Resources Policy 69 (2020) 101859

Contents lists available at ScienceDirect

Resources Policy

journal homepage: http://www.elsevier.com/locate/resourpol

A random walk through the trees: Forecasting copper prices using decision learning methods

Juan D. Díaz [a], Erwin Hansen [b,*], Gabriel Cabrera [c]

[a] Department of Management Control and Information Systems, Faculty of Economics and Business, University of Chile, Diagonal Paraguay 257, Of. 2103c, Santiago, Chile

[b] Department of Business Administration, Faculty of Economics and Business, University of Chile, Diagonal Paraguay 257, Of. 1204, Santiago, Chile

[c] Faculty of Economics and Business, University of Chile, Diagonal Paraguay 257, Santiago, Chile

ARTICLE INFO

Keywords:
Copper price
Forecasting
Decision learning methods
Tree-based methods
Random walk

ABSTRACT

We investigate the accuracy of copper price forecasts produced by three decision learning methods. Prior evidence (Liu et al. Resources Policy, 2017) shows that a regression tree, a simple decision learning model, can be used to predict copper prices for both short-term and long-term horizons (several days and several years, respectively). We contribute to this literature by evaluating more sophisticated decision learning methods based on trees: random forests and gradient boosting regression trees. Our results indicate that random forests and gradient boosting regression trees significantly outperform regression trees at forecasting copper prices. Our analysis also reveals that a random walk process, recognized in the literature as one of the most useful models for forecasting copper prices, yields competitive out-of-sample forecasts as compared to these decision learning methods.

1. Introduction

The ability to forecast copper prices is of interest to governments and practitioners alike. Governments of major copper-producing economies such as Chile[1] (where income from copper accounts for around a third of GDP and 50 percent of total exports) aim to predict the medium-run price of copper in order to prepare the national budget for the coming years. The Chilean government sets this future price annually on the basis of individual forecasts provided by an expert committee. From the investor's perspective, copper is one of the most traded commodities on the major futures trading exchanges such as the London Metal Exchange, the New York Commodity Exchange, and the Shanghai Futures Exchange (see Sanchéz-Lasheras et al., 2015). As such, accurate predictions of future copper prices are an indispensable foundation from which to establish profitable investment strategies. Finally, forecasts of copper prices and their volatility are a critical factor that must be taken into consideration when evaluation (the viability of) mining projects. The approval or rejection of a mining project depends on these critical predictions (see Dehghani and Ataee-pour, 2012; and Dehghani et al., 2014).

In a recent paper, Li et al. (2017) propose forecasting copper prices using a regression tree (RT) model, which is a simple decision learning method. Li et al. (2017) conclude that this machine-learning method produces accurate and reliable short-term and long-term copper price forecasts (for horizons of several days and years, respectively), with mean absolute errors (MAE) of out-of-sample predictions below 4 percent and root-mean-square errors (RMSE) of out-of-sample predictions below 8 percent, regardless of the forecast horizon. This is a surprising result, because short-term forecasts are typically more accurate than long-term forecasts.

In this article, we contribute to the literature on copper-price forecasting by analyzing two additional decision learning approaches: the random forest (RF) method (Ho, 1995; Amit; Geman, 1997; Breiman, 2001) and the gradient boosting regression trees (GBRT) method (Chapter 10 of Hastie et al., 2009). From a theoretical standpoint, these two machine-learning methods offer better prediction accuracy than the simple regression trees employed by Li et al. (2017); while from an applied point of view, they have been successfully used to produce

* Corresponding author.
E-mail addresses: juadiaz@fen.uchile.cl (J.D. Díaz), ehansen@fen.uchile.cl (E. Hansen), gcabrerag@fen.uchile.cl (G. Cabrera).

[1] Chile is the world's leading copper producer and accounts for roughly a third of total production worldwide. For example, in 2019, Chile produced 5789 thousand metric tons of the 20,830 thousand metric tons produced worldwide (see Commodity Markets Outlook, World Bank Report, April 2020, https://openknowledge.worldbank.org/bitstream/handle/10986/33624/CMO-April-2020.pdf).

https://doi.org/10.1016/j.resourpol.2020.101859
Received 3 July 2020; Received in revised form 14 August 2020; Accepted 4 September 2020
Available online 22 September 2020

forecasts in areas ranging from empirical finance to medical science. Therefore, we expect them to produce significantly more accurate predictions in our application as well. Besides evaluating the forecasting ability of RFs, GBRTs, and RTs (Li et al., 2017), we also analyze a random walk (RW) model. Prior literature (see, *e.g.*, Engel and Valdés, 2001) shows that the best predictor of future copper prices is often its current price. To the best of our knowledge, this is the first study to evaluate the ability of RF and GBRT decision-learning models to forecast copper prices.

We perform a comprehensive empirical analysis and compare the forecasting performance of the three decision-learning methods (RT, RF, and GBRT) with that of the RW model. We implement a genuine forecasting test to evaluate the accuracy of these methods. Indeed, we determine the accuracy of the models' forecasts by considering how well they perform on new data that were not used to initially fit them. In our first assessment of the models, we employ the same empirical framework as Liu et al. (2017). More specifically, we use the same target variable (the price of copper), the same set of predictors (lags of the price of copper, the lagged prices of gold, silver, crude oil, natural gas, lean hogs, and coffee, as well as the lagged value of the Dow Jones index), the same sample period, and the same training and testing windows. However, unlike Liu et al. (2017), our results show that the longer the range of the forecasts, the greater their error, regardless of the method under consideration, as we would expect. Our results also reveal that both the RF and the GBRT models produce better predictions of the price of copper than the RT model developed by Liu et al. (2017). Indeed, we report significant reductions in RMSE and MAE across all the specifications and forecast horizons we evaluate. Moreover, upon using an updated database, all our findings remain valid. Hence, our results provide evidence in favor of the implementation of more sophisticated decision-learning methods to forecast copper prices.

