Received: 26 April 2022 | Revised: 8 August 2022 | Accepted: 14 August 2022

DOI: 10.1002/for.2914

RESEARCH ARTICLE

WILEY

El Niño, La Niña, and forecastability of the realized variance of agricultural commodity prices: Evidence from a machine learning approach

Matteo Bonato[1,2] | **Oğuzhan Çepni**[3,4] | **Rangan Gupta**[5] | **Christian Pierdzioch**[6]

[1]Department of Economics and Econometrics, University of Johannesburg, Auckland Park, Johannesburg, South Africa

[2]IPAG Business School, 184 Boulevard Saint-Germain, Paris, 75006, France

[3]Copenhagen Business School, Department of Economics, Porcelænshaven 16A, Frederiksberg DK-2000, Denmark

[4]Central Bank of the Republic of Turkey, Haci Bayram Mah. Istiklal Cad. No:10 06050, Ankara, Turkey

[5]Department of Economics, University of Pretoria, Private Bag X20, Hatfield, 0028

[6]Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, Hamburg, 22008, Germany

Correspondence

Rangan Gupta, Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa.
Email: rangan.gupta@up.ac.za

Christian Pierdzioch, Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, Hamburg, 22008, Germany.
Email: macroeconomics@hsu-hh.de

Abstract

We examine the predictive value of El Niño and La Niña weather episodes for the subsequent realized variance of 16 agricultural commodity prices. To this end, we use high-frequency data covering the period from 2009 to 2020 to estimate the realized variance along realized skewness, realized kurtosis, realized jumps, and realized upside and downside tail risks as control variables. Accounting for the impact of the control variables as well as spillover effects from the realized variances of the other agricultural commodities in our sample, we estimate an extended heterogeneous autoregressive (HAR) model by means of random forests to capture in a purely data-driven way potentially nonlinear links between El Niño and La Niña and the subsequent realized variance. We document such nonlinear links, and that El Niño and La Niña increase forecast accuracy, especially at longer forecast horizons, for several of the agricultural commodities that we study in this research.

KEYWORDS

agricultural commodities, El Niño and La Niña, forecasting, random forests, realized variance

1 | INTRODUCTION

It is well established that the so-called El Niño-Southern Oscillation (ENSO), an irregularly periodic variation in winds and sea surface temperatures over the tropical eastern Pacific Ocean, tends to influence the climate of much of the tropics and subtropics (Trenberth et al., 2007). The warming phase of sea temperature is known as El Niño and the corresponding cooling phase as La Niña. Each of these two phases can last several months, and usually, they occur every few years with intensities varying per phase. Understandably, the ENSO is an important source of inter-annual variability in weather and climate patterns in many parts of the world (Shabbar & Khandekar, 1996). Not surprisingly, quite a few recent studies, following the early work of Brunner (2002), have highlighted the significant impact of the ENSO on prices of agricultural commodities (see, for example, Bastianin et al., 2018; Ubilava, 2012a, 2012b, 2014, 2017, 2018; Ubilava & Holt, 2013).[1]

All the above mentioned studies, however, have analyzed the impact of the ENSO on the first moment of agricultural commodity prices and/or returns and have primarily been in-sample-based structural analyses. An exception is the study by Ubilava (2018), who reports the results of an out-of-sample forecasting experiment based on nonlinear (smooth-transition autoregressive) models. Against this backdrop, the objective of our research is to shed light on the ability of the ENSO to forecast the realized variance of the returns of agricultural commodity prices, where we separate out the El Niño and La Niña phases. Economically, forecasting realized variance is of paramount importance because agricultural commodities have experienced enormous price swings since 2008, resulting in both high and low volatility regimes (Greb & Prakash, 2015). In this regard, Johnson (2011) argues that increased volatility primarily stems from extreme weather events (albeit factors such as the production of biofuels, market speculation, rising demand, and declines in food stocks cannot be ruled out). Naturally, accurate forecasts of the volatility of movements of the prices of agricultural commodities are of key importance for policy authorities, who need to ensure food security, as well as for market participants and traders, given that volatility is a key input for investment decisions. In terms of forecasting, the usefulness of in-sample results is limited because, from a statistical perspective, in-sample predictability does not necessarily translate into a good out-of-sample forecasting performance of a specific predictor, besides the fact that it is out-of-sample forecasting that tends to constitute a more stringent test of the appropriateness of an econometric model being studied and a predictor being under scrutiny (Campbell, 2008). Hence, we conduct an out-of-sample forecasting experiment to identify the role of the ENSO for the realized variance of returns of the prices of 16 highly traded agricultural commodities.[2]

At this stage, it is important to describe the underlying economic channel via which we conjecture the link between the ENSO and the realized variance of the returns of the agricultural commodities prices. Recent studies (Balcilar et al., 2021; Bouri et al., 2021; Demirer et al. 2022; Qin et al., 2020) have indicated that the ENSO can successfully serve as an empirical proxy for the theoretical concept of rare disaster risks, as coined by Rietz (1988), Barro, (2006, 2009), and Gabaix (2012). The rise of rare disaster risks associated with the ENSO is likely to make the path of future aggregate demand and aggregate production less predictable. Facing the enhanced uncertainty emanating from this more intense unpredictability, risk-averse commodity producers will prefer to hold physical inventory when facing uncertain aggregate demand conditions. Increases in inventories, in turn, will increase the convenience yield for holding physical inventory and eventually will amplify the variance of returns of agricultural commodity prices. See Bakas and Triantafyllou (2018, 2020) for an analysis of the implications of uncertainty shocks and uncertainty due to pandemics for the volatility of commodity prices.

As far as the econometric model is concerned, we forecast weekly realized variance (RV) using an extended version of the heterogeneous autoregressive (HAR)-RV model of Corsi (2009). Our extended HAR-RV model incorporates the role of not only the El Niño and La Niña events associated with the ENSO but also a rich array of control variables, including cross-market RVs, realized jumps, realized tail risks, realized skewness, and realized kurtosis, over the sample period from 9/30/2009 to 5/13/2020. We measure RV by computing the sum of squared 5-min intraday returns of the various agricultural commodity prices over a week (following Andersen & Bollerslev, 1998). In this regard, it is essential to note that using RV to measure the variance of returns of agricultural commodity prices has the advantage that we can build our empirical analysis on an observable and unconditional metric of "volatility" (unlike in the case of generalized autoregressive conditional heteroscedastic [GARCH] and stochastic volatility [SV]) models that researchers have traditionally used to model and forecast agricultural commodity price volatility),[3] which is a latent process. Besides this, as pointed out by McAleer and Medeiros (2008), because intraday data contain rich information, studying RV leads to more accurate estimates and forecasts of (weekly) realized variance. At the same time, our decision to study the HAR-RV model extended to include various additional control variables is based on much significant recent research on forecasting RV of returns of agricultural commodity prices (see, for example, Degiannakis et al., 2022; Luo et al., 2022; Marfatia et al., 2022; Tian et al., 2017a, 2017b; Yang et al., 2017).[4] The popularity of the HAR-RV model stems from the fact that, while its basic structure is rather simple, it can capture long-memory and multi-scaling properties of the returns of agricultural commodity RV, which have been reported by Gil-Alana et al. (2012) and Živkov et al. (2019). In addition, the HAR-RV model, which uses RV from different time resolutions to forecast the RV of the returns of agricultural commodity prices, has a solid theoretical foundation in the form of the so-called heterogeneous market hypothesis (Müller et al., 1997) according to which different groups of participants who differ in their sensitivity to information flows at different time horizons populate the markets for agricultural commodities.

Econometrically, for our forecasting application, we use a machine learning technique known as random forests (Breiman, 2001) to compute forecasts of overall RV, as well as "good" (weekly sum of squared positive intraday returns) and "bad" (weekly sum of squared negative intraday returns) RV, given the observation made by Giot et al. (2010) that market agents care not only about the level of volatility but also of its nature, with traders typically making the distinction between upside and downside volatilities. The use of random forests in our forecasting application has several advantages. First, in a completely data-driven way, random forests are tailored to capture the links between RV and an arbitrarily large number of predictors. In our forecasting experiment, we study as many as 26 predictors, given that combination of information from predictors tends to matter in this literature (Tian et al., 2017b). Second, random forests automatically capture not only linear but also potential nonlinear links between RV and El Niño and La Niña phases (which is likely to be a pertinent feature of the data given the nonlinear evolution of the ENSO in itself Hall et al., 2001) besides the other predictors (Luo et al., 2022; Tian et al., 2017a), as well as any interaction effects among the predictors. Finally, unlike the ordinary-least-squares (OLS) technique commonly used to estimate HAR-RV models, random forests always produce forecasts of RV that are non-negative.

In earlier studies on forecasting the RV of the returns of agricultural commodity prices based on variants and extensions of the HAR-RV model, researchers have primarily used derived metrics from intraday data associated with the price of agricultural commodities itself[5]; to the best of our knowledge, ours is the first paper to analyze, in addition, the role of weather- and climate-related risks emanating from the ENSO in forecasting the RV of 16 important agricultural commodities using a machine learning approach. One must realize that accurate forecasts of agricultural commodity prices and their RV is of tremendous importance from the perspective of both investors and policymakers. As pointed out by Girardi (2015), Bruno et al. (2016), Aït Youcef (2019), and (Ouyang & Zhang, 2020), agricultural markets have become increasingly financialized since 2008, resulting in increased holdings of institutional investors. In light of this, proper modeling and forecasting of RV is required as a key input to investment decisions, portfolio allocation, risk management, and evaluation of hedging performance. At the same time, because agricultural commodities represent a major component of household consumption, volatility of their prices is likely to have a pronounced impact on food security, which affects primarily the poorer part of the population (Ordu et al., 2018). Hence, it is necessary to develop models that accurately forecast the RV of the returns of agricultural commodity prices, such that policy institutions can prepare for periods of large price fluctuations, and, in turn, to design preventative policies.

We organize the remainder of our paper as follows. In Section 2, we describe the data and methodologies used in our empirical analysis. In Section 3, we describe our methods. In Section 4, we summarize our empirical results. In Section 5, we conclude.

2 | DATA

2.1 | Agricultural commodities and realized variance

Intraday commodity futures prices are obtained from https://www.kibot.com/. Futures data are in continuous format: close to expiration of a contract the position is rolled over to the next available contract provided that activity has increased. Intraday prices are sampled at a 5-min frequency. Three groups of agricultural commodities are covered in our dataset: grains, softs, and livestock, and, in general, are considered to fall within the category of highly traded agricultural commodities, as identified by the Food and Agriculture Organization (FAO) of the United Nations (UN).[6] Individual commodities along with ticker and trading exchange are reported in Table 1.

