357	0.1055	0.0864	0.1136	0.1010	0.1301	0.0827

Panel C: Futures Upside RV (El Niño and La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.0533	0.0032	0.0102	0.0393	0.0106	0.0134	0.0128
Window length = 60	0.1807	0.0060	0.0114	0.0171	0.0201	0.0063	0.0016
Window length = 72	0.2214	0.0263	0.0364	0.0186	0.0066	0.0472	0.0292
Window length = 96	0.2847	0.0135	0.0326	0.0229	0.1442	0.1936	0.0782
Window length = 120	0.3162	0.0477	0.0733	0.0343	0.0821	0.1282	0.0658

Panel D: Futures Downside RV (El Niño)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.2710	0.0041	0.0225	0.0106	0.0086	0.0434	0.0574
Window length = 60	0.0560	0.0106	0.0298	0.0224	0.0411	0.0454	0.0393
Window length = 72	0.2463	0.1200	0.2237	0.2001	0.3406	0.6335	0.7758
Window length = 96	0.2147	0.1212	0.2447	0.3576	0.5843	0.9916	0.9729
Window length = 120	0.1684	0.0538	0.0585	0.0428	0.1032	0.0623	0.3510

Panel E: Futures Downside RV (La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.2406	0.1667	0.2421	0.4474	0.0424	0.0513	0.0483
Window length = 60	0.0164	0.0844	0.0665	0.0249	0.0008	0.0055	0.0027
Window length = 72	0.2914	0.1460	0.1349	0.0762	0.0065	0.0006	0.0004
Window length = 96	0.0882	0.1237	0.1200	0.0665	0.1919	0.3756	0.1083
Window length = 120	0.1228	0.1277	0.1302	0.0547	0.0141	0.0258	0.0377

Panel F: Futures Downside RV (El Niño and La Niña)						
Forecast horizon	h = 1	h = 3	h = 6	h = 12	h = 24	h = 36	h = 48
Window length = 48	0.1381	0.0011	0.0101	0.0317	0.0019	0.0203	0.0201
Window length = 60	0.0170	0.0359	0.0430	0.0298	0.0035	0.0041	0.0013
Window length = 72	0.3039	0.1486	0.1882	0.0822	0.0046	0.0001	0.0002
Window length = 96	0.1072	0.0984	0.1174	0.0554	0.2025	0.4249	0.1153
Window length = 120	0.1366	0.0746	0.0512	0.0230	0.0147	0.0111	0.0199

Note: The results for upside and downside variances are based on the futures and the EQSOI data. This table reports results (*p*-values) of the Clark–West test for an equal mean squared prediction error (MSPE) for alternative forecast horizons and alternative lengths of the rolling-estimation window used to compute forecasts. The baseline HAR–RV model is the benchmark model, and the HAR–RV that features El Niño and/or La Niña phases as predictors is the rival model. The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The parameter *h* denotes the forecast horizon (in months). The *p*-values are based on robust standard errors.

5. Conclusions

Our empirical findings show that both El Niño and La Niña phases contain predictive value for the realized variance (and volatility) of heating oil price movements. Across the various model configurations that we have studied, however, the predictive value of La Niña phases appears to be stronger than the predictive value of El Niño phases in terms of both statistical significance of the test results and in terms of forecast accuracy, which is not surprising given that the latter is associated with the cooling of temperatures and, hence, greater demand for heating oil and a stronger impact on volatility via trading in the market.

Our results are useful for investors who need to obtain reliable information on the response of the future path of heating oil price fluctuations to changes in climate patterns for portfolio management and hedging decisions. Our results also have implications from the context of "sustainability"; in this case related to the environment. Note that La Niña phases lead to a greater demand for heating oil, and hence drive the volatility of heating oil price. Now, the rise in heating oil demand is expected to cause an increase in the emissions of carbon dioxide, which is known to result in global warming [49–52]. The implication of this is that the warming of the ocean's surface is faster than the subsurface, which is likely to induce occurrences of the El Niño phenomena [53]. In this regard, the National Oceanic and Atmospheric Management Climate Prediction Center states that global warming makes the extreme weather more severe, including an increase in the frequency of the super El Niño phenomenon. Therefore, there is a close feedback relationship between the El Niño phenomenon and the movements in prices in the heating oil market, resulting from its high demand due to La Niña phases. Naturally, policy authorities need to ensure stability in the price of the heating oil market, which would only be possible if the demand is kept from fluctuating wildly in the wake of La Niña events, and this would entail investing more in greener, more sustainable and renewable energy sources such as solar, wind, hydro, ocean, geothermal, biomass and hydrogen.