Turning now to the RW model, we find that its forecasts are the most accurate. It outperforms the three decision-learning models described above in almost all of the cases under evaluation. This evidence supports prior literature (Engel and Valdés, 2001) stating that the best predictor of future copper prices is its current price. From an economic point of view, this evidence suggests that regardless of the forecasting horizon, it is very unlikely future copper prices can be predicted from prior movements or trends in copper, gold, silver, crude oil, natural gas, lean hogs, or coffee prices, or based on the lagged values of the Dow Jones index. This is consistent with the spirit of the weak form of the efficient market hypothesis, because any attempt to use the set of predictors under consideration to forecast future copper price movements, either by implementing an RT, RF, or GBRT, is futile. Importantly, upon reiterating our analysis using an updated database of commodity prices, all our results remain valid.

Although our paper provides new empirical evidence regarding the ability of tree-based (RF and GBRT) decision-learning methods to forecast copper prices, it adds to a broader body of literature focusing on the ability of other decision-learning methods to predict commodity prices. For example, Dehghami (2018) assesses the predictive accuracy of the gene expression programming (GEP) algorithm when using it to forecast commodity prices, while Dehghani and Bogdanovi (2018) study copper predictability using a procedure combining traditional time-series analysis and the BAT algorithm. These studies highlight the superior forecasting accuracy of the decision-learning methods they evaluate as compared to the more traditional time-series and multiple-linear-regression methods.

Our paper is also related to studies using several Artificial Neural Network (ANN) methods (another type of decision-learning approach) to forecast copper prices. For instance, Sanchez-Lasheras et al. (2015) compare the predictive performance of two kinds of ANN: a multilayer perceptron ANN and an Elman Neural Network. The authors compare these models with standard ARIMA models and find that ANN models produce better forecasts of copper prices than the ARIMA models. Wang et al. (2019) introduce a hybrid predictive technique combining a complex network with traditional ANNs. Specifically, they combine a price volatility network (PVN) with three types of ANN and find that the hybrid method outperforms traditional ANN methods when forecasting copper prices. Additionally, Shi et al. (2011) study an empirical forecasting model that combines a wavelet method with an ANN. As compared with ARIMA models, the dynamic wavelet-ANN models produce better forecasts of copper prices. Furthermore, Garcia and Kristjanpoller (2019) present an adaptative approach to forecast copper price volatility by considering ARIMA and GARCH methods, as well as ANN, Fuzzy Inference Systems (FIS), and hybrid combinations of these. The authors report that an Adaptative-GARCH-FIS model produces the most accurate forecasts of copper volatility. Finally, Liu et al. (2020) introduce a hybrid decision (deep) learning method that combines variational mode decomposition (VMD) and long short-term memory network methods (LSTM) to predict zinc, copper, and aluminum prices. The VDM-LSTM model outperforms ARIMA models and other benchmark models.

Our work is also related to the literature aiming to forecast copper prices and volatility using ARIMA and GARCH models, respectively. For instance, Kriechbaumer et al. (2014) combine wavelet analysis (in which a time series is decomposed into its frequency and time domain) with ARIMA models to forecast monthly aluminium, copper, lead, and zinc prices. The combined method offers better forecasting accuracy than ARIMA methods. Moreover, Rubaszek et al. (2020) evaluate the ability of AR and VAR models, as well as a non-linear threshold VAR model to forecast the monthly price of aluminium, copper, nickel, and zinc. The authors find that commodity prices are mean-reverting and that univariate models produce better forecasts than multivariate models. Finally, Bundic and Moretto (2015) forecast monthly copper prices using a broad set of economic predictors, including autoregressive terms, using a dynamic-averaging and selection (DMA/DMS) approach. The authors find that several economic variables can contribute to forecasts of prices up to 6 months in advance. Beyond that range, a random walk benchmark produces superior forecasts.

Additionally, a related branch of the time-series literature focuses on forecasting copper volatility rather than prices. Methodologically, most of these studies use GARCH models or heterogeneous autoregressive (HAR) models. As is well known from the seminal work of Bolleslev (1986), GARCH models are well suited to the conditional dynamic of financial asset volatilities. By combining volatility estimates of the same asset at different frequencies, HAR models have also shown good forecasting ability. Notable studies using GARCH models to forecast copper volatility include those by Li and Li (2015), Khalifa et al. (2011), Hammoudeh and Yuan (2008), and Smith and Bracker (2003). Conversely, Gong and Li (2018), Todorova (2015), Todorova et al. (2014) and Lyocsa et al. (2017) have used HAR models to forecast realized copper volatilities. Meanwhile, Díaz et al. (2020) have added to the set of economic predictors used by Bundic and Moretto (2015) to forecast realized copper volatility. These authors find that fundamental copper market variables such as excess demand and several proxies for uncertainty (VIX index, EPU index, and geopolitical risk indices) are important predictors of monthly copper volatility in the short-run.

Lastly, our evidence regarding the predictive power of the RW model is akin to that of Buncic and Moretto (2015). These authors evaluated a macroeconomic factor-based model for predicting monthly copper returns. In results similar to ours, they find that the RW model produces competitive forecasts of copper returns (as compared to the macro-factor model) six to 12 months ahead. Engel and Valdés (2001) report similar evidence after comparing a set of 18 time-series models using quarterly and yearly data. In contrast to these studies, we rely on decision-learning models to compute our forecasts. Additionally, we perform our empirical analysis on daily, instead of monthly, quarterly, or yearly data.

The remainder of this paper is structured as follows. In section 2, we describe the three decision-learning methods used in the empirical analysis. In section 3, we present our database and provide descriptive statistics. In section 4, we report our empirical results and discuss their key implications. Finally, we conclude in section 5.