We use the classical estimator of RV, that is, the sum of squared intraday returns (Andersen and Bollerslev, 1998), expressed as

$$RV_t^d = \sum_{i=1}^{M} r_{t,i}^2, \tag{1}$$

where $r_{t,i}$ denotes the intraday $M \times 1$ return vector and $i = 1, \ldots, M$ is the number of intraday returns. We calculate the weekly RV by aggregating the daily RV_t^d over a trading week (i.e., from Monday to Friday):

$$RV_t = \sum_{i=Monday}^{Friday} RV_t^i \tag{2}$$

Table 2 reports summary statistics of the weekly RV for the 16 agricultural commodities in our sample. The start and end dates of our sample period are 9/30/2009 and 5/13/2020. After matching all dates across the 16 commodities, which are trading over different exchanges and different features, for example, open out cry versus

TABLE 1 Commodity futures traded in the United States (sectors, tickers, and exchanges)

Commodity	Ticker	Exchange
Grains		
Corn	C	CME Group
Soybeans	S	CME Group
Chicago wheat	W	CME Group
Soybean oil	BO	CME Group
Soybean meal	SM	CME Group
Rough rice	RR	CME Group
Oats	O	CME Group
Softs		
Coffee	KC	ICE
Cotton	CT	ICE
Sugar	SB	ICE
Cocoa	CC	ICE
Lumber	LB	CME Group
Orange juice	OJ	ICE
Livestock		
Feeder cattle	GF	CME Group
Lean hogs	HE	CME Group
Live cattle	LE	CME Group

Abbreviations: CME, Chicago Mercantile Exchange Group; ICE, Intercontinental Exchange.

TABLE 2 Summary statistics of realized variances

Commodity	Ticker	Mean	Median	Max
Soybean oil	BO	0.0008	0.0007	0.0044
Corn	C	0.0012	0.0009	0.0122
Cocoa	CC	0.0012	0.0011	0.0071
Cotton	CT	0.0013	0.0009	0.0134
Feeder cattle	GF	0.0006	0.0005	0.0090
Lean hogs	HE	0.0016	0.0009	0.0405
Coffee	KC	0.0017	0.0014	0.0094
Lumber	LB	0.0021	0.0017	0.0135
Live cattle	LE	0.0005	0.0004	0.0076
Oats	O	0.0028	0.0023	0.0176
Orange juice	OJ	0.0022	0.0018	0.0089
Rough rice	RR	0.0013	0.0010	0.0236
Soybean	S	0.0008	0.0006	0.0137
Sugar	SB	0.0017	0.0015	0.0106
Soybean meals	SM	0.0011	0.0009	0.0127
Chicago wheat	W	0.0015	0.0013	0.0133

globex, we are left with 459 observations.[7] Figure 1 plots the realized variances. A noticeable feature of the realized variances, a feature that is well-known from the volatilities of other financial-market returns data, is that the time series display occasional large peaks.

We also study whether El Niño and La Niña events help to predict downward ("bad," *RVB*) and upward ("good," *RVG*) realized variance (that is, the semi-variances). The bad and good RV are statistics that capture a potential sign asymmetry in the realized-variance process. Like Barndorff-Nielsen et al. (2010), we estimate bad and good realized variance as follows:

$$RVB_t = \sum_{i=1}^{M} r_{t,i}^2 \mathbf{1}_{[(r_{t,i})<0]}, \tag{3}$$

$$RVG_t = \sum_{i=1}^{M} r_{t,i}^2 \mathbf{1}_{[(r_{t,i})>0]}, \tag{4}$$

where **1** denotes the indicator function.

2.2 | ENSO data

As far as the weekly observations on the measure of the ENSO intensity are concerned, we use sea surface temperature anomalies (SSTA) for the "Niño 3.4" region—the region between 5°N–5°S and 120°W–170°W. While other measures of ENSO intensity, such as the Southern Oscillation Index (SOI) or the Equitorial SOI (EQSOI) anomalies, have sometimes been used in the literature, the use of SSTA is more prevalent (Atems & Sardar, 2021; Atems et al., 2020) and is in line with the metric of the ENSO used by Ubilava (2018) in his forecasting exercise. This metric of the ENSO intensity also renders it possible to perform the analysis at a high frequency, that is, weekly, because SOI and EQSOI are only available as monthly data for the period of our empirical exercise. The data on weekly SSTA come from the National Weather Service Climate Prediction Center (National Oceanic and Atmospheric Administration [NOAA]).[8]

An El Niño event is defined as three consecutive months of SSTA of 0.5°C (0.9°F) or higher, for which we define a dummy that take on a value of 1 during such months, and 0 otherwise. A La Niña episode, in turn, is defined as three consecutive months of SSTA of −0.5°C (−0.9°F) or less. Again, we define a dummy that take on a value of 1 in such months, and 0 otherwise.[9] Finally, we multiply the SSTA data with these dummies to create two series that separate out the El Niño and La Niña events, both of which we consider in our forecasting analyses. Figure 2 plots the resulting El Niño and La Niña events.

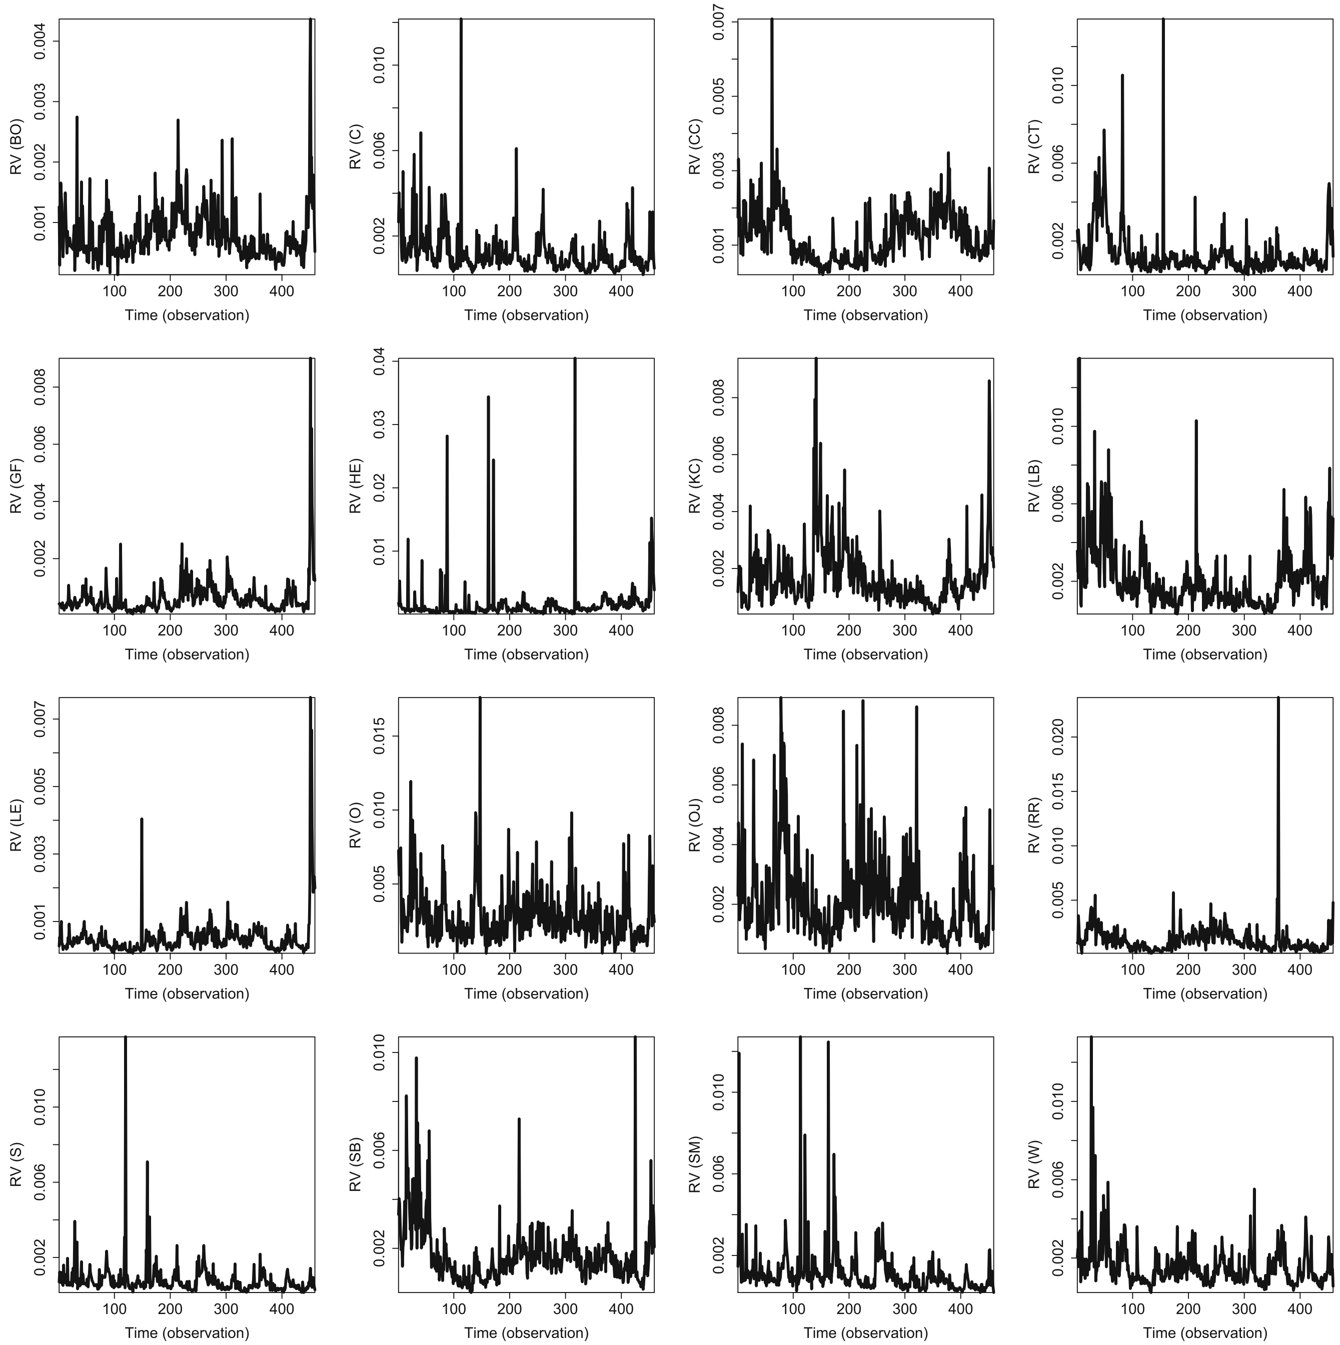

FIGURE 1 Realized variances

3 | METHODS

3.1 | Random forests

In order to study the predictive role of extreme weather events (that is, in our case, El Niño and La Niña events) for the realized variance of the returns of agricultural commodity prices, we use random forests. Random forests have the advantage that they account in a completely data-driven way for potential nonlinear links between realized variance and El Niño and La Niña events. In addition, random forests account for potential interaction effects between predictors, and they provide a natural modeling platform the number of predictors is relatively large. In our case, the number of predictors is relatively large because we account for the potential interconnectedness of agricultural commodity markets in the form of spillover effects of realized variance from other agricultural commodities.