In future research, it will be interesting to build on our empirical research and examine how the returns and the realized variance of the prices of non-energy (for example, agricultural) commodities respond to El Niño and La Niña phases, and whether these phases help to forecast price movements. Such a forecasting exercise for the returns and the variance should be particularly useful for policymakers in developing countries, and in countries whose exports consist, to a large extent, of agricultural commodities.

Appendix A

Table A1. *k*-th Order Nonparametric Causality-in-quantiles Test Results.

Quantile	Heating Oil Returns				Heating Oil Volatility (Squared Returns)			
	Spot		Futures		Spot		Futures	
	El Niño	La Niña	El Niño	La Niña	El Niño	La Niña	El Niño	La Niña
0.05	1.89 *	1.98 **	2.00 **	2.29 **	1.70 *	1.84 *	2.21 **	2.24 **
0.10	2.71 ***	2.61 ***	3.40 ***	3.23 ***	2.63 ***	2.58 ***	3.26 ***	3.07 ***
0.15	3.82 ***	3.51 ***	4.32 ***	4.33 ***	3.57 ***	3.40 ***	4.07 ***	4.08 ***
0.20	4.44 ***	4.35 ***	4.78 ***	4.86 ***	4.08 ***	3.91 ***	4.81 ***	5.10 ***
0.25	4.66 ***	4.41 ***	5.77 ***	5.39 ***	4.67 ***	4.51 ***	5.36 ***	5.11 ***
0.30	4.93 ***	5.08 ***	5.80 ***	5.50 ***	4.79 ***	4.84 ***	5.75 ***	5.95 ***
0.35	5.19 ***	5.21 ***	5.86 ***	5.90 ***	5.08 ***	5.39 ***	6.10 ***	6.19 ***
0.40	5.22 ***	5.22 ***	5.86 ***	6.03 ***	5.26 ***	5.34 ***	5.95 ***	6.01 ***
0.45	5.26 ***	5.28 ***	5.91 ***	5.84 ***	5.38 ***	5.39 ***	5.84 ***	6.03 ***
0.50	5.53 ***	5.40 ***	5.98 ***	5.90 ***	5.38 ***	5.31 ***	5.93 ***	5.81 ***
0.55	5.09 ***	5.25 ***	5.91 ***	5.62 ***	5.09 ***	5.09 ***	5.87 ***	5.73 ***
0.60	4.86 ***	5.03 ***	5.59 ***	5.43 ***	4.91 ***	4.88 ***	5.50 ***	5.34 ***
0.65	4.71 ***	4.87 ***	5.16 ***	5.31 ***	4.66 ***	4.86 ***	5.48 ***	5.30 ***
0.70	4.15 ***	4.45 ***	4.86 ***	5.22 ***	4.26 ***	4.64 ***	5.12 ***	5.13 ***
0.75	3.96 ***	4.17 ***	4.57 ***	4.90 ***	4.00 ***	4.17 ***	4.66 ***	5.10 ***
0.80	3.59 ***	3.72 ***	4.28 ***	4.45 ***	3.48 ***	3.76 ***	4.14 ***	4.23 ***
0.85	3.09 ***	3.42 ***	3.55 ***	3.57 ***	3.02 ***	3.32 ***	3.61 ***	3.66 ***
0.90	2.79 ***	2.87 ***	2.85 ***	3.09 ***	2.69 ***	2.87 ***	3.13 ***	3.21 ***
0.95	2.22 **	1.98 **	2.24 **	2.25 **	2.06 **	2.03 **	2.10 **	2.12 **