2. Methodology: Decision learning methods

In this section, we describe the decision learning methods considered in the forecasting exercise below. These models are also known in the machine learning literature as tree-based methods (Hastie et al., 2009; James et al., 2013). We begin by defining the regression trees (RT) used by Liu et al. (2017) since they constitute the benchmark model for our analysis. We subsequently introduce random forests (RF) and gradient boosted regression trees (GBRT).

2.1. Regression trees

A regression tree is the adaptation of a decision tree model (Breiman et al., 1984) for the regression case. Specifically, to forecast the price of copper, p, using a set of predictors, $X_1, X_2, \ldots, X_k$, a regression tree model is built by partitioning the possible values of the predictors into M different and disjoint regions, *e.g.*, $H_1, H_2, \ldots, H_M$. In order to predict the copper price of an observation belonging to any given region H_m, the simple average of the copper price values for the training observations in H_m, *i.e.*, $\widehat{p}_{H_m}$, is used.

In principle, the construction of regions $H_1, H_2, \ldots, H_M$ should minimize the residual sum of squares (*RSS*), which is given by the expression $RSS = \sum_{m=1}^{M} \sum_{i \in H_m} (p_i - \widehat{p}_{H_m})^2$. However, in practice, given that it is impossible to assess every possible split in predictor space due to limited computational resources, this strategy is not feasible. To overcome this issue, one alternative is to build the regions by growing a binary tree following a top-down strategy. This approach, also known as recursive binary splitting, begins at the top of the tree and successively splits the values of the predictors into two new branches further down the tree. Importantly, the cutoff value for the partition space of one predictor is optimally selected by minimizing the residual sum of squares of the training observations that are relevant to that specific partition (*i.e.*, not considering future binary splits of other predictors).

Once a large tree has been grown through recursive binary splitting (*e.g.*, with a maximum depth or minimum node size as the stopping criterion), the selection of the best subtree, *i.e.*, that of the partition (of the predictor space) which attained the lowest *RSS*, is performed by implementing a cost-complexity pruning process. This approach uses cross-validation to prune the original large tree and select the optimal subtree that leads to the lowest average forecast error. For comprehensive revisions of the regression tree methods, see Chapter 9 in Hastie et al. (2009), and Chapter 8 in James et al. (2013).

2.2. Random Forests

Random forests (Ho, 1995; Amit and Geman, 1997; Breiman, 2001) offer better prediction accuracy than regression trees. These models extend the single regression tree by building multiple regression trees based on random subsets of training data and predictors. Forecasts based on random forests average the resulting predictions of the multiple fitted regression trees.

The statistical procedure underpinning random forests is bootstrap aggregation (typically called bagging), which was developed by Breiman (1996). As is well known, bagging increases the prediction accuracy of any statistical method (*e.g.*, a single regression tree) by taking many bootstrapped subsamples from the training observations and fitting separate prediction models (regression trees) for each of those random subsamples. In bagging, the final prediction is the average (aggregation) of the resulting predictions, which reduces the variance. When bootstrap aggregation is applied to regression trees, these are often named bagged trees.

In addition to the advantages of bagging, the random forests method further improves prediction accuracy by reducing the correlation in predictive power between the regression trees that are fitted separately with the bootstrapped subsamples. This reduction is achieved when each time a bootstrapped subsample is considered, the regression tree is built using only a random sample of predictors (from the full set of predictors) at each binary split. Intuitively, a random forest fits multiple decorrelated regression trees, with each tree predicting with high variance, but low bias. Thus, averaging the predictions of these models reduces the overall variance. Note that if the forecasts of the multiple regression trees are too highly correlated (*i.e.*, if the predictive ability is similar across the regression trees), then it is not possible to reduce the variance when aggregating the predictions.

Random forest models have three adjustable (tuning) parameters: the number of predictors to randomly select for each bootstrapped subsample (we use 8); the total number of regression trees to be aggregated (in our case 1500); and the size (or depth) of each regression tree (we employ trees with 5 terminal nodes). A detailed description of the statistical theory behind random forests can be found in chapter 15 of Hastie et al. (2009), while chapter 8 of James et al. (2013) focuses on applications of this model.

2.3. Gradient Boosting Regression Trees

A gradient boosting regression tree (GBRT, Hastie et al., 2009, chapter 10) also provides better prediction accuracy than a single regression tree, but it employs a different approach to that of the random forest method. Intuitively, unlike random forests that reduce variance by aggregating complex decorrelated trees, this method sequentially decreases the bias of a simple regression tree with higher bias but low variance.

More formally, GBRTs are based on boosting (Freund and Schapire, 1996; Schapire, 2003), which is another general statistical technique for improving predictions generated from learning models (*e.g.*, regression trees). In GBRTs, an ensemble (aggregation) of multiple regression trees is also used to forecast the outcome variable of interest, but the main difference is that each tree is constructed using the knowledge acquired from previously fitted models. Specifically, each regression tree is built sequentially using the current model's residuals as the outcome variable. Then, this new regression tree is aggregated in the fitted model to update the residuals. The rationale behind this approach is that the new regression tree that is aggregated to the ensemble model corrects the errors of the current ensemble and, therefore, reduces the *RSS*. In GBRTs, the way the contemporary ensemble model is updated makes use of the same principle as that underpinning the gradient descent algorithm used in numerical optimization, but with the *RSS* as the objective function.

GBRTs have three tuning parameters: the number of trees to be considered in the ensemble model, the depth of each tree (the number of splits), and a shrinkage parameter related to the learning rate of the ensemble model. As Hastie et al. (2009) point out in their tenth chapter, although parameter configurations depend on the specificities of each case study, a value of 0.01 for the learning rate and regression trees with just one split work well. Finally, we use a cross-validation procedure to select the number of regression trees to be aggregated in the model.