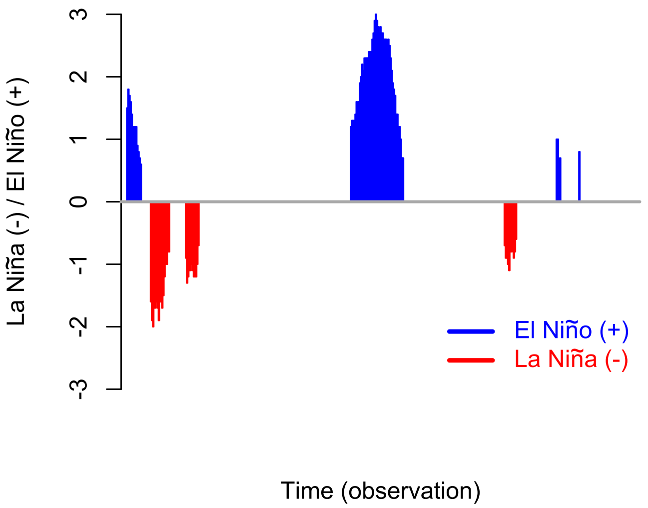

FIGURE 2 El Niño and La Niña events

The nucleus of a random forest is an individual regression tree (for a textbook exposition, see the textbook by Hastie et al. (2009); our notation follows the notation they use in their textbook). The basic idea of a regression tree, T, is to form branches that partition the space of predictors, $x=(x_1,x_2,...)$, into l non-overlapping regions, R_l. At the top level, a regression tree selects the first partition by searching across the set of predictors and corresponding potential partitioning points, p, to form two half-planes, $R_1(s,p)=\{x_s|x_s\le p\}$ and $R_2(s,p)=\{x_s|x_s>p\}$ to find the minimum of the standard squared-error loss objective function:

$$\min_{s,p}\left\{\min_{\overline{RV}_1}\sum_{x_s\in R_1(s,p)}\left(RV_i-\overline{RV}_1\right)^2 + \min_{\overline{RV}_2}\sum_{x_s\in R_2(s,p)}\left(RV_i-\overline{RV}_2\right)^2\right\}, \tag{5}$$

where the subscript i denotes those observations of realized variance, RV, that belong to a half-plane (we have dropped for convenience the time index), and $\overline{RV}_k=\text{mean}\{RV_i|x_s\in R_k(s,p)\}, k=1,2$ is a short-hand notation for the half-plane-specific mean of RV. Finding the minimum of the objective function defined in Equation (5) can be achieved by searching over all combinations of s and p, and, for any given combination, computing the loss-minimizing half-plane-specific means of realized variance. Once the solution to the minimization problem has been found, a researcher is equipped with information on the top-level optimal splitting predictor, its corresponding optimal splitting point, and the two half-plane specific means of realized variance (simple regression tree has two terminal nodes).

Next, the same minimization problem as given in Equation (5) for the top level of the regression tree is applied to find the optimal splitting predictors and splitting points for the two top-level half-planes. Solving these minimization problems gives up to two second-level optimal splitting predictors along with the corresponding optimal splitting points and four second-level region-specific means of realized variance. Upon recursively applying this partitioning scheme in a top-down and binary way, one obtains an increasingly complex, pyramid-like, hierarchical regression tree. The resulting regression tree is grown until it reaches a maximum number of terminal nodes set in advance by a researcher, or the terminal nodes contain a minimum number of observations.

Having grown a regression tree, a researcher can use its nodes, branches, and terminal leaves to trickle down the predictors of realized variance from the top level of the regression tree to its terminal leaves. Depending on which terminal leave is reached, the corresponding region-specific mean of realized variance forms the prediction of realized variance. Hence, when a regression tree has L regions, this prediction of realized variance simply can be computed as

$$T\left(\mathbf{x}_i,\{R_l\}_1^L\right)=\sum_{l=1}^{L}\overline{RV}_l\mathbf{1}(\mathbf{x}_i\in R_l), \tag{6}$$

where $\mathbf{1}$ denotes the indicator function.

While growing a large regression tree renders it possible to form increasingly granular predictions of realized variance, its increasingly complex hierarchical structure eventually results in an overfitting and data-sensitivity problem. This problem can be addressed by growing a random forest that consists of a large number of individual regression trees (Breiman, 2001). A random forest is grown by drawing a large number of bootstrap samples from the data and fitting a random regression tree to every bootstrap sample. The characteristic feature of a random regression tree is that it uses for every splitting step a random subset of the predictors and, in this way, mitigates the impact of influential predictors on tree building. Importantly, the fact that a random forest consists of a large number of random trees reduces the correlation of predictions obtained from the individual random regression trees. Finally, computing the average of the decorrelated predictions a researcher obtains from the individual random regression trees stabilizes the random-forest-based prediction of realized variance.

In our empirical analysis, we use the R language and environment for statistical computing (R Core Team, 2019), where we use the R add-on package "randomForestSRC" (Ishwaran & Kogalur, 2007, 2021; Ishwaran et al., 2008) to estimate random forests. We use random forests that consist of 500 individual random regression trees, where we set the minimum terminal node size to five observations. Moreover, we randomly select one third of the total number of predictors for splitting (which is the industry standard). We sample without replacement in our baseline scenario, but also shall present results for the case of sampling with replacement as a robustness check.

3.2 | Forecasting

Bootstrapping random regression trees has the further advantage that we can use the hold-out data of the bootstrap (also called out-of-bag data) for testing the predictive value of El Niño and La Niña events for realized variance. The out-of-bag data are those data that are not included in a given bootstrap sample. It follows that the out-of-bag data, which are not used for growing the random regression tree corresponding to a given bootstrap sample, render it possible to form out-of-sample forecasts of realized variance from that random regression tree. Because we draw a large number of bootstrap samples to grow a random forest, and any given bootstrap sample gives us a different sample of out-of-bag data, we can average the out-of-sample forecasts of realized variance across the various sampled out-of-bag data. The averaged out-of-bag forecasts form our final out-of-sample forecasts of realized variance.

Our approach to compute out-of-sample forecasts of realized variance based on the bootstrap-based out-of-bag data differs from the classic approach to use, for example, a recursive or rolling estimation window to compute out-of-sample forecasts. The logic underlying such a classic approach is to estimate a forecasting model on data included in a recursive or rolling estimation window, then to update the data to obtain an out-of-sample forecast, and finally to expand or shift the estimation window until the end of the sample period is reached. Such a classic approach, however, most likely will fail to detect a predictive value of El Niño and La Niña events for realized variance in the case of our data. The plot of our data in Figure 2 demonstrates that El Niño and La Niña events are relatively rare weather events (relative to the length of the sample period that we study) and that these events are, in case of La Niña, mostly centered at the beginning of our sample period. Hence, if we used a classic recursive or rolling estimation-window approach to forecast realized variance, a forecasting model that features La Niña events as a predictor would give, almost by construction, forecasts that are hardly discernible from forecasts obtained from a model that does not include La Niña events in the vector of predictors. In contrast, using out-of-bag data to compute out-of-sample predictions of realized variance has the advantage that we can use all El Niño and La Niña events, including those at the very beginning of our sample period, to form out-of-sample forecasts of realized variance. A forecasting experiment based on out-of-bag data, thus, has a greater power to recover a predictive value of El Niño and La Niña events for forecasting the realized variance of the returns of the agricultural commodity prices in our sample of data.

In order to formally study the predictive value of El Niño and La Niña events for realized variance, we use the test proposed by Clark and West (2007). In order to implement their test, we compute the quantity $f_{t+h} = (RV_{t+h} - \widehat{RV}_{A,t+h})^2 - [(RV_{t+h} - \widehat{RV}_{B,t+h}) - (\widehat{RV}_{A,t+h} - \widehat{RV}_{B,t+h})]$, where a hat denotes a forecast or realized variance at some forecast horizon h and the subscripts A and B denote to competing forecasting models. We shall present results for several short- and long-term forecast horizons by setting $h = 1, 4, 8, 12, 16$.[10] Model B is the larger model, that is, the model that features, in our case, El Niño and La Niña events as additional predictors of realized variance. We then regress the quantity f_{t+h} on a constant.[11] The Clark–West test rejects the null hypothesis of no difference in forecasting performance of Models A and B 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 study the significance of the t statistic).

In addition, we use variants of the R^2-statistic to study fit of the estimated models. To this end, we report an in-sample R^2 statistic and an out-of-bag R^2 statistic. In addition, we use the out-of-bag R^2 statistic to compare the forecasting performance of the estimated models. Specifically, we compute the out-of-bag statistic $R^2_{A,B} = 1 - \sum(RV_{t+h} - \widehat{RV}_{B,t+h})^2 / \sum(RV_{t+h} - \widehat{RV}_{A,t+h})^2$, where we use the out-of-bag forecasts to compute the sums. Model B (the larger model) outperforms Model A when we observe $R^2_{A,B} > 0$.

3.3 | Predictors

In addition to El Niño and La Niña events, we consider three groups of predictors. In total, our prediction model features 26 predictors, and random forests are ideally suited to capture the potentially complex links between the realized variance of movements of agricultural commodity prices and such a relatively large number of predictors.

Our first group of predictors consists of the current realized variance, the monthly realized variance, RV_m, and the quarterly, RV_q realized variance. The monthly realized variance is computed as the average realized variance from period $t-4$ to period $t-1$, and the quarterly realized variance is computed as the average realized variance from period $t-13$ to period $t-1$. The predictors in our first group form the constituent elements of the so-called heterogeneous autoregressive realized variance (HAR-RV) model of Corsi (2009). The HAR-RV model is one of the most popular models in the literature on modeling and forecasting realized variance.