Note: *, **, and *** indicates rejection of the null hypothesis of no-causality due to El Niño or La Niña phases derived from the EQSOI (when controlling for La Niña or El Niño in the model) for a specific quantile at 10%, 5%, and 1% level of significance, with critical values of 1.645, 1.96 and 2.575 respectively corresponding to the standard normal test-statistic of the k-th order nonparametric causality-in-quantiles test of Balcilar et al. (2018) [46].
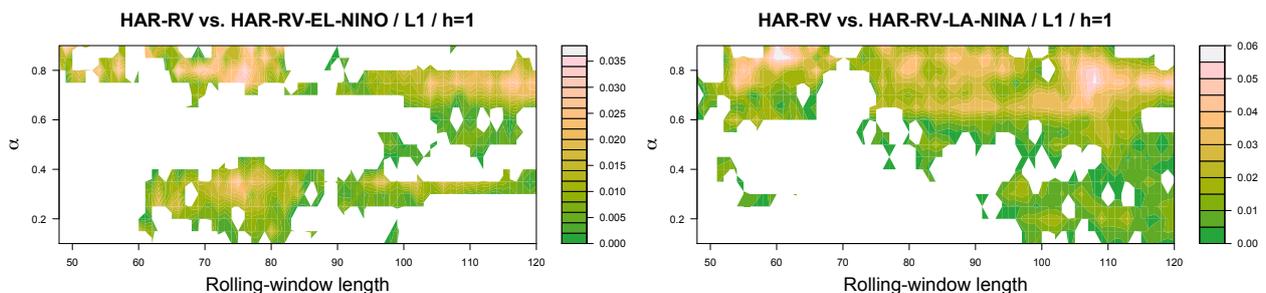