3. Data and descriptive statistics

Following Liu et al. (2017), we build our database by collecting daily copper prices and a set of predictors covering the period between January 2008 and May 2020. As predictors, we consider lags of the price of copper, the lagged prices of a set of other commodities (mostly minerals commodities), and the lagged value of a stock market index in the USA (Dow Jones). We include gold, silver, crude oil, natural gas, lean hogs, and coffee in the set of commodities used as predictors. This particular selection of predictors of copper prices allows us to remain as close as possible to the work of Liu et al. (2017), our benchmark paper in this forecasting exercise. The authors discuss the economic rationale behind the selection of these predictors in their article and refer to prior evidence (see, *e.g.*, Charlot and Marimoutou, 2014; Joseph and Kunding, 1999; and Gargano and Timmermann, 2014). We download all the data from the website www.investing.com.

Table 1 reports descriptive statistics (panel A) and contemporaneous correlations (panel B) among the variables. The average price of the copper in the sample is US$ 3.22, with a standard deviation of US$ 0.65. The contemporaneous correlations between the price of copper and that of the other commodities are positive and statistically significant. The correlations between the price of copper and that of the other metals and energy commodities are stronger (*e.g.*, 0.56 with gold, 0.76 with silver, 0.74 with oil) than the correlation with the Dow Jones index (0.01).

4. Empirical results

In this section, we describe our empirical framework and report our forecasting results.

4.1. Framework

We implement two exercises to evaluate the out-of-sample forecasting accuracy of the different methods described in this paper. Throughout this article, we refer to these exercises as exercise A and exercise B. Specifically, in exercise A, we use the sample period beginning in January 2008 and ending in December 2015, while in exercise B, we employ the period from January 2008 until May 2020. The first time period is chosen so as to replicate the results of Liu et al. (2017) as closely as possible.

In each exercise (A and B), we consider several forecast horizons, which we express in terms of days ahead (h): 1 day ($h = 1$), 1 week ($h = 5$), 1 month ($h = 20$), 6 months ($h = 120$), 1 year ($h = 250$), and 2 years ($h = 500$)[2] ahead. Moreover, in each exercise, we also consider different augmentations of the set of predictors, by adding lags to the current values of the predictors. Specifically, we study the following configurations of the set of predictors (D): 1 (only the current values of the predictors); 5 (from the current value until the fourth lag of the predictors); 10 (from the current value until the ninth lag of the predictors); 20 (from the current values until the nineteenth lag of the predictors); 40 (from the current values until the thirty-ninth lag of the predictors); and 80 (from the current values until the seventy-ninth lag of the predictors).

More precisely, if p_t denotes the price of copper on day t and X_t is the vector collecting the values of the set of predictors on that day t, then for a given combination of h and D, we model p_{t+h} from X_t, X_{t-1}, X_{t-2}, …, X_{t-D-1}. For example, if $h = 1$ and $D = 1$, we fit the models with the outcome (response) variable p_{t+1} and predictors X_t. Alternatively, in the case of $h = 5$ and $D = 5$, we adjust the models with the outcome variable p_{t+5} and predictors X_t, X_{t-1}, X_{t-2}, X_{t-3}, and X_{t-4}; while if $h = 250$ and $D = 10$, we estimate the models with the outcome variable p_{t+250} and predictors X_t, X_{t-1}, X_{t-2}, X_{t-3}, …, X_{t-9}. In total, each exercise is composed of 36 cases.

In each exercise, we split the sample into training and testing subsamples. The training subsample is used to fit the model, while the testing subsample serves to evaluate the model's forecasting accuracy. Regardless of the values of h and D, we set the initial size of the training subsample in 607 observations, as in Liu et al. (2017), for exercises A and B. Note that, for different values of h and D, the number of observations in the testing subsamples varies from 677 observations, when $h = 500$ and $D = 80$ in exercise A, to 3135 observations, when $h = 1$ and $D = 1$ in exercise B.

Given h and D for each exercise, we make predictions using a rolling training sample, *i.e.*, each forecast in the testing sample uses a fixed number of previous observations for model estimation[3]. More precisely, to predict the copper price for the first date in the testing sample, we begin by estimating the model using the initial training sample. Then, to forecast the price of copper on the second date in the testing sample, we move forward the estimation window by one period. Subsequently, to predict the price of copper for the next date in the testing subsample, we move forward the model estimation window by one period, and so on. We stop once the price of copper has been forecasted for all of the dates in the testing sample. Lastly, we repeat this predictive approach using a rolling training sample for the four models considered in this article: regression trees, random forests, gradient boosted decision trees, and random walks[4].

A crucial difference between our forecasting exercise and the one in Liu et al. (2017) refers to the out-of-sample periods considered in the forecasting evaluation. While we only use information up to t to forecasting $t + h$ periods ahead, Liu et al. (2017) seem to be using information up to $t + h$*-1*. This difference in the forecasting framework explains differences in the magnitude of the reported forecasting errors, especially at the longer horizons.[5] We believe that the evaluation of the forecast accuracy of these methods must be based on a genuine forecast exercise, where the determination of the accuracy of forecasts must consider how well the models perform on new data that were not used when fitting the models.

For a combination of h and D, we consider the root-mean-square error (RMSE) and the mean absolute error (MAE) as measures of forecast error to evaluate the out-of-sample forecasting accuracy of the four methods. More formally, the RMSE and the MAE are given by the following expressions:

$$RMSE = \sqrt{\frac{1}{(T^* - t^* + 1)} \sum_{s=t^*}^{T^*} \left(\frac{100\%(p_s - \widehat{p}_s)}{p_s} \right)^2}, \ and\ MAE = \frac{1}{(T^* - t^* + 1)} \sum_{s=t^*}^{T^*} \left| \frac{100\%(p_s - \widehat{p}_s)}{p_s} \right| \quad (1)$$

where t^* is the start date of the testing subsample, T^* is the end date of the testing subsample, p_s is the actual price of copper on date s of the testing period and $\widehat{p}_s$ is the predicted price of copper on date s of the

[2] Liu et al. (2017) also consider a 4-year forecasting horizon. We exclude that horizon from the analysis because of the reduced number of periods available for the out-of-sample evaluation given our sample size and forecasting framework. As we explain below, our forecasting framework differs from the one in Liu et al. (2017).