Our second group of predictors consists of the realized variances of the respective 15 other agricultural commodities. We use this group of predictors to control for volatility spillovers across the markets for agricultural commodities. Cross-market volatility spillovers have been extensively studied (see, for example, Lahiani et al., (2013); Beckmann & Czudaj, (2014); Hernandez et al., (2014); Al-Maadid et al., (2017); Etienne et al., (2017); Chen & Wang, (2018); Bonato, (2019)).

Our third group of predictors consists of other predictors that have been widely studied in the literature on the modeling of realized variance: realized jumps, $JUMPS$, realized upside and downside tail risks, TR_u and TR_d, and realized skewness, RSK, as well as realized kurtosis, RKU. Following Amaya et al. (2015), we use RSK to capture the asymmetry of the returns distribution of agricultural commodity prices, while RKU accounts for extremes. We compute RSK as

$$RSK_t^d = \frac{\sqrt{M}\sum_{i=1}^{M} r_{(i,t)^3}}{RV_t^{3/2}}, \tag{7}$$

and RKU as

$$RKU_t^2 = \frac{M\sum_{i=1}^{M} r_{(i,t)^4}}{RV_t^2}. \tag{8}$$

The scaling of RSK and RKU by $(M)^{1/2}$ and M provides the corresponding daily skewness and kurtosis values. Weekly realized skewness and kurtosis are derived as weekly averages over the Monday to Friday aggregation.

In order to calculate the realized jumps, we utilize the formula derived by Barndorff-Nielsen and Shephard (2004) suggesting that realized variance converges into discontinuous (jump) and permanent components as

$$\lim_{M\to\infty} RV_t^2 d = \int_{t-1}^{t} \sigma^2(s)ds + \sum_{j=1}^{N_t} k_{t,j}^2, \tag{9}$$

where N_t is the number of jumps within day t and $k_{t,j}$ is the jump size. Equation (9) shows that RV_t is a consistent estimator of the jump contribution plus the integrated variance $\int_{t-1}^{t}\sigma^2(s)ds$. Using the asymptotic properties, Barndorff-Nielsen and Shephard ((2004), (2006)) further demonstrate that

$$\lim_{M\to\infty} BV_t^2 d = \int_{t-1}^{t} \sigma^2(s)ds, \tag{10}$$

where BV_t^d is the daily realized bipolar variation defined as

$$BV_t^d = \mu_1^{-2}\left(\frac{M}{M-1}\right)\sum_{i=2}^{M}|r_{t,i-1}||r_{i,t}| = \frac{\pi}{2}\sum_{i=2}^{M}|r_{t,i-1}||r_{i,t}|, \tag{11}$$

where

$$\mu_a = E(|Z|^a), Z\sim N(0,1), a>0. \tag{12}$$

Hence, using the continuous component of realized variance, we define the consistent estimator of the pure daily jump contribution as the following equation:

$$J_t^d d = RV_t^d - BV_t^d. \tag{13}$$

By applying the formal test estimator proposed by Barndorff-Nielsen and Shephard (2006), we test the significance of the jumps using the following test statistic:

$$JT_t = \frac{RV_t^d - BV_t^d}{(v_{bb} - v_{qq})\frac{1}{N}TP_t^d}, \tag{14}$$

where TP_t^d is the daily Tri-Power Quarticity

$$TP_t^d = M\frac{M}{M-2}\left(\frac{\Gamma(0.5)}{2^{2/3}\Gamma(7/6)}\right)\sum_{i=3}^{M}|r_{t,i}|^{4/3}|r_{t,i-1}|^{4/3}|r_{t,i-2}|^{4/3}, \tag{15}$$

which converges to Integrated Quarticity

$$IQ_t^d \to \int_{t-1}^{t} \sigma^4(s)ds, \tag{16}$$

even in the presence of jumps. We use the notation $v_{bb} = \left(\frac{\pi}{2}\right) + \pi - 3$ and $v_{qq} = 2$. It should also be noted that, for each t, $JT_t \sim N(0,1)$ as $M\to\infty$.

As can be seen in Equation (13), the jump contribution to RV_t^d is either positive or null. In order to avoid obtaining negative empirical contributions, we redefine the jump measure as (see also Zhou & Zhu, (2012)):

$$RJ_t^d = \max\left(RV_t^d - BV_t^d; 0\right). \tag{17}$$

The weekly jump contribution is obtained as average over the trading week (Monday to Friday) of the daily jump contribution. Last, we consider the Hill tail risk estimator (Hill, 1975). We consider $X_{t,i}$ the set of reordered intraday returns $r_{t,i}$ in such a way that

$$X_{t,i} \geq X_{t,j} \text{ for } i < j. \tag{18}$$

The Hill positive tail risk estimator (that is, our predictor TR_u) is then defined as

$$H_t^{up} = \frac{1}{k}\sum_{j=1}^{k} \ln\left(X_{t,i}\right) - \ln\left(X_{t,k}\right) \tag{19}$$

and the negative tail risk estimator (our predictor TR_d) as

$$H_t^{down} = \frac{1}{k}\sum_{j=n-k}^{n} \ln\left(X_{t,i}\right) - \ln\left(X_{t,n-k}\right) \tag{20}$$

where k the observation denoting the chosen α tail interval. We calculate the weekly tail risk as usual by averaging the daily tail risk over the Monday–Friday trading week.

4 | EMPIRICAL RESULTS

In order to illustrate the mechanics of a regression tree, we begin our empirical analysis by presenting in the upper left panel of Figure 3 an example of a stylized regression tree estimated on data for soybean oil (BO), where the forecast horizon is $h=1$. The regression tree is particularly simple because it features only three levels (in addition to the top level).

The example regression tree uses as its top-level splitting variable the realized monthly variance RV_m, where the optimal splitting point is shown next to the arrows. The arrow that points to the left applies in case the realized monthly variance takes on a value smaller than the splitting point. The arrow points in the direction of the lagged realized variance, RV_m, which in this example is one of the second-level splitting variables. The arrow that points in the right direction applies in case the realized monthly variance, RV_m, exceeds the splitting point. In this case, the arrow points in the direction of El Niño, which is used as the other second-level splitting variable.

For the optimal half-planes identified for the two second-level splitting variables, the third-level splitting variables and the third-level splitting points are identified. In this example, the third-level splitting variables are the realized variance of oats (O), the realized quarterly realized variance of BO (RV_q), and the downward tail risk. In the left-hand branch of the tree, the realized variance of oats is used twice as a splitting variable, which one is the relevant one depends on whether the realized variance of soybean oil assumes a value above or below its second-level splitting value. Hence, this example nicely illustrates the potentially complicated "path-dependency" that arises due to the hierarchical structure of a regression tree.

Finally, the regression tree in this stylized example reaches its terminal nodes, which are shown as green-colored cells in the figure, where the number of observations associated with a terminal node is also shown in the figure. Summing up across the terminal nodes gives the number of observations that are assigned to the various terminal nodes (that is, $72+28+68+\ldots+1=272$). The observations that are trickled down to the terminal nodes are the out-of-bag data, which for sampling without replacement are sampled using a 0.632 bootstrap of the total number of observations.

The upper right panel of Figure 3 plots the cumulated error rate (in this example, for soybean oil) implied by a random forest as a function of the number of random regression trees. The key point to take home from the plot is that the cumulated error rate attains a bottom plateau way before the maximum number of trees is reached. Hence, our decision to set the maximum of random regression trees that form a random forest to 500 should suffice to capture the key properties of the data.

The middle and lower panels of Figure 3 plot examples of partial dependence functions. The partial dependence functions in the two middle panels visualize the dependence of the realized variance of soybean oil at a forecast horizon of $h=1$ on El Niño (right-hand panel) and La Niña (left-hand panel) periods. Similarly, the lower panels plot partial dependence function for the realized variance of cocoa, in this example for a forecast horizon of $h=8$. The thin red lines are the smoothed plus/minus two-standard errors bands. The partial dependence functions illustrate a clear nonlinear dependence of the realized variance in the case of soybean oil on El Niño events and on La Niña events in the case of cocoa. The realized variance is more or less insensitive to small realizations of the two extreme weather events and starts increasing (decreasing) as an El Niño (a La Niña) weather event becomes more pronounced. At high realizations of El Niño and La Niña, in turn, the partial dependence functions tend flatten again. The partial