Figure A1. *Cont.*
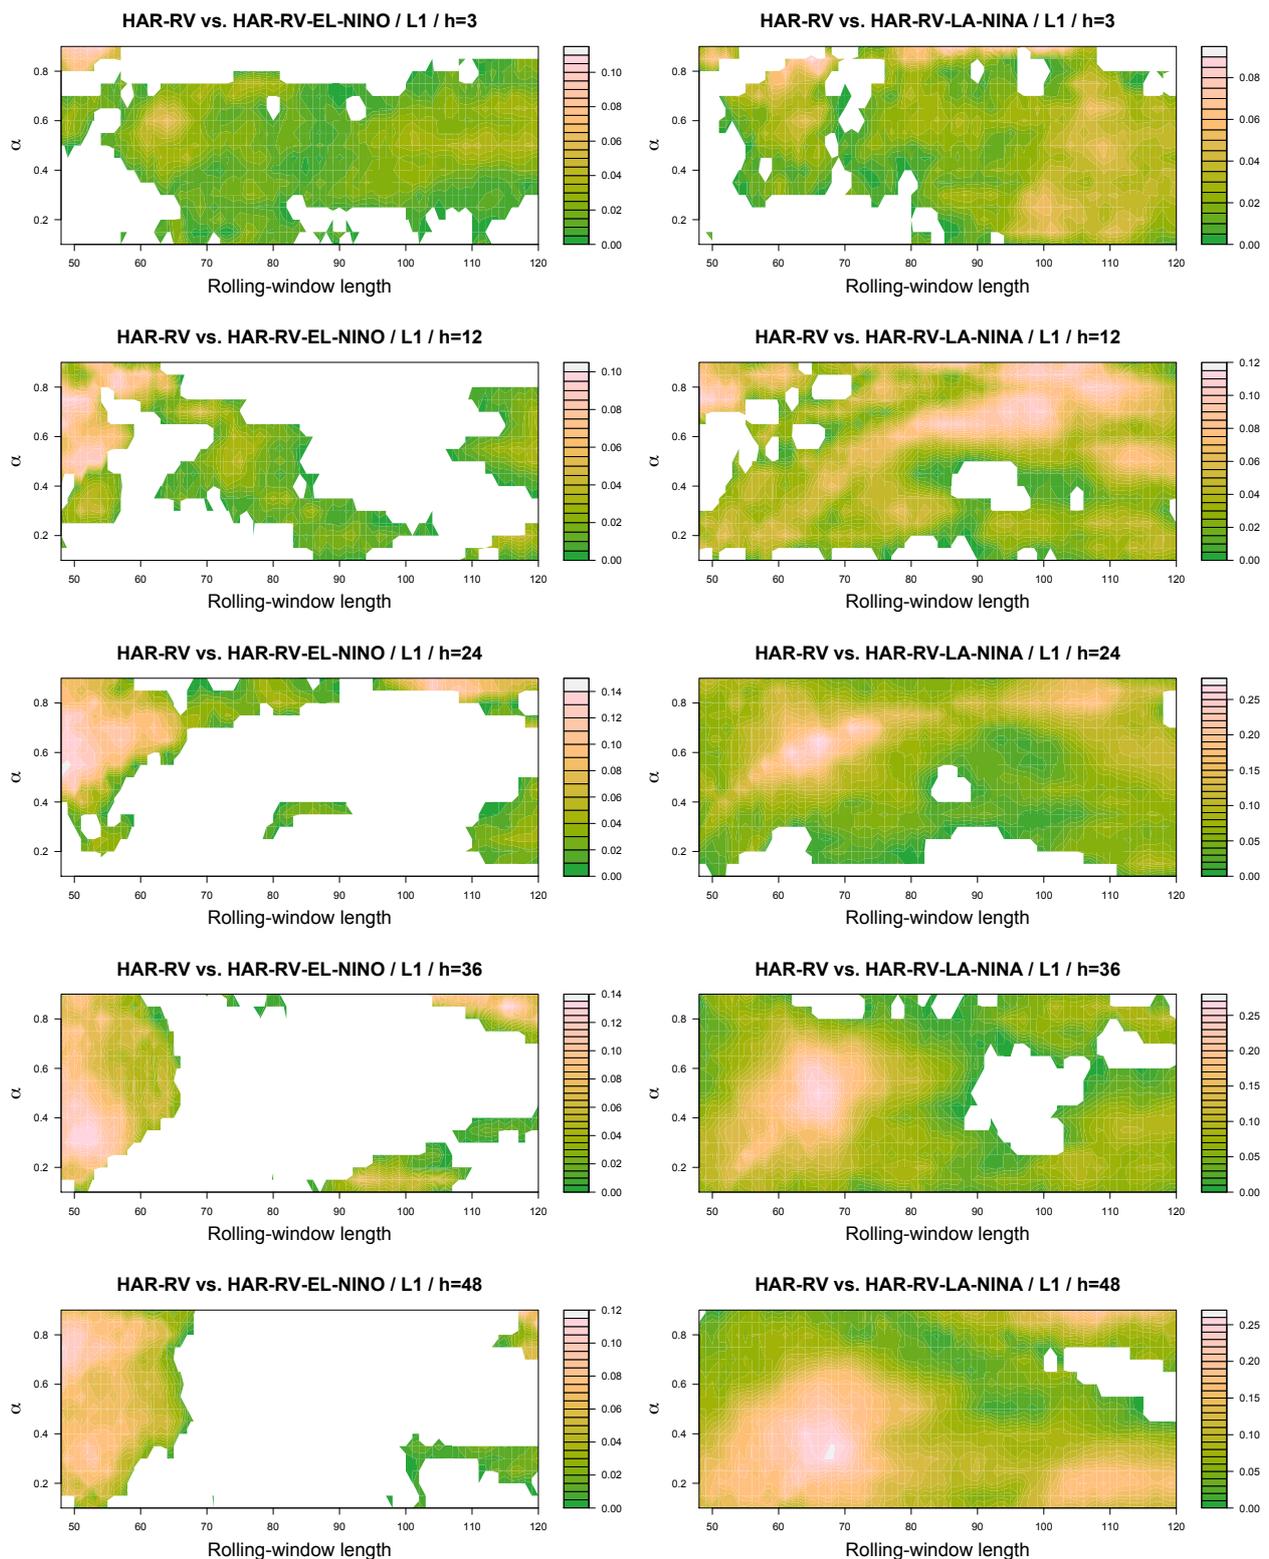
Figure A1. Results for the HARQ-RV Model (Futures; Lin-Lin Loss Function). Note: This figure reports the out-of-sample relative loss criterion for a lin-lin loss function (L1) and an extended HARQ-RV models that features either El Niño or La Niña phases as an additional predictor. The baseline model is the benchmark HARQ-RV model. White spaces indicate areas where the out-of-sample relative loss criterion takes on negative values (the baseline