[3] Alternatively, predictions can be based on an expanding training sample, *i.e.*, each forecast is produced by a model estimated with all of the previous observations.

[4] Specifically, we consider random walk model without drift.

[5] In the appendix, we report the forecast exercise using the approach employed in Liu et al. (2017). When we use the same forecasting framework, we replicate their results. In particular, as in Liu et al. (2017), the estimated mean absolute errors are always below 4% and the root-mean-squared errors are always below 8%, regardless of h and D.

Table 1
Descriptive Statistics and Correlations among variables.

Panel A: Descriptive Statistics

Variable	Mean	Std. Dev.	Min	Max
Copper	3.22	0.65	1.25	4.62
Gold	1304.37	267.94	704.90	1888.70
Silver	22.31	7.86	8.79	48.58
Oil	84.41	21.14	33.87	145.20
Gas	4.29	1.91	1.91	13.57
Lean Hogs	81.01	16.68	44.52	133.30
Coffee	162.58	45.68	101.50	304.90
Dow Jones	13229.73	2926.48	6547.05	18312.39

Panel B: Correlations

	Copper	Gold	Silver	Oil	Gas	Lean hogs	Coffee	Dow Jones
Copper	1.00							
Gold	0.56*	1.00						
Silver	0.77*	0.86*	1.00					
Oil	0.75*	0.30*	0.47*	1.00				
Gas	0.14*	-0.55*	-0.27*	0.43*	1.00			
Lean hogs	0.39*	0.53*	0.40*	0.51*	-0.17*	1.00		
Coffee	0.65*	0.58*	0.75*	0.34*	-0.14*	0.43*	1.00	
Dow Jones	0.01	0.28*	-0.02	0.01	-0.29*	0.46*	-0.04	1

Notes: Panel A reports descriptive statistics (mean, standard deviation, minimum, and maximum values) for the futures prices of several commodities and the index value of the Dow Jones index. Panel B reports pairwise correlations between the variables. * indicates a statistically significant pairwise correlation at the 95 percent confidence level. The sample period considered is from January 2008 until May 2020.

testing period using a specific model. For a combination of h and D, we report four RMSE values (and four MAE values), one for each of the models under consideration (regression trees, random forests, gradient boosted regression trees, and random walk). Naturally, the model with the lowest RMSE (MAE) is the best in terms of out-of-sample forecasting accuracy.

4.2. Results

Table 2 shows the results for exercise A, which use the sample mirroring that employed by Liu et al. (2017) covering the period between January 2008 and December 2015. As mentioned above, we compare and evaluate the models in terms of out-of-sample forecasting accuracy on the basis of the root-mean-square error (RMSE) and the mean absolute error (MAE). We report the values of these forecast error measures for several horizons: the 1-day, 1-week, 1-month, 6-month, 1-year, and 2-year forecast horizons. Liu et al. (2017) use similar evaluation metrics and forecast horizons, so our empirical results are comparable with theirs.

As explained above, the set of predictors considered in the forecasting exercise included lags of the price of copper, the lagged values of a number of other commodity prices (gold, silver, oil, gas, lean hogs, and coffee) and lagged values of a US stock index (the Dow Jones index). As discussed in section 4.1, the parameter D determines the number of lags considered in the set of predictors in each forecasting exercise. Table 2 shows the forecasting results for $D = 1$, 5, 10, 20, 40, and 80.

The first panel in Table 2 shows the results for the case in which $D = 1$. For the 1-day forecast horizon, the RMSE for the four forecasting models under evaluation, *i.e.*, RT, RF, GBRT, and RW, are 3.16%, 1.43%, 1.56%, and 1.25%, respectively. Comparing these values, we observe that both the RF and GBRT models provide significantly more accurate forecasts than the RT model used by Liu et al. (2017). At around 50%, the improvement in forecast accuracy for both models is considerable. When the RW model is included in the comparison, the overall forecasting results improve further. The RW model produces an RMSE of 1.25%, which is 74% lower than the value of the benchmark RT model. We obtain similar results when comparing the models using the MAE. In this case, the RT model produces an MAE of 2.27%, while the RF and GBRT models produce values of 1.01% and 1.11%, respectively. These values again indicate a sizeable and significant improvement in forecasting accuracy of around 50%. The RW model produces even better forecasting results by delivering an MAE of 0.86%, which represents a 62% improvement in forecasting accuracy as compared with the RT model. Overall, this evidence shows that more general tree-based regression models (RF and GBRT) outperform the RT model. Importantly, considering all of the models evaluated, our results indicate that the RW model delivers the best forecasting performance in the very short-term.

Although our results show that the forecast error increases as the forecast horizon increases, regardless of the method under analysis, it is noteworthy that the conclusions from the first panel of results remain qualitatively valid across all the other forecast horizons under evaluation, but with smaller quantitative differences in forecasting errors among the methods. For instance, for the 1-year horizon, we find that the RMSE of the RT model is 19.70%, while the RMSE of the RF and GBRT models are 17.80% and 17.25%, respectively, indicating an improvement in forecasting accuracy of around 12%. The RW model produces an RMSE of 11.84%, which is the lowest among all the models under consideration. Turning to the MAE metrics, the RT model provides an MAE of 16.41%, while the other models produce MAE values of 16.44% (RF), 15.98% (GBRT), and 9.85% (RW), respectively. Again, we observe substantial improvements in the accuracy of future copper price predictions using the RW model.