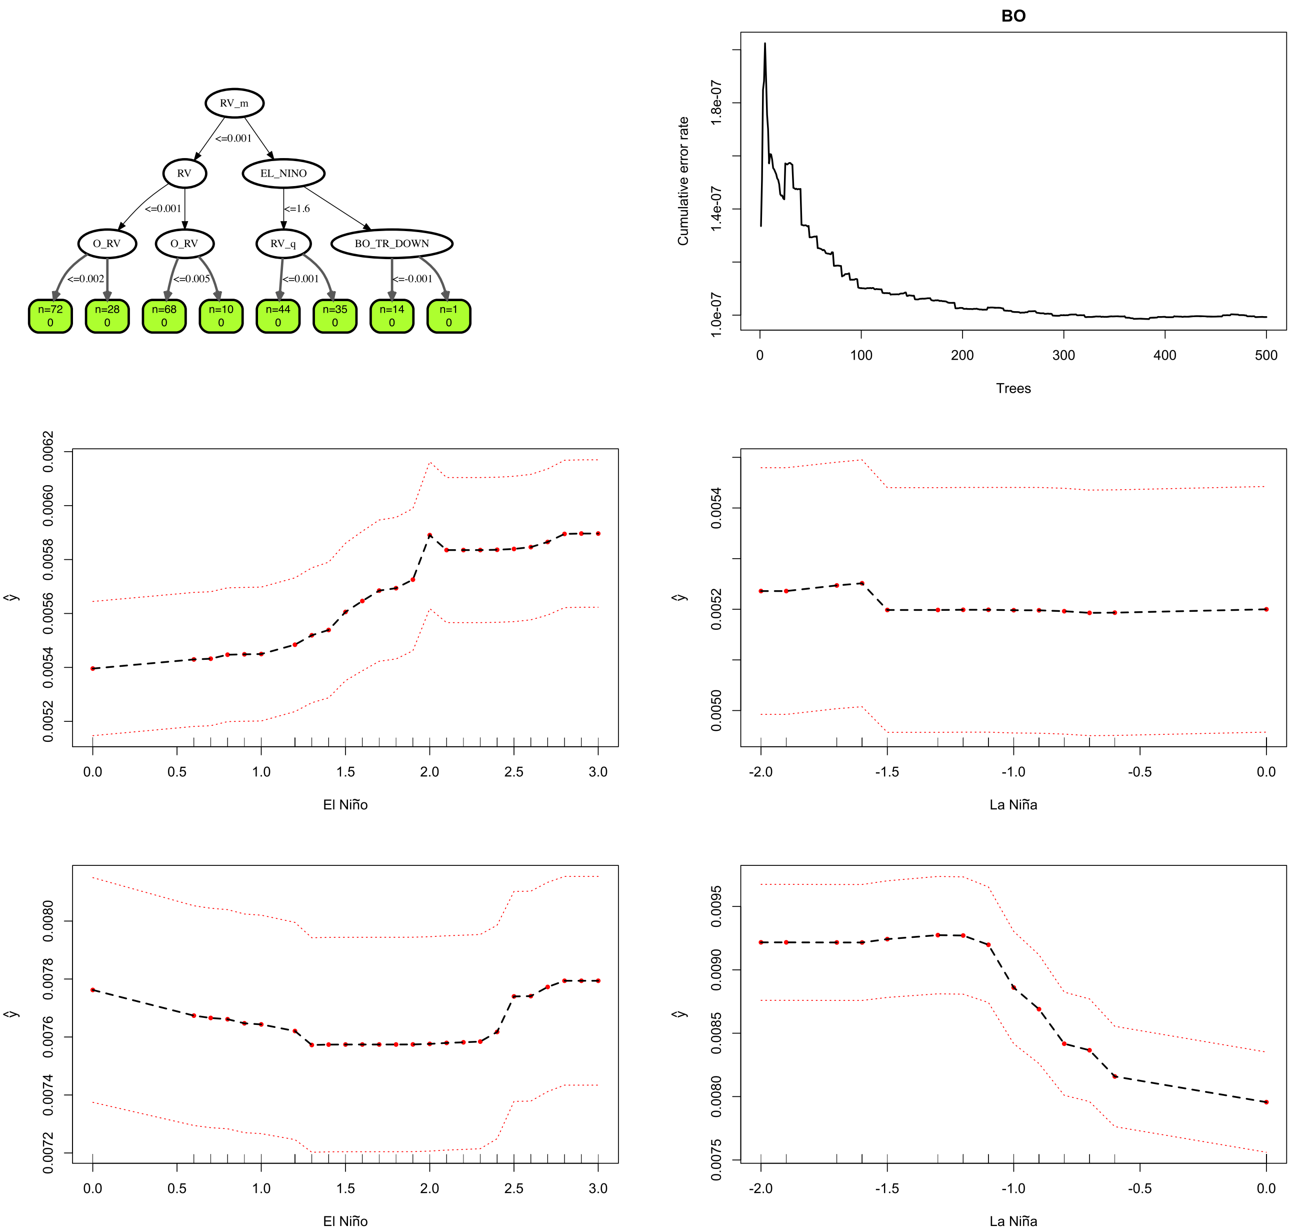

FIGURE 3 Illustration of tree mechanics. *Note*: Top-left panel: Example of a single tree estimated on data for BO. Top-right panel: Cumulative error rate as a function of the number of trees in a random forest estimated on data for BO. Middle-left panel: Partial-dependence function showing the dependence of RV of BO on El Niño. Middle-right panel: Partial-dependence function showing the dependence of RV of BO on La Niña. Lower-left panel: Partial-dependence function showing the dependence of RV of CC on El Niño. Middle-right panel: Partial-dependence function showing the dependence of RV of CC on La Niña. Red points/black dashed lines: partial values. Dashed red lines: smoothed error band (plus/minus two-standard errors)

dependence functions further show that the data do not support a nonlinear dependence of the realized variance on El Niño and La Niña events in all cases. In the plotted example, the partial dependence function with respect to La Niña events is more or less flat in the case of soybean oil, while the partial dependence function of the realized variance of cocoa with respect to El Niño events shows a marked albeit insignificant U-shaped pattern.

Table 3 informs about the fit of the estimated random forests, where we measure the model fit in terms of the familiar R^2 statistic. We present, however, two versions of the R^2 statistic. The first version (panel A) informs about the in-sample model fit. The second version (panel B), in contrast, informs about the more interesting out-of-bag model fit, that is, how the out-of-bag predictions of the model fit the corresponding out-of-bag realizations of the realized variance of our agricultural commodities. We report results for five different forecast horizons. Three key results emerge. First, the model fit shows a noticeable extent of cross-sectional heterogeneity, that is, "not all agricultural commodities are alike." Second, the in-sample and out-of-sample R^2 statistic tend to increase in the forecast horizon; that is, the random forests seem to have a stronger predictive power for the realized variance of agricultural commodity price movements at the longer forecast horizons. And, third, the out-of-bag R^2 statistic is smaller than the in-sample R^2 statistic, a result that is reminiscent of the well-known argument, as outlined earlier, that the ultimate test of a forecasting model is primarily its out-of-sample performance.

TABLE 3 In-sample and out-of-bag R^2 statistics

Commodity	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$
Panel A: In-sample results					
BO	0.7855	0.8694	0.8685	0.8653	0.8591
C	0.7234	0.8343	0.8480	0.8492	0.8563
CC	0.8322	0.9273	0.9417	0.9443	0.9412
CT	0.7636	0.8913	0.8972	0.8915	0.8931
GF	0.8302	0.8772	0.8890	0.8507	0.8421
HE	0.4213	0.6730	0.7620	0.8080	0.8323
KC	0.8033	0.8543	0.8704	0.8890	0.8947
LB	0.8279	0.9084	0.9317	0.9313	0.9318
LE	0.6669	0.8414	0.8501	0.7386	0.7424
O	0.7445	0.8004	0.8135	0.8310	0.8372
OJ	0.7647	0.8461	0.8574	0.8632	0.8854
RR	0.5637	0.8110	0.8561	0.9103	0.9223
S	0.6596	0.7745	0.8299	0.8580	0.8679
SB	0.7940	0.9134	0.9331	0.9391	0.9403
SM	0.7213	0.8249	0.8516	0.8555	0.8693
W	0.7407	0.8522	0.8654	0.8837	0.8772
Panel B: Out-of-bag results					
BO	0.2563	0.5370	0.5160	0.4965	0.4832
C	0.1733	0.4372	0.4432	0.4418	0.4503
CC	0.4736	0.7424	0.7787	0.7890	0.7732
CT	0.3799	0.6188	0.6349	0.6461	0.6396
GF	0.4349	0.5490	0.5704	0.4739	0.4149
HE	−0.1136	0.0593	0.1691	0.2517	0.3319
KC	0.3977	0.5232	0.5418	0.5475	0.6034
LB	0.4124	0.6636	0.7412	0.7397	0.7430
LE	0.1998	0.4373	0.4309	0.2461	0.1746
O	0.1119	0.3171	0.3547	0.3833	0.3676
OJ	0.2451	0.4586	0.5233	0.5588	0.5781
RR	0.1372	0.4052	0.5667	0.6749	0.7098
S	0.1033	0.2654	0.3656	0.4644	0.4867
SB	0.3692	0.6936	0.7605	0.7757	0.7701
SM	0.1187	0.3641	0.4501	0.4719	0.5029
W	0.1687	0.4754	0.5127	0.5585	0.5257

Note: Panel A depicts the in-sample R^2 statistic computed as $R^2_{in} = 1 - \sum_t (RV_t - \widehat{RV}_t)^2 / \sum_t (RV_t - \overline{RV}_t)^2$, where $\hat{y}_t$ denotes an in-sample forecast of RV and $\overline{RV}_t$ denotes the mean of RV.
Panel B depicts the out-of-bag R^2 statistic computed as $R^2_{oob} = 1 - \sum_t \left(RV_t - \widehat{RV}_t^{oob}\right)^2 / \sum_t (RV_t - \overline{RV}_t)^2$, where $\widehat{RV}_t^{oob}$ denotes an out-of-bag forecast of RV. Forecasts are computed using 500 regression trees, where sampling is done without replacement.
Abbreviations: BO, soybean oil; C, corn; CC, cocoa; CT, cotton; GF, feeder cattle; HE, lean hogs; KC, coffee; LB, lumber; LE, live cattle; O, oats; OJ, orange juice; RR, rough rice; S, soybeans; SB, sugar; SM, soybean meal; W, Chicago wheat.

It is for this reason that we focus in the following on the contribution of El Niño and La Niña events to the out-of-bag/out-of-sample of the random forest models. To this end, we use the test developed by Clark and West (2007). Table 4 summarizes the results. Because random forests are random by construction, we re-estimate the random forests for every agricultural commodity in our sample 25 times and then average over the resulting p-values. The key result of applying the Clark–West test is that the significance of the test results tends to strengthen as the forecast horizon increases. While the test results are insignificant for the majority of agricultural commodities for the two short forecast horizons, $h=1$ and $h=4$, almost all test results are significant for the three longer forecast horizons. Hence, a prediction model that includes El Niño and La Niña events in the vector of predictors tends to produce systematically more accurate forecasts then a prediction model that does not take into account the predictive value of El Niño and La Niña events for the subsequent realized variance of movements of agricultural commodity prices at approximately a quarterly and longer forecast horizon. In other words, our results suggest that El Niño and La Niña events, for the majority of agricultural commodities, do not have a systematic effect in the short term on the subsequent realized variance, but rather that they unfold their impact and predictive value in the medium to long term.

Another result that emerges from eye-balling Table 4 is that, as already witnessed by the results summarized in Table 3, there is again a certain extent of cross-sectional heterogeneity of the results in that the Clark–West test yields insignificant results even for $h=16$ for lean hogs (HE) and coffee (KC). In contrast, for some agricultural commodities like sugar (SB), we observe significant test results already for an intermediate forecast horizon ($h=8$), and for few agricultural commodities like rough rice (CC), the test result is even significant for $h=4$. Furthermore, tests are significant, for example, for soybean oil (BO) for all forecast horizons. The reason might be that El Niño conditions bring dry weather to the Pacific Rim, threatening palm oil crop conditions in countries

TABLE 4 Results of the Clark–West test

Commodity	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$
BO	0.0115	0.0193	0.0406	0.0412	0.0187
C	0.0969	0.1665	0.0562	0.0214	0.0148
CC	0.0226	0.0101	0.0036	0.0029	0.0008
CT	0.1348	0.0477	0.0148	0.0124	0.0190
GF	0.0604	0.0200	0.0350	0.0297	0.0223
HE	0.2791	0.2308	0.1191	0.1075	0.0942
KC	0.1123	0.0765	0.1439	0.1414	0.0990
LB	0.1344	0.0878	0.0269	0.0136	0.0076
LE	0.1220	0.0396	0.0199	0.0264	0.0656
O	0.1464	0.1082	0.0690	0.0544	0.0304
OJ	0.0732	0.0156	0.0064	0.0078	0.0013
RR	0.2098	0.1386	0.0599	0.0020	0.0014
S	0.1134	0.1300	0.0994	0.0670	0.0290
SB	0.0859	0.0021	0.0018	0.0038	0.0015
SM	0.1904	0.1088	0.0591	0.0398	0.0108
W	0.1117	0.1144	0.0840	0.0487	0.0171

Note: This table depicts the results (p-values; based on robust standard errors) of the Clark–West test for an equal mean-squared prediction error (MSPE). The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The benchmark model does not include El Niño and La Niña as predictors. Out-of-bag forecasts are computed using 500 regression trees. The random forests are re-estimated 25 times, and the p-values are averages over the 25 p-values obtained in this way.
Abbreviations: BO, soybean oil; C, corn; CC, cocoa; CT, cotton; GF, feeder cattle; HE, lean hogs; KC, coffee; LB, lumber; LE, live cattle; O, oats; OJ, orange juice; RR, rough rice; S, soybeans; SB, sugar; SM, soybean meal; W, Chicago wheat.