model performs better than the extended model). The rolling-estimation window is varied from 48, 49, 50, ... to 120 observations. The results are based on the SOI data. The parameter *h* denotes the forecast horizon (in months).
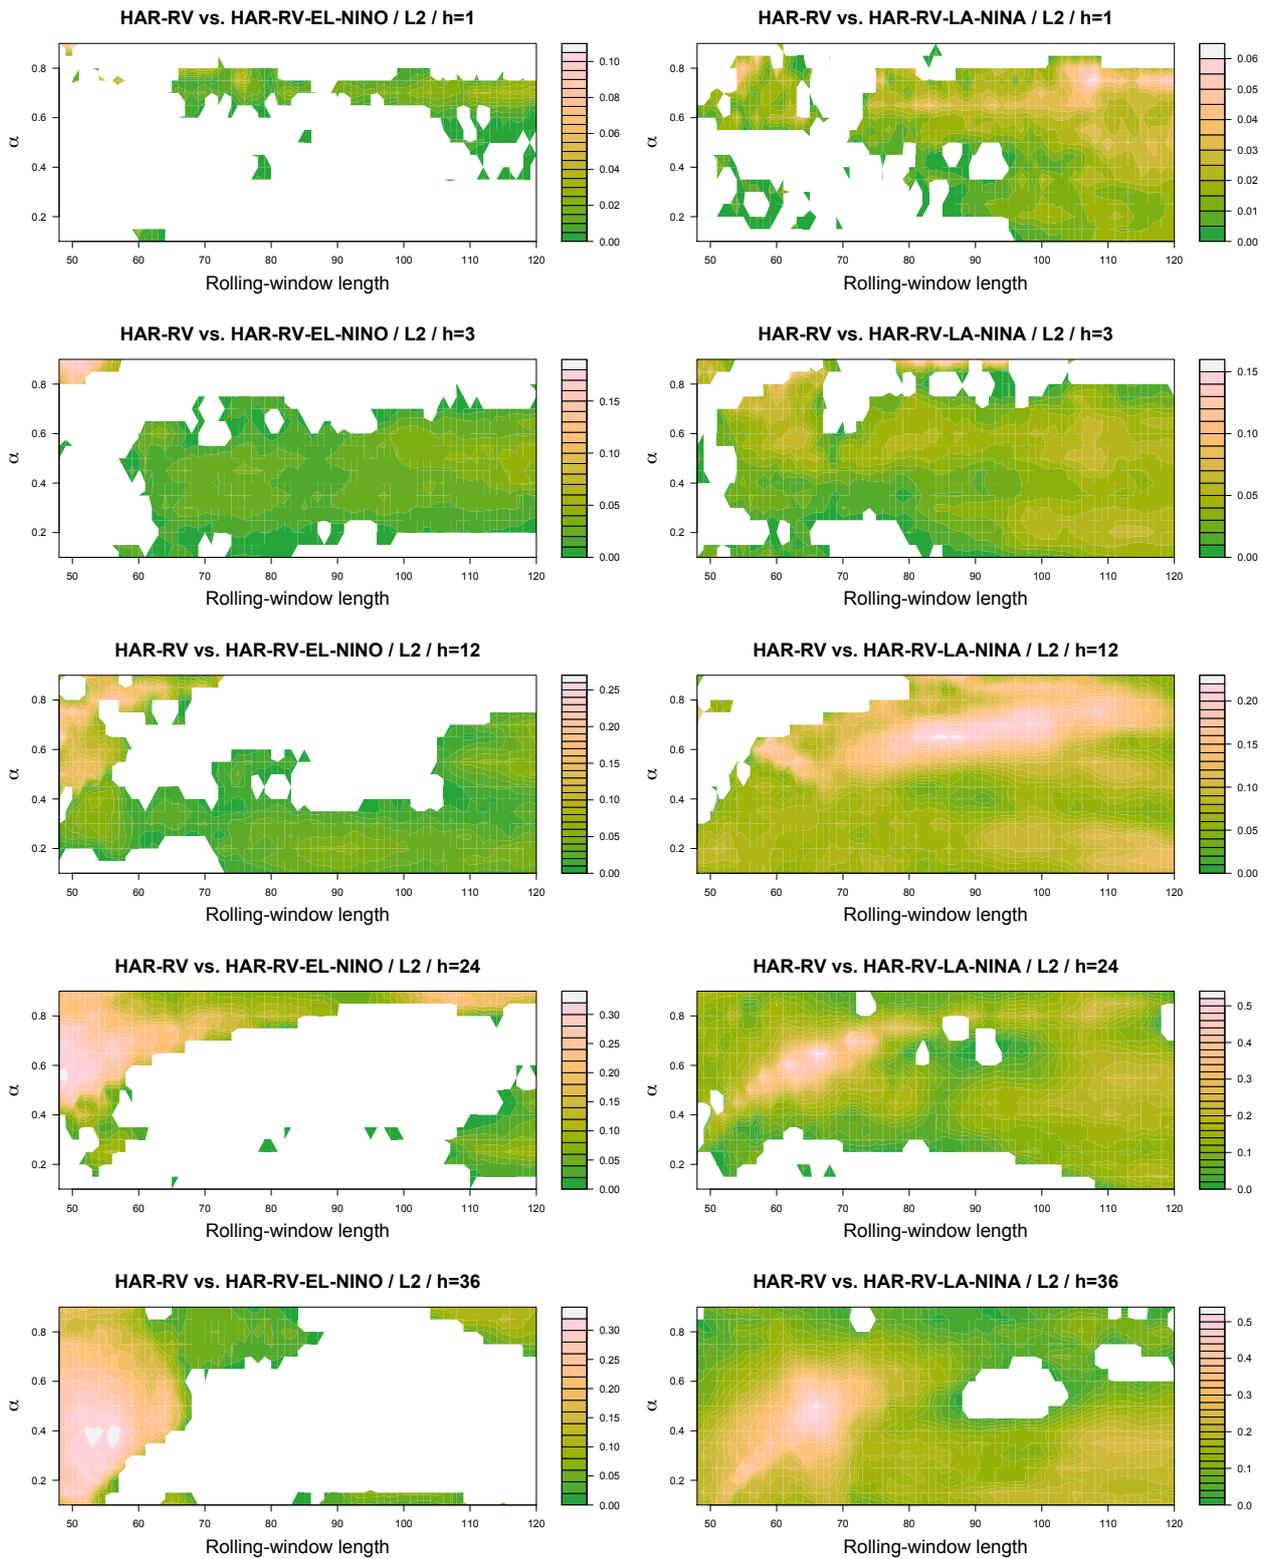
Figure A2. *Cont.*
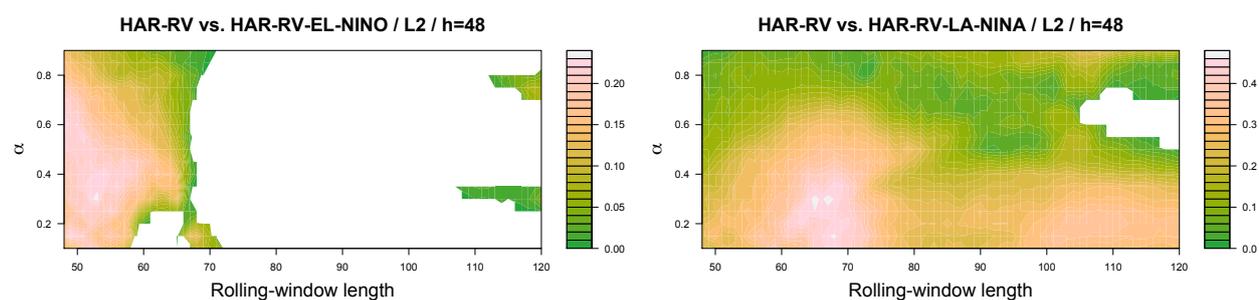
Figure A2. Results for the HARQ-RV Model (Futures; Quad-Quad Loss Function). Note: This figure reports the out-of-sample relative loss criterion for a quad-quad loss function (L2) and an extended HARQ-RV models that features either El Niño or La Niña phases as an additional predictor. The baseline model is the benchmark HARQ-RV model. White spaces indicate areas where the out-of-sample relative loss criterion takes on negative values (the baseline model performs better than the extended model). The rolling-estimation window is varied from 48, 49, 50, ... to 120 observations. The results are based on the SOI data. The parameter *h* denotes the forecast horizon (in months).