The conclusions for the results in the remaining panels, *i.e.*, for $D = 5$, 10, 20, 40, and 80, are almost identical to those obtained with $D = 1$. Again, for all of the prediction horizons considered, we observe that the RF, GBRT, or RW models provide significant improvements in forecasting accuracy over the RT model, regardless of the forecast error measure employed. The RW model consistently performs the best for out-of-sample predictions in most of the cases, and the RF model performs second best. The improvement in prediction accuracy varies between models, as well as according to the set of predictors included in each model (D), and the forecast horizons (h) considered.

Table 3 shows the results of exercise B. In this exercise, we use an updated database covering the period from January 2008 until May 2020. We observe that the use of this extended database does not affect the validity of most of the conclusions obtained in exercise A. In general, we find that short-term forecasts are more accurate than long-term forecasts, with more sophisticated tree-based models tending to yield more accurate forecasts of the price of copper than those generated from

Table 2
Copper price out-of-sample forecasting evaluation for alternative models and horizons. (Short sample, 2008:01-2015:12).

	RMSE				MAE			
	RT	**RF**	**GBM**	**RW**	**RT**	**RF**	**GBM**	**RW**
				D = 1				
1 day	3.16	1.43	1.56	1.25	2.27	1.01	1.11	0.86
1 week	4.49	3.36	3.48	2.83	3.38	2.51	2.61	2.12
1 month	7.76	6.80	6.47	5.28	5.87	5.12	4.99	4.06
6 months	13.37	13.06	12.10	10.19	11.20	11.10	10.05	8.20
1 year	19.70	17.80	17.25	11.84	16.41	16.44	15.98	9.85
2 years	20.18	19.52	18.89	16.89	17.69	17.19	16.54	15.21
				D = 5				
1 day	3.16	1.60	1.54	1.25	2.28	1.17	1.11	0.86
1 week	4.48	3.44	3.48	2.82	3.38	2.57	2.61	2.12
1 month	7.95	6.44	6.70	5.26	6.12	4.92	5.20	4.04
6 months	13.09	13.01	11.91	10.13	10.90	10.78	9.90	8.16
1 year	20.62	16.08	16.58	11.76	16.99	14.67	14.98	9.80
2 years	20.74	20.15	18.85	16.94	18.20	17.53	16.36	15.27
				D = 10				
1 day	3.17	1.74	1.52	1.25	2.28	1.28	1.09	0.86
1 week	4.53	3.53	3.41	2.81	3.38	2.62	2.57	2.11
1 month	8.15	6.51	6.78	5.26	6.20	4.96	5.24	4.05
6 months	12.86	12.99	11.79	10.01	10.74	10.75	9.88	8.10
1 year	19.34	15.73	16.24	11.68	15.78	14.12	14.44	9.74
2 years	20.82	20.42	19.08	16.97	18.03	17.58	15.98	15.31
				D = 20				
1 day	3.16	1.87	1.51	1.25	2.27	1.37	1.08	0.86
1 week	4.50	3.61	3.38	2.79	3.34	2.69	2.53	2.09
1 month	8.95	6.67	6.97	5.26	6.57	4.98	5.21	4.04
6 months	12.49	12.84	11.42	9.76	10.57	10.70	9.57	7.96
1 year	17.66	15.06	15.46	11.58	14.60	13.18	13.61	9.65
2 years	20.61	20.78	19.41	17.04	17.40	17.66	15.57	15.38
				D = 40				
1 day	3.15	1.97	1.48	1.23	2.26	1.46	1.06	0.84
1 week	4.51	3.65	3.28	2.74	3.34	2.75	2.46	2.06
1 month	9.66	6.68	7.26	5.28	6.98	5.12	5.48	4.06
6 months	12.38	12.53	10.63	9.57	10.63	10.50	9.00	7.81
1 year	17.73	14.58	15.35	11.34	14.91	12.66	13.51	9.46
2 years	22.30	21.01	20.33	17.25	19.14	17.84	16.74	15.65
				D = 80				
1 day	3.15	1.99	1.45	1.21	2.25	1.49	1.04	0.83
1 week	4.69	3.63	3.25	2.70	3.43	2.76	2.43	2.03
1 month	9.03	6.34	6.61	5.19	6.52	4.94	5.12	4.00
6 months	13.20	11.94	10.53	9.48	10.98	10.13	8.80	7.69
1 year	14.86	13.22	13.35	11.29	12.05	11.03	11.24	9.36
2 years	23.20	21.92	21.36	17.46	20.38	19.32	18.40	15.82

Notes: the table reports Root-Mean-Square Errors (RMSE) and Mean Absolute Errors (MAE) obtained from a Regression Tree (RT) model, a Random Forest model (RF), a Gradient-Boosting Regression Tree (GBRT) model, and a Random Walk (RW) aiming to predict Copper Prices. The predictions are computed for 1-day, 1-week, 1-month, 6-month, 1-year, and 2-year forecast horizons. The forecasts also consider the inclusion of D lags of the Price of Copper and D lags of a set of additional variables included in the forecasting exercise (see section 3 for a list of these variables). We evaluate the cases in which $D = 1$, 5, 10, 20, 40, and 80. The training period for the models rans from January 2008 until October 2011; while the out-of-sample forecasting period covers from November 2011 until December 2015, depending on the forecast horizon.

a simple regression tree. Specifically, both the RF and the GBRT models considerably reduce forecast errors, as compared with the RT model, regardless of the set of predictors included (D) and for all the forecast horizons (h) under consideration except for the very long-term (2 years), where we observe that the forecasting performance of the four models become similar. This exercise also reveals that the RW model produces more accurate predictions of copper prices for the 1-day, 1-week, 1-month, 6-month, and 1-year forecast horizons, independently of the sets of predictors used.

Overall, two key points emerge from these empirical results. Firstly, as compared to the RT model used by Liu et al. (2017), more sophisticated tree-based methods such as the RF and GBRT models appear to produce considerably better forecasts of copper prices. In general, the RF model outperforms the GBRT model. Secondly, our evidence suggests that the RW model provides very competitive forecasts of short-, medium-, and long-term copper prices. From an economic point of view, this last result is consistent with the weak version of the efficient market hypothesis. In other words, it is unlikely future copper prices can be predicted from the past movements or trends of copper prices (or from those of any of the other predictors considered here).