TABLE 5 Results of model comparisons based on the out-of-bag R^2 statistic

Commodity	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$
BO	0.0105	0.0668	0.0923	0.0667	0.0404
C	−0.0019	−0.0016	0.0030	0.0141	0.0008
CC	0.0002	0.1208	0.1224	0.1232	0.0883
CT	0.0105	0.0548	0.1012	0.1409	0.1807
GF	0.0018	0.0386	0.1222	0.1009	0.1156
HE	−0.0043	−0.0042	−0.0102	0.0033	0.0122
KC	0.0010	0.0119	0.0121	−0.0061	0.0186
LB	0.0000	−0.0001	0.0245	0.0149	0.0327
LE	−0.0111	0.0214	0.0544	0.0362	0.0480
O	0.0032	−0.0076	0.0082	0.0039	0.0153
OJ	0.0066	0.0200	0.0527	0.0408	0.0685
RR	−0.0001	0.0102	0.0389	0.0935	0.1090
S	0.0021	0.0015	−0.0073	0.0245	0.0061
SB	0.0023	0.0745	0.1315	0.1201	0.0964
SM	−0.0069	−0.0033	0.0115	0.0133	0.0300
W	−0.0148	0.0051	0.0021	0.0176	0.0295

Note: This table depicts the results of model comparisons based on the out-of-bag R^2 statistic. A positive value of the statistic shows that the rival model outperforms the benchmark model. The benchmark model does not include El Niño and La Niña as predictors. Out-of-bag forecasts are computed using 500 regression trees. The random forests are re-estimated 25 times, and the out-of-bag R^2 statistics are averaged across the 25 simulation runs.
Abbreviations: BO, soybean oil; C, corn; CC, cocoa; CT, cotton; GF, feeder cattle; HE, lean hogs; KC, coffee; LB, lumber; LE, live cattle; O, oats; OJ, orange juice; RR, rough rice; S, soybeans; SB, sugar; SM, soybean meal; W, Chicago wheat.

such as Indonesia and Malaysia. Hence, this reduction in global vegetable oil supplies increases the demand for soybean oil, and eventually results in larger fluctuations in soybean oil prices. Taken together, our results demonstrate again that "not all agricultural commodities are alike."

Table 5 summarizes the results that we obtain when we use the out-of-bag R^2 statistic to compare the estimated forecasting models. As for the Clark–West test, we re-estimate 25 times the random forests for our agricultural commodities and then average the out-of-bag R^2 statistics across simulations. We observe that the overwhelming majority of the statistics is positive, and we find cross-sectional heterogeneity with regard to the magnitude of the statistic. Moreover, the statistics take on small values for the short forecast horizon, but their magnitude tends to increase in the forecast horizon, thus corroborating the results for the Clark–West test.

The next step is to assess the robustness of our results. Table 6 summarizes the results of two robustness checks. As a first robustness check, we switch from sampling without replacement to sampling with replacement when bootstrapping the data. As second robustness check, we consider the case of random splitting. Random splitting implies that not all possible splitting points for a candidate splitting predictor are used for tree building, but rather that only a random subsample of splitting points is considered. Random splitting speeds up tree building and is a popular alternative to the "classic" (that is, deterministic) approach to tree building. The key message to take home from the two robustness checks is that, as in our baseline scenario laid out in Table 4, the predictive value of El Niño and La Niña events for the subsequent realized variance tends to strengthen when we move from a shorter to a longer forecast horizon and that for some agricultural commodities, the test results already are significant at the short forecast horizons.

As a final exercise, we study in Table 7 the predictive value of El Niño and La Niña events for the realized bad and realized good variance of movements of agricultural commodity prices. Corroborating the results reported in Tables 4 and 6, we observe a general tendency that, despite a certain extent of cross-sectional heterogeneity,

TABLE 6 Results of robustness checks

Commodity	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$
Panel A: Sampling with replacement					
BO	0.0280	0.0118	0.0234	0.0203	0.0208
C	0.0750	0.0363	0.0486	0.0120	0.0339
CC	0.0642	0.0073	0.0022	0.0014	0.0009
CT	0.1433	0.0245	0.0167	0.0101	0.0144
GF	0.1092	0.0258	0.0306	0.0275	0.0205
HE	0.2347	0.1555	0.1069	0.1009	0.0691
KC	0.0898	0.0867	0.0945	0.0849	0.1151
LB	0.0896	0.0875	0.0088	0.0331	0.0054
LE	0.1932	0.0290	0.0208	0.0457	0.0608
O	0.1047	0.0734	0.0325	0.0582	0.0189
OJ	0.0820	0.0129	0.0180	0.0123	0.0012
RR	0.1705	0.1194	0.0992	0.0114	0.0002
S	0.1605	0.1232	0.1366	0.1074	0.0177
SB	0.0310	0.0011	0.0004	0.0030	0.0025
SM	0.1561	0.0951	0.1282	0.0541	0.0137
W	0.1757	0.0513	0.0462	0.0072	0.0196
Panel B: Random splitting (10 random spilts)					
BO	0.0561	0.0283	0.0398	0.0310	0.0287
C	0.1399	0.0741	0.0481	0.0329	0.0050
CC	0.0624	0.0138	0.0041	0.0030	0.0011
CT	0.0271	0.0228	0.0261	0.0252	0.0265
GF	0.0745	0.0339	0.0433	0.0474	0.0251
HE	0.0699	0.2254	0.1087	0.1222	0.0709
KC	0.0797	0.1024	0.0480	0.0701	0.0712
LB	0.1397	0.0449	0.0134	0.0291	0.0305
LE	0.0831	0.0859	0.0152	0.0955	0.0563
O	0.0981	0.1431	0.0913	0.0428	0.0122
OJ	0.0740	0.0146	0.0098	0.0067	0.0017
RR	0.1380	0.1233	0.0259	0.0021	0.0006
S	0.1770	0.2157	0.1422	0.0513	0.0326
SB	0.0190	0.0035	0.0011	0.0031	0.0022
SM	0.1304	0.0862	0.0555	0.0263	0.0337
W	0.2004	0.0698	0.0582	0.0269	0.0053

Note: This table depicts the results (*p*-values; based on robust standard errors) of the Clark–West test for an equal mean-squared prediction error (MSPE). The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The benchmark model does not include El Niño and La Niña as predictors. Out-of-bag forecasts are computed using 500 regression trees. The random forests are re-estimated 25 times, and the *p*-values are averages over the 25 *p*-values obtained in this way.
Abbreviations: BO, soybean oil; C, corn; CC, cocoa; CT, cotton; GF, feeder cattle; HE, lean hogs; KC, coffee; LB, lumber; LE, live cattle; O, oats; OJ, orange juice; RR, rough rice; S, soybeans; SB, sugar; SM, soybean meal; W, Chicago wheat.

TABLE 7 Bad and good realized volatility

Commodity	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$
Panel A: Bad realized volatility					
BO	0.1354	0.0228	0.0362	0.0225	0.0783
C	0.0782	0.3570	0.2347	0.1415	0.1524
CC	0.1487	0.0102	0.0043	0.0029	0.0005
CT	0.1196	0.1074	0.0427	0.0247	0.0318
GF	0.0936	0.0308	0.0166	0.0130	0.0088
HE	0.2334	0.1487	0.0083	0.0371	0.0264
KC	0.1078	0.0398	0.0608	0.0491	0.0221
LB	0.2104	0.0135	0.0162	0.0195	0.0198
LE	0.2431	0.1705	0.0284	0.0117	0.0619
O	0.1732	0.2671	0.1473	0.0883	0.1237
OJ	0.1099	0.0525	0.0345	0.0523	0.0085
RR	0.1978	0.1870	0.1398	0.0839	0.0212
S	0.1351	0.3515	0.2133	0.1160	0.0713
SB	0.1065	0.0057	0.0036	0.0008	0.0003
SM	0.1203	0.2359	0.2962	0.2307	0.2052
W	0.1570	0.0776	0.0180	0.0420	0.0056
Panel B: Good realized volatility					
BO	0.1604	0.1663	0.0536	0.0448	0.0304
C	0.0680	0.0886	0.0859	0.0319	0.0141
CC	0.0994	0.0078	0.0070	0.0048	0.0020
CT	0.0136	0.0268	0.0191	0.0230	0.0191
GF	0.0512	0.0224	0.0222	0.0286	0.0329
HE	0.2015	0.1205	0.0234	0.0405	0.0401
KC	0.2845	0.1687	0.1863	0.0528	0.0874
LB	0.1671	0.0730	0.0637	0.0422	0.0812
LE	0.1569	0.0234	0.0207	0.0990	0.0542
O	0.1034	0.0898	0.1218	0.0835	0.0490
OJ	0.1322	0.0348	0.0294	0.0466	0.0016
RR	0.1084	0.0068	0.0020	0.0006	0.0005
S	0.0668	0.0686	0.0929	0.0393	0.0166
SB	0.1161	0.2308	0.1013	0.1201	0.0144
SM	0.0387	0.0745	0.0336	0.0399	0.0476
W	0.2658	0.0567	0.0442	0.0213	0.0161

Note: This table depicts the results (*p*-values; based on robust standard errors) of the Clark–West test for an equal mean-squared prediction error (MSPE). The alternative hypothesis is that the rival model has a smaller MSPE than the benchmark model. The benchmark model does not include El Niño and La Niña as predictors. Out-of-bag forecasts are computed using 500 regression trees. The random forests are re-estimated 25 times, and the *p*-values are averages over the 25 *p*-values obtained in this way.
Abbreviations: BO, soybean oil; C, corn; CC, cocoa; CT, cotton; GF, feeder cattle; HE, lean hogs; KC, coffee; LB, lumber; LE, live cattle; O, oats; OJ, orange juice; RR, rough rice; S, soybeans; SB, sugar; SM, soybean meal; W, Chicago wheat.

the predictive value of El Niño and La Niña events for the realized bad and the realized good variance strengthens as the forecast horizon increases. The test results for the realized bad and realized good variance, thereby, underscore the robustness of our main result.