References

1. Trenberth, K.E.; Jones, P.D.; Ambenje, P.; Bojariu, R.; Easterling, D.; Tank, K.A.; Parker, D.; Rahimzadeh, F.; Renwick, J.A.; Rusticucci, M.; et al. (Eds.) *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*; Cambridge University Press: Cambridge, UK, 2007; pp. 235–336.
2. Qin, M.; Qiu, L-H.; Tao, R.; Umar, M.; Su, C-W.; Jiao, W. The inevitable role of El Niño: A fresh insight into the oil market. *Econ.-Res.-Ekon. IstražIvanja* **2020**, *33*, 1943–1962. [CrossRef]
3. Grove, R.H.; Chappell, J. *El Niño?History and Crisis*; The White Horse Press: Cambridge, UK, 2000.
4. Davis, M. *Late Victorian Holocausts: El Niño Famines and the Making of the Third World*; Verso: London, UK, 2001.
5. Changnon, S.A. Impacts of 1997–98 El Niño generated weather in the United States. *Bull. Am. Meteorol. Soc.* **1999**, *80*, 1819–1827. [CrossRef]
6. Brunner, A.D. El Niño and World Primary Commodity Prices: Warm Water or Hot Air? *Rev. Econ. Stat.* **2002**, *84*, 176–183. [CrossRef]
7. Peersman, G. International food commodity prices and missing (dis)inflation in the euro area. *Rev. Econ. Stat.* **2018**, 1–43. [CrossRef]
8. De Winne, J.; Peersman, G. *Economic Consequences of Global Agricultural Price Shifts Caused by Harvest and Weather Disruptions*; Working Paper No. 21/1012; Ghent University, Faculty of Economics and Business Administration: Ghent, Belgium, 2021.
9. Paldam, M. Inflation and political instability in eight Latin American countries 1946-83. *Public Choice* **1987**, *52*, 143–168. [CrossRef]
10. Bittencourt, M.; Gupta, R.; Makena, P.; Stander, L. Socio-Political Instability and Growth Dynamics. *Econ. Syst.* Forthcoming.
11. Tol, R.S.J. The economic effects of climate change. *J. Econ. Perspect.* **2009**, *23*, 29–51. [CrossRef]
12. Hsiang, S.M.; Meng, K.C.; Cane, M.A. Civil conflicts are associated with the global climate. *Nature* **2011**, *476*, 438–441. [CrossRef]
13. Dell, M.; Jones, B.F.; Olken, B.A. What do we learn from the weather? The new climate-economy literature. *J. Econ. Lit.* **2014**, *52*, 740–798. [CrossRef]
14. De Winne, J.; Peersman, G. The impact of food prices on conflict revisited. *J. Bus. Econ. Stat.* **2021**, *39*, 547–560. [CrossRef]
15. Bouri, E.; Gupta, R.; Pierdzioch, C.; Salisu, A.A. El Niño and forecastability of oil-price realized volatility. *Theor. Appl. Climatol.* **2021**, *144*, 1173–1180. [CrossRef]
16. Demirer, R.; Gupta, R.; Suleman, M.T.; Wohar, M.E. Time-varying rare disaster risks, oil returns and volatility. *Energy Econ.* **2018**, *75*, 239–248. [CrossRef]
17. Gkillas, K.; Floros, C.; Suleman, M.T. Quantile dependencies between discontinuities and time-varying rare disaster risks. *Eur. J. Financ.* **2020**, 1–31. 1080/1351847X.2020.1809487. [CrossRef]
18. Caporin, M.; Rossi, E.; Santucci de Magistris, P. Volatility jumps and their economic determinants. *J. Financ. Econom.* **2016**, *14*, 29–80. [CrossRef]
19. Gupta, R.; Suleman, M.T.; Wohar, M.E. Exchange rate returns and volatility: the role of time-varying rare disaster risks. *Eur. J. Financ.* **2019**, *25*, 190–203. [CrossRef]
20. Gupta, R.; Suleman, M.T.; Wohar, M.E. The role of time-varying rare disaster risks in predicting bond returns and volatility. *Rev. Financ. Econ.* **2019**, *37*, 327–340. [CrossRef]
21. Campbell, J.Y. Viewpoint: Estimating the equity premium. *Can. J. Econ.* **2008**, *41*, 1–21. [CrossRef]
22. Bampinas, G.; Panagiotidis, T. On the Relationship between Oil and Gold before and after Financial Crisis: Linear, Nonlinear and Time-Varying Causality Testing. *Stud. Nonlinear Dyn. Econom.* **2015**, *19*, 657–668. [CrossRef]
23. Bampinas, G.; Panagiotidis, T. Oil and stock markets before and after financial crises: A local Gaussian correlation approach. *J. Futur. Mark.* **2017**, *37*, 1179–1204 [CrossRef]
24. Corsi, F. A simple approximate long-memory model of realized volatility. *J. Financ. Econom.* **2009**, *7*, 174–196. [CrossRef]