5. Conclusions

This article assesses the accuracy of copper-price forecasts from three tree-based decision-learning methods. In an earlier study, Liu at al. (2017) concluded that regression tree models produce accurate and reliable short-term and long-term copper-price forecasts (with horizons several days and years ahead, respectively). We compare the forecasting performance of the model presented by Liu et al. (2017) with that of two other tree-based models: random forests (RF) and gradient boosting regression trees (GBRT). We also evaluate the performance of a random walk (RW) model that prior literature (see, *e.g.*, Engel and Valdés, 2001) has identified as a competitive model for forecasting copper prices.

Using daily data from January 2008 until May 2020, our empirical analysis shows that both the RF and the GBRT models produce more accurate forecasts of copper prices than the RT model developed by Liu et al. (2017). We document significant improvements in root-mean-square errors (RMSE) and mean absolute errors (MAE),

Table 3
Copper price out-of-sample forecasting evaluation for alternative models and horizons. (Full sample, 2008:01-2020:06).

	RMSE				MAE			
	RT	**RF**	**GBM**	**RW**	**RT**	**RF**	**GBM**	**RW**
				D = 1				
1 day	2.94	1.32	1.45	1.15	2.16	0.93	1.01	0.79
1 week	4.12	3.15	3.21	2.62	3.12	2.35	2.40	1.97
1 month	7.17	6.44	6.15	5.09	5.43	4.82	4.69	3.89
6 months	14.41	13.98	13.03	10.37	11.93	11.75	10.75	8.20
1 year	23.51	22.51	21.49	14.41	19.44	19.17	18.12	12.04
2 years	23.23	22.61	23.58	23.68	19.75	19.38	20.14	19.42
				D = 5				
1 day	2.93	1.50	1.44	1.15	2.16	1.10	1.01	0.79
1 week	4.17	3.23	3.22	2.62	3.13	2.42	2.42	1.96
1 month	7.29	6.07	6.34	5.08	5.55	4.59	4.86	3.88
6 months	14.16	13.49	13.00	10.34	11.66	11.27	10.68	8.18
1 year	23.92	21.14	21.27	14.38	19.74	17.68	17.75	12.02
2 years	23.23	22.78	23.57	23.70	19.91	19.49	20.10	19.45
				D = 10				
1 day	2.95	1.63	1.42	1.15	2.17	1.20	1.00	0.79
1 week	4.26	3.30	3.19	2.61	3.18	2.48	2.40	1.96
1 month	7.49	6.11	6.43	5.08	5.65	4.63	4.93	3.88
6 months	14.33	13.48	12.95	10.28	11.75	11.24	10.58	8.14
1 year	23.46	21.10	21.18	14.35	19.25	17.55	17.60	11.99
2 years	23.41	22.86	23.79	23.72	20.01	19.52	20.10	19.47
				D = 20				
1 day	2.97	1.74	1.42	1.15	2.18	1.28	1.00	0.79
1 week	4.29	3.36	3.19	2.60	3.19	2.52	2.40	1.95
1 month	7.91	6.20	6.54	5.08	5.77	4.67	4.91	3.88
6 months	14.64	13.37	12.78	10.15	11.74	11.15	10.33	8.07
1 year	22.87	21.02	21.31	14.32	18.74	17.33	17.59	11.96
2 years	22.59	23.08	23.47	23.77	18.73	19.50	19.35	19.51
				D = 40				
1 day	2.96	1.83	1.40	1.13	2.16	1.36	0.99	0.78
1 week	4.30	3.41	3.13	2.57	3.20	2.56	2.37	1.93
1 month	8.28	6.24	6.74	5.09	5.97	4.75	5.09	3.88
6 months	14.40	13.16	12.60	10.06	11.43	10.93	10.13	8.00
1 year	24.04	21.53	22.09	14.25	19.98	17.52	18.17	11.89
2 years	23.28	22.95	22.86	23.87	19.19	19.14	18.45	19.65
				D = 80				
1 day	2.96	1.86	1.37	1.12	2.16	1.40	0.97	0.78
1 week	4.52	3.42	3.12	2.54	3.29	2.58	2.36	1.91
1 month	8.52	6.14	6.35	5.04	6.14	4.72	4.87	3.85
6 months	15.24	12.85	12.68	10.02	12.22	10.78	10.28	7.94
1 year	22.88	20.66	21.98	14.27	18.45	16.42	17.61	11.88
2 years	24.17	23.02	23.17	24.04	19.44	19.18	18.62	19.78

Notes: the table reports Root-Mean-Square Errors (RMSE) and Mean Absolute Errors (MAE) obtained from a Regression Tree (RT) model, a Random Forest model (RF), a Gradient-Boosting Regression Tree (GBRT) model, and a Random Walk (RW) aiming to predict Copper Prices. The predictions are computed for 1-day, 1-week, 1-month, 6-month, 1-year, and 2-year forecast horizons. The forecasts also considered the inclusion of D lags of the Price of Copper and D lags of a set of additional variables included in the forecasting exercise (see section 3 for a list of these variables). We evaluate the cases in which $D = 1$, 5, 10, 20, 40, and 80. The training period for the models rans from January 2008 until October 2011; while the out-of-sample forecasting period covers from November 2011 until May 2020, depending on the forecast horizon.

favoring both the RF and the GBRT models over the RT model. Moreover, our evidence shows that the RW model is very competitive when compared with the tree-based methods. From an economic standpoint, this evidence of the RW model's predictive power can be interpreted as a corroboration of the weak form of the efficient market hypothesis for the copper market. Indeed, it is unlikely future copper prices can be predicted from the past movements or trends of the predictors under consideration, irrespective of whether an RT, RF, or GBRT model is implemented.