5 | CONCLUDING REMARKS

We have studied the predictive value of El Niño and La Niña events for the realized variances of movements of 16 agricultural commodity prices. In order to do so, we have set up an empirical prediction model that includes in addition to El Niño and La Niña events the main elements of the popular HAR-RV model, cross-market spillover effects of realized variances, and realized jumps, realized skewness, realized kurtosis, and realized up and down tail risks. Random forests provide a natural modeling platform to estimate a predictive model with such a relatively large number of predictors. Random forests capture in a purely data-driven way potential interaction effects between the various predictors and, importantly, a potential nonlinear dependence of the subsequent realized variance on a predictor variable. We have shown how random forests work and how they can be used to study the predictive role of El Niño and La Niña events for forecasting the realized variances of movements of agricultural commodity prices at short- and long-term forecast horizons. The main result that emerges from our empirical analysis is that, while "not all agricultural commodities are alike" and the test results our significant already at a short or an intermediate forecast horizon for some agricultural commodities, there is a general tendency that the evidence of predictive value of El Niño and La Niña events strengthens at the longer term forecast horizons that we considered in our empirical research.

The El Niño and La Niña events are global climatological phenomena that occur periodically and are characterized by unusually warm/cold ocean temperatures in the equatorial pacific. A typical El Niño event causes heavy rainfall or flooding in the East Pacific (Argentina, Chile) and drought-like conditions in the west (Australia, Indonesia, India, and the Philippines). Due to the change in precipitation patterns, El Niño has important effects on crop production and, indirectly, on food prices and inflation. Considering that a significant El Niño effect would create upside inflation risks, our results suggest that central banks should analyze in detail any transitory/permanent effects of El Niño and La Niña events on food prices to adjust their policy stance more timely. Hence, our results provide useful insights to academics as well for potential causes of increases in inflation, as is being currently witnessed worldwide, by identifying which agricultural commodities have historically been the most sensitive to El Niño and La Niña events.

Because the variability of price fluctuations is a key input to investment decisions and portfolio choices in general, fund managers can generate sophisticated investment ideas focusing on how the volatility profile of the returns of agricultural commodity prices changes during El Niño and La Niña events, and invest in options position or structured products to profit from ensuing volatility patterns. Therefore, understanding which commodities are likely to be most affected by El Niño and La Niña events has the potential to enhance the performance of trading and investment strategies by capturing any El Niño- and La Niña-induced changes in risk premia ahead of the event.

As part of future analysis, it is interesting to investigate in detail the differential impacts of El Niño and La Niña events on the forecastability of economic activity and inflation of major agricultural commodity exporters, by accounting for historical data on precipitation patterns across these countries and their agricultural dependencies, given the existence of some in-sample evidence in this regard (see, for example, Cashin et al., (2017)). Furthermore, given that traditionally, the ENSO is measured at a monthly frequency, forecasting of daily or weekly realized variances can be performed based on HAR-RV models that are estimated using the reverse unrestricted-MIDAS approach (see, for instance, Foroni et al., (2018)).

ACKNOWLEDGMENTS

We would like to thank an anonymous referee for many helpful comments. However, any remaining errors are solely ours. Open Access funding enabled and organized by Projekt DEAL.

DECLARATION OF INTEREST

The authors declare not conflict of interest.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Oğuzhan Çepni https://orcid.org/0000-0003-0711-8880
Rangan Gupta https://orcid.org/0000-0001-5002-3428

ENDNOTES

[1] Of course, the role of the ENSO in affecting yields of agricultural commodities is also well recognized (see, for instance, Adams et al., 1999; Cadson et al., 1996; Hansen et al., 1998; Iizumi et al., 2014; Schlenker & Roberts, 2006; Tack & Ubilava, 2013, 2015; among others).

[2] The agricultural commodities that we study belong to the categories of Grains, Softs, and Livestock. Specifically, we study the following agricultural commodities: soybean oil, corn, cocoa, cotton, feeder cattle, lean hogs, coffee, lumber, live cattle, oats, orange juice, rough rice, soybean, sugar, soybean meals, and Chicago wheat.

[3] The reader is referred to Giot and Laurent (2003), Egelkraut and Garcia (2006), Elder and Jin (2007), Anderluh and Borovkova (2008), Triantafyllou et al. (2015), and Li et al. (2017) in this regard.

[4] Chatziantoniou et al. (2021) also use a HAR-RV model to forecast the monthly RV, rather than daily values of the same derived from intraday data, of agricultural commodities based on the volatility of oil. Unlike widespread in-sample evidence of volatility spillovers between agricultural commodities and the oil market (Luo & Ji, 2018), the authors could not detect out-of-sample forecasting gains emanating from both monthly and daily (which involved a mixed data sampling [MIDAS]) metrics of oil-price volatility.

[5] Most of the papers have focused on various forms of realized jumps as predictors (Luo et al., (2022), also incorporated the role of speculation), and we also include the role of extreme risks via the usage of realized skewness, realized kurtosis, and realized tail risks, which we believe are important controls when analyzing rare disaster risks as proxied by the ENSO.

[6] See https://www.fao.org/faostat/en/#home for details.

[7] We consolidate the data across commodities by date because we account in our empirical analysis for variance-spillover effects across agricultural commodities (see Section 3.3). In this regard, it should also be noted that we drop days with less than 20 observations to obtain meaningful tail risk statistics.

[8] Data source: https://www.cpc.ncep.noaa.gov/data/indices/.

[9] SSTA between 0.5°C (0.9°F) and −0.5°C (−0.9°F) are referred to as neutral ENSO events.

[10] We construct the data matrix in a way such that it has the same dimension for all forecast horizons.

[11] It follows that we treat the benchmark model as a restricted version of the larger model, but we note that comparing forecasts obtained from different random forests is complicated by their nonlinear complex structure. We, therefore, do not only use formal tests, but also out-of-bag R^2 statistics to asses relative model performance.

REFERENCES

Adams, R., Chen, C., McCall, B., & Weiher, R. (1999). The economic consequences of ENSO events for agriculture. *Climate Research*, *13*(3), 165–172.

Aït Youcef, C. (2019). How index investment impacts commodities: A story about the financialization of agricultural commodities. *Economic Modelling*, *80*, 23–33.

Al-Maadid, A., Corporale, G. M., Spagnolo, F., & Spagnolo, N. (2017). Spillovers between food and energy prices and structural breaks. *International Economics*, *150*, 1–18.

Amaya, D., Christoffersen, P., Jacobs, K., & Vasquez, A. (2015). Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics*, *118*, 135–167.

Anderluh, J., & Borovkova, S. (2008). Commodity volatility modelling and option pricing with a potential function approach. *European Journal of Finance*, *14*(2), 91–113.

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

Atems, B., McGraw, E., Maresca, M., & Ma, B. (2020). The impact of El Niño-Southern Oscillation on U.S. Agricultural Stock Returns. *Water Resources and Economics*, *32*, 100157.

Atems, B., & Sardar, N. (2021). Exploring asymmetries in the effects of El Niño-Southern Oscillationon U.S. food and agricultural stock prices. *Quarterly Review of Economics and Finance*, *81*, 1–14.

Bakas, D., & Triantafyllou, A. (2018). The impact of uncertainty shocks on the volatility of commodity prices. *Journal of International Money and Finance*, *87*, 96–111.

Bakas, D., & Triantafyllou, A. (2020). Commodity price volatility and the economic uncertainty of pandemics. *Economics Letters*, *193*, 109283.

Balcilar, M., Bouri, E., Gupta, R., & Pierdzioch, C. (2021). El Niño, La Niña, and the forecastability of the realized variance of heating Oil price movements. *Sustainability*, *13*, 7987.

Barndorff-Nielsen, O. E., Kinnebrouk, S., & Shephard, N. (2010). Measuring downside risk: Realised semivariance. In T. Bollerslev, J. Russell, & M. Watson (Eds.), *Volatility and time series econometrics: Essays in honor of Robert F. Engle*: Oxford University Press, pp. 117–136.

Barndorff-Nielsen, O. E., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, *2*, 1–37.

Barndorff-Nielsen, O. E., & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, *4*, 1–30.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, *121*(3), 823–866.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review*, *99*(1), 243–264.

Bastianin, A., Lanza, A., & Manera, M. (2018). Economic impacts of El Niño Southern Oscillation: Evidence from the Colombian Coffee Market. *Agricultural Economics*, *49*(5), 623–633.

Beckmann, J., & Czudaj, R. (2014). Volatility transmission in agricultural futures markets. *Economic Modelling*, *36*, 541–546.

Bonato, M. (2019). Realized correlations, betas and volatility spillover in the agricultural commodity market: What has changed? *Journal of International Financial Markets, Institutions and Money*, *62*(C), 184–202.

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

Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32.

Brunner, A. D. (2002). El Niño and world primary commodity prices: Warm water or hot air? *Review of Economics and Statistics*, *84*(1), 176–183.

Bruno, V. G., Büyükşahin, B., & Robe, M. A. (2016). The financialization of food? *American Journal of Agricultural Economics*, *99*(1), 243–264.

Cadson, R., Toyed, D., & Taylor, S. (1996). Midwestern corn yield and weather in relation to extremes of the southern oscillation. *Journal of Production Agriculture*, *9*(3), 347–352.

Campbell, J. Y. (2008). Viewpoint: Estimating the equity premium. *Canadian Journal of Economics*, *41*, 1–21.

Cashin, P., Mohaddes, K., & Raissi, M. (2017). Fair weather or foul? The macroeconomic effects of El Niño. *Journal of International Economics*, *106*, 37–54.

Chatziantoniou, I., Degiannakis, S., Filis, G., & Lloyd, T. (2021). Oil price volatility is effective in predicting food price volatility. Or is it? *Energy Journal*, *42*(6), 2. https://doi.org/10.5547/01956574.42.6.icha

Chen, Q., & Weng, X. (2018). Information flows between the us and china's agricultural commodity futures markets—Based on VAR-BEKK-Skew-t model. *Emerging Markets Finance and Trade*, *54*(1), 71–87.

Clark, T. D., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, *138*, 291–311.