25. Andersen T.G.; Bollerslev T. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *Int. Econ. Rev.* **1998**, *39*, 885–905. [CrossRef]
26. Sadorsky P. Modeling and forecasting petroleum futures volatility. *Energy Econ.* **2006**, *28*, 467–488. [CrossRef]
27. Nomikos N.K.; Pouliasis P.K. Forecasting petroleum futures markets volatility: The role of regimes and market conditions. *Energy Econ.* **2011**, *33*, 321–337. [CrossRef]
28. Arouri, M.; Lahiani, A.; Levy, A.; Nguyen, D. Forecasting the conditional volatility of oil spot and futures prices with structural breaks and long memory models. *Energy Econ.* **2012**, *34*, 283–293. [CrossRef]
29. Kang S.H.; Yoon S. Modeling and forecasting the volatility of petroleum futures prices'. *Energy Econ.* **2012**, *36*, 354–362. [CrossRef]
30. Wang Y.; Wu C. Forecasting energy market volatility using GARCH models: Can multivariate models beat univariate models? *Energy Econ.* **2012**, *34*, 2167–2181. [CrossRef]
31. Kumar, D. Forecasting energy futures volatility based on the unbiased extreme value volatility estimator. *IIMB Manag. Rev.* **2017**, *29*, 294–310. [CrossRef]
32. Fałdziński, M.; Fiszeder, P.; Orzeszko, W. Forecasting Volatility of Energy Commodities: Comparison of GARCH Models with Support Vector Regression. *Energies* **2021**, *14*, 6. [CrossRef]
33. Asai, M.; Gupta, R.; McAleer, M. The impact of jumps and leverage in forecasting the co-volatility of oil and gold futures. *Energies* **2019**, *12*, 3379. [CrossRef]
34. Asai M.; Gupta, R.; McAleer, M. Forecasting volatility and co-volatility of crude oil and gold futures: Effects of leverage, jumps, spillovers, and geopolitical risks. *Int. J. Forecast.* **2020**, *36*, 933–948. [CrossRef]
35. Müller, U.A.; Dacorogna, M.M.; Davé, R.D.; Olsen, R.B.; Pictet, O.V. Volatilities of different time resolutions—Analyzing the dynamics of market components. *J. Empir. Financ.* **1997**, *4*, 213–239. [CrossRef]
36. Halkos, G.E.; Tsirivis, A.S. Effective energy commodity risk management: Econometric modeling of price volatility. *Econ. Anal. Policy* **2019**, *63*, 234–250. [CrossRef]
37. Baba, Y.; Engle, R.F.; Kraft, D.F.; Kroner, K.F. *Mulitvariate Simultaneous Generalized ARCH*; Mimeo: New York, NY, USA; Department of Economics, University of California: San Diego, CA, USA, 1990.
38. Cashin, P.; Mohaddes, K.; Raissi, M. Fair weather or foul? The macroeconomic effects of El Niño. *J. Int. Econ.* **2017**, *106*, 37–54. [CrossRef]
39. Clark, T.D.; West, K.D. Approximately normal tests for equal predictive accuracy in nested models. *J. Econom.* **2007**, *138*, 291–311. [CrossRef]
40. *R Core Team R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2019. Available online: https://www.R-project.org/ (accessed on 21 May 202).
41. Haugom, E.; Ray, R.; Ulrich, C.J.; Veka, S. A parsimonious quantile regression model to forecast day-ahead value-at-risk. *Financ. Res. Lett.* **2016**, *16*, 196–207. [CrossRef]
42. Gkillas 2020, K.; Gupta, R.; Pierdzioch, C. Forecasting realized gold volatility: Is there a role of geopolitical risks? *Financ. Res. Lett.* **2020**, *35*, 101280. [CrossRef]
43. Pierdzioch C.; Risse M.; Rohloff, S. The international business cycle and gold-price fluctuations. *Q. Rev. Econ. Financ.* **2014**, *54*, 292–305. [CrossRef]
44. Pierdzioch C.; Risse M.; Rohloff, S. Fluctuations of the real exchange rate, real interest rates, and the dynamics of the price of gold in a small open economy. *Empir. Econ.* **2016**, *51*, 1481–1499. [CrossRef]
45. Elliott, G.; Komunjer I.; Timmermann, A. Estimation and testing of forecast rationality Uunder flexible loss. *Rev. Econ. Stud.* **2005**, *72*, 1107–1125. [CrossRef]
46. Balcilar, M.; Gupta R.; Nguyen D.K.; Wohar, M.E. Causal effects of the United States and Japan on Pacific-Rim stock markets: nonparametric quantile causality approach. *Appl. Econ.* **2018**, *50*, 5712–5727. [CrossRef]
47. Belloni, A.; Chernozhukov, V.; Chetverikov, D.; Fernandez-Val, I. Conditional quantile processes based on series or many regressors. *J. Econom.* **2019**, *213*, 4–29. [CrossRef]
48. Giot, P.; Laurent, S.; Petitjean, M. Trading activity, realized volatility and jumps. *J. Empir. Financ.* **2010**, *17*, 168–175. [CrossRef]
49. Ramanathan, R. A multi-factor efficiency perspective to the relationships among world GDP, energy consumption and carbon dioxide emissions. *Technol. Forecast. Soc. Chang.* **2006**, *73*, 483–494. [CrossRef]
50. Alkhathlan, K.; Javid, M. Carbon emissions and oil consumption in Saudi Arabia. *Renew. Sustain. Energy Rev.* **2015**, *48*, 105–111. [CrossRef]
51. Agbanike, T.F.; Nwani, C.; Uwazie, U.I.; Anochiwa, L.I.; Onoja, T.C.; Ogbonnaya, I.O. Oil price, energy consumption and carbon dioxide (CO_2) emissions: Insight into sustainability challenges in Venezuela. *Lat. Am. Econ. Rev.* **2019**, *28*, 1–26. [CrossRef]
52. Kashif, M.; Nimra, R. Energy consumption and environmental quality in South Asia: Evidence from panel non-linear ARDL. *Environ. Sci. Pollut. Res. Int.* **2019**, *26*, 29307–29315.
53. Cai, W.J.; Borlace, S.; Lengaigne, M.; Rensch, P.V.; Collins, M.; Vecchi, G.; Timmermann, A.; Santoso, A.; McPhaden, M.J.; Wu, L.X.; et al. Increasing frequency of extreme El Niño events due to greenhouse warming. *Nat. Clim. Chang.* **2014**, *4*, 111–116. [CrossRef]