Author statement

Juan D. Díaz: Conceptualization, Methodology, Formal analisis, Investigation, Writing - Original Draft, Writing - Review & Editing

Erwin Hansen: Conceptualization, Methodology, Formal analisis, Investigation, Writing - Original Draft, Writing - Review & Editing

Gabriel Cabrera: Software, Formal analisis, Investigation

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.resourpol.2020.101859.

Appendix(Not for Publication): Replication of Liu et al. (2017)

(Sample, 2008:01-2015:12).

	RMSE				MAE			
	RT	RF	GBM	RW	RT	RF	GBM	RW
D = 1								
1 day	3.16	1.43	1.56	1.25	2.27	1.01	1.11	0.86
1 week	4.04	2.16	2.59	1.25	3.05	1.54	1.95	0.86
1 month	5.12	2.89	3.18	1.25	3.68	1.90	2.39	0.86
6 months	4.93	2.64	3.18	1.21	3.67	1.78	2.37	0.83
1 year	5.19	2.86	3.18	1.21	3.50	1.82	2.33	0.82
2 years	4.37	2.56	2.92	1.09	3.02	1.70	2.20	0.75
D = 5								
1 day	3.16	1.60	1.54	1.25	2.28	1.17	1.11	0.86
1 week	4.03	2.15	2.54	1.25	3.04	1.57	1.90	0.86
1 month	5.20	2.63	3.17	1.24	3.77	1.85	2.39	0.85
6 months	4.87	2.51	3.12	1.21	3.57	1.81	2.36	0.83
1 year	5.55	2.49	3.26	1.21	3.73	1.73	2.38	0.82
2 years	4.29	2.28	2.91	1.09	2.95	1.65	2.19	0.75
D = 10								
1 day	3.17	1.74	1.52	1.25	2.28	1.28	1.09	0.86
1 week	3.98	2.18	2.42	1.25	2.99	1.58	1.83	0.86
1 month	5.07	2.53	3.08	1.23	3.67	1.82	2.29	0.85
6 months	4.97	2.38	3.05	1.21	3.63	1.74	2.30	0.83
1 year	5.54	2.32	3.11	1.21	3.63	1.64	2.28	0.82
2 years	4.51	2.17	2.83	1.08	3.10	1.60	2.15	0.75
D = 20								
1 day	3.16	1.87	1.51	1.25	2.27	1.37	1.08	0.86
1 week	3.97	2.15	2.32	1.24	2.98	1.57	1.73	0.85
1 month	5.20	2.43	3.00	1.23	3.63	1.76	2.17	0.84
6 months	4.72	2.21	2.70	1.22	3.44	1.67	2.08	0.83
1 year	5.47	2.18	2.84	1.21	3.63	1.56	2.09	0.82
2 years	4.54	1.95	2.60	1.07	3.09	1.49	1.99	0.74
D = 40								
1 day	3.2	2.0	1.5	1.2	2.26	1.46	1.06	0.84
1 week	4.0	2.1	2.2	1.2	2.95	1.57	1.62	0.84
1 month	6.1	2.3	2.9	1.2	3.87	1.70	2.13	0.84
6 months	4.3	2.1	2.5	1.2	3.10	1.59	1.89	0.83
1 year	5.3	2.1	2.7	1.2	3.52	1.50	1.92	0.82
2 years	5.5	1.8	2.3	1.1	3.57	1.36	1.75	0.74
D = 80								
1 day	3.15	1.99	1.45	1.21	2.25	1.49	1.04	0.83
1 week	3.98	2.06	2.06	1.21	2.95	1.55	1.55	0.83
1 month	5.79	2.12	2.62	1.21	3.68	1.60	1.92	0.83
6 months	4.60	2.05	2.30	1.21	3.20	1.53	1.71	0.82
1 year	4.86	2.00	2.59	1.20	3.22	1.48	1.84	0.81
2 years	5.24	1.72	2.16	1.05	3.42	1.33	1.64	0.73

Notes: the table reports Root-Mean-Square Errors (RMSE) and Mean Absolute Errors (MAE) obtained from a Regression Tree (RT) model, a Random Forest model (RF), a Gradient-Boosting Regression Tree (GBRT) model, and a Random Walk (RW) aiming to predict Copper Prices. The predictions are computed for 1-day, 1-week, 1-month, 6-month, 1-year, and 2-year forecast horizons. The forecasts also considered the inclusion of D lags of the Price of Copper and D lags of a set of additional variables included in the forecasting exercise (see section 3 for a list of these variables). We evaluate the cases in which $D = 1, 5, 10, 20, 40,$ and 80. The training period for the models rans from January 2008 until October 2011; while the out-of-sample forecasting period covers from November 2011 until December 2015.

References

Amit, Y., Geman, D., 1997. Shape quantization and recognition with randomized trees. Neural Comput. 9 (7), 1545–1588.

Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 1984. Classification and Regression Trees. Wadsworth International Group, Belmont, CA, USA.

Breiman, L., 1996. Bagging predictors. Mach. Learn. 24 (2), 123–140.

Breiman, L., 2001. Random forests. Mach. Learn. 45 (1), 5–32.

Buncic, D., Moretto, C., 2015. Forecasting copper prices with dynamic averaging and selection models. N. Am. J. Econ. Finance 33, 1–38.

Charlot, P., Marimoutou, V., 2014. On the relationship between the prices of oil and the precious metals: revisiting with a multivariate regime-switching decision tree. Energy Econ. 44, 456–467.

Dehghani, H., 2018. Forecasting copper price using gene expression programming. J. Min. Environ. 9 (2), 349–360.

Dehghani, H., Ataee-pour, M., 2012. Determination of the effect of operating cost uncertainty on mining project evaluation. Resour. Pol. 37 (1), 109–117.

Dehghani