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

Degiannakis, S., Filis, G., Klein, T., & Walther, T. (2022). Forecasting realized volatility of agricultural commodities. *International Journal of Forecasting*, *38*(1), 74–96. https://doi.org/10.1016/j.ijforecast.2019.08.011

Demirer, R., Gupta, R., Nel, J., & Pierdzioch, C. (2022). Effect of rare disaster risks on crude oil: Evidence from El Niño from over 140 years of data. *Theoretical and Applied Climatology*, *147*(1), 691–699.

Egelkraut, T. M., & Garcia, P. (2006). Intermediate volatility forecasts using implied forward volatility: The performance of selected agricultural commodity options. *Australian Journal of Agricultural and Resource Economics*, *31*(3), 508–528.

Elder, J., & Jin, H. J. (2007). Long memory in commodity futures volatility: A wavelet perspective. *Journal of Futures Markets*, *27*(5), 411–437.

Etienne, X. L., Trujillo-Barrera, A., & Hoffmann, L. A. (2017). Volatility spillover and time varying conditional correlation between DDGS, corn, and soybean meal markets. *Agricultural and Resource Economics Review*, *46*(3), 529–554.

Foroni, C., Guérin, P., & Marcellino, M. (2018). Using low frequency information for predicting high frequency variables. *International Journal of Forecasting*, *34*(4), 774–787.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, *127*(2), 645–700.

Gil-Alana, L. A., Cunado, J., & de Gracia, F. P. (2012). Persistence, long memory, and unit roots in commodity prices. *Canadian Journal of Agricultural Economics*, *60*(4), 451–468.

Giot, P., & Laurent, S. (2003). Market risk in commodity markets: A VaR approach. *Energy Economics*, *25*(5), 435–457.

Giot, P., Laurent, S., & Petitjean, M. (2010). Trading activity, realized volatility and jumps. *Journal of Empirical Finance*, *17*(1), 168–175.

Girardi, D. (2015). Financialization of food. Modelling the time-varying relation between agricultural prices and stock market dynamics. *International Review of Applied Economics*, *29*(4), 482–505.

Greb, F., & Prakash, A. (2015). Has price volatility changed? Food Outlook, Rome, Italy: Food and Agriculture Organization of the United Nations.

Hall, A., Skalin, J., & Terasvirta, T. (2001). A nonlinear time series model of El Niño. *Environmental Modelling and Software*, *16*(2), 139–146.

Hansen, J., Hodges, A., & Jones, J. (1998). ENSO influences on agriculture in the Southeastern United States. *Journal of Climate*, *11*(3), 404–411.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference and prediction* (2nd ed.). New York, NY: Springer.

Hernandez, M. A., Ibarra, R., & Trupkin, D. R. (2014). How far do shocks move across borders? Examining volatility transmission in major agricultural futures markets. *European Review of Agricultural Economics*, *41*(2), 301–325.

Hill, B. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, *3*, 1163–1173.

Iizumi, T., Lou, J. J., Challis, A. J., Samurai, G., Yokohama, M., Satsuma, H., & Yamagata, T. (2014). Impacts of El Niño Southern Oscillation on the global yields of major crops. *Nature Communications*, *5*, 1–7.

Ishwaran, H., & Kogalur, U. B. (2007). Random survival forests for R. *R News*, *7*(2), 25–31.

Ishwaran, H., & Kogalur, U. B. (2021). Fast unified random forests for survival, regression, and classification (RF-SRC). R package version 2.12.1.

Ishwaran, H., Kogalur, U. B., Blackstone, E. H., & Lauer, M. S. (2008). Random survival forests. *Annals of Applied Statistics*, *2*(3), 841–860.

Johnson, T. (2011). Food price volatility and insecurity. Council on Foreign Relations (CFR), Available for download from: https://www.cfr.org/backgrounder/foodprice-volatility-and-insecurity

Lahiani, A., Nguyen, D. K., & Vo, T. (2013). Understanding return and volatility spillovers among major agriculture commodities. *Journal of Applied Business Research*, *29*(6), 1781–1790.

Li, N., Ker, A., Sam, A. G., & Aradhyula, S. (2017). Modeling regime-dependent agricultural commodity price volatilities. *Agricultural Economics*, *48*(6), 683–691.

Luo, J., & Ji, Q. (2018). High-frequency volatility connectedness between the US crude oil market and China's agricultural commodity markets. *Energy Economics*, *76*, 424–438.

Luo, J., Klein, T., Ji, Q., & Hou, C. (2022). Forecasting realized volatility of agricultural commodity futures with infinite hidden Markov HAR models. *International Journal of Forecasting*, *38*(1), 51–73. https://doi.org/10.1016/j.ijforecast.2019.08.007

Marfatia, H. A., Luo, J., & Ji, Q. (2022). Forecasting the volatility of agricultural commodity futures: The role of co-volatility and oil volatility. *Journal of Forecasting*, *41*(2), 383–404. https://doi.org/10.1002/for.2811

McAleer, M., & Medeiros, M. C. (2008). Realized volatility: A review. *Econometric Reviews*, *27*, 10–45.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., & Pictet, O. V. (1997). Volatilities of different time resolutions—Analyzing the dynamics of market components. *Journal of Empirical Finance*, *4*, 213–239.

Ordu, B. M., Oran, A., & Soytas, U. (2018). Is food financialized? Yes, but only when liquidity is abundant. *Journal of Banking and Finance*, *95*, 82–96.

Ouyang, R., & Zhang, X. (2020). Financialization of agricultural commodities: Evidence from China. *Economic Modelling*, *85*, 381–389.

Qin, M., Qiu, L.-H., Tao, R., Umar, M., Su, C.-W., & Jiao, W. (2020). The inevitable role of El Nino: A fresh insight into the oil market. *Economic Research*, *33*, 1943–1962.

R Core Team. (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/

Rietz, T. A. (1988). The equity risk premium: A solution. *Journal of Monetary Economics*, *22*(1), 117–131.

Schlenker, W., & Roberts, M. (2006). Nonlinear effects of weather on corn yields. *Applied Economic Perspectives and Policy*, *28*(3), 391–398.

Shabbar, A., & Khandekar, M. (1996). The impact of El Niño-Southern Oscillation on the temperature field over Canada. *Atmosphere-Ocean*, *34*(2), 401–416.

Tack, J. B., & Ubilava, D. (2013). The effect of El Niño Southern Oscillation on U.S. corn production and downside risk. *Climatic Change*, *121*(4), 689–700.

Tack, J. B., & Ubilava, D. (2015). Climate and agricultural risk: Measuring the effect of ENSO on U.S. crop insurance. *Agricultural Economics*, *46*(2), 245–257.

Tian, F., Yang, K., & Chen, L. (2017a). Realized volatility forecasting of agricultural commodity futures using long memory and Regime switching. *Journal of Forecasting*, *36*(4), 421–430.

Tian, F., Yang, K., & Chen, L. (2017b). Realized volatility forecasting of agricultural commodity futures using the HAR model with timevarying sparsity. *International Journal of Forecasting*, *33*(1), 132–152.

Trenberth, K. E., Jones, P. D., Ambenje, P., Bojariu, R., Easterling, D., Tank, K. A., Parker, D., Rahimzadeh, F., Renwick, J. A., Rusticucci, M., Soden, B., & Zhai, P. (2007). Observations: Surface and atmospheric climate change, *Climate change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 235–336). Cambridge, UK: Cambridge University Press.

Triantafyllou, A., Dotsis, G., & Sarris, A. H. (2015). Volatility forecasting and time-varying variance risk premiums in grains commodity markets. *Journal of Agricultural Economics*, *66*(2), 329–357.

Ubilava, D. (2012a). El Niño, La Niña, and World coffee price dynamics. *Agricultural Economics*, *43*(1), 17–26.

Ubilava, D. (2012b). Modeling nonlinearities in the U.S. soybean-to-corn price ratio: A smooth transition autoregression approach. *Agribusiness: an International Journal*, *28*(1), 29–41.

Ubilava, D. (2014). El Niño Southern Oscillation and the fishmeal-soya bean meal price ratio: Regime-dependent dynamics revisited. *European Review of Agricultural Economics*, *41*(4), 583–604.

Ubilava, D. (2017). The ENSO effect and asymmetries in wheat price dynamics. *World Development*, *96*, 490–502.

Ubilava, D. (2018). The Role of El Niño Southern oscillation in commodity price movement and predictability. *American Journal of Agricultural Economics*, *100*(1), 239–263.

Ubilava, D., & Holt, M. T. (2013). El Niño Southern Oscillation and its effects on world vegetable oil prices: Assessing asymmetries using Smooth Transition Models. *Australian Journal of Agricultural and Resource Economics*, *57*(2), 273–297.

Yang, K., Tian, F., Chen, L., & Li, S. (2017). Realized volatility forecast of agricultural futures using the HAR models with bagging and combination approaches. *International Review of Economics and Finance*, *49*, 276–291.

Zhou, H., & Zhu, J. Q. (2012). An empirical examination of jump risk in asset pricing and volatility forecasting in China's equity and bond markets. *Pacific-Basin Finance Journal*, *20*(5), 857–880.

Živkov, D., Njegić, J., & Pećanac, M. (2019). Multiscale interdependence between the major agricultural commodities. *Agricultural Economics – Czech*, *65*(2), 82–92.

AUTHOR BIOGRAPHIES

Matteo Bonato is Senior Research Associate at the Department of Economics and Econometrics of the University of Johannesburg and at the IPAG Business School, Paris, France. His research focuses on volatility modeling and forecasting using high-frequency data.

Oğuzhan Çepni is a PhD fellow at Copenhagen Business School (CBS). Before joining CBS, he was an Economist at the Central Bank of Turkey with more than 6 years of experience. His primary research interests include applied macroeconomics, climate finance, empirical asset pricing, and forecasting.

Rangan Gupta is a Professor at the Department of Economics, University of Pretoria. He holds a PhD in Economics from the University of Connecticut, USA. He conducts research in the fields of macroeconomics and financial economics. His research has appeared in high impact factor international journals such as Energy Economics, International Journal of Forecasting, Journal of Banking and Finance, Journal of Forecasting, Journal of International Money and Finance, Macroeconomic Dynamics.

Christian Pierdzioch Professor of Economics, Helmut-Schmidt-University Hamburg, Germany, Department of Economics. Research interests: macroeconomics, monetary economics, forecasting, and financial econometrics.

How to cite this article: Bonato, M., Çepni, O., Gupta, R., & Pierdzioch, C. (2023). El Niño, La Niña, and forecastability of the realized variance of agricultural commodity prices: Evidence from a machine learning approach. *Journal of Forecasting*, *42*(4), 785–801. https://doi.org/10.1002/for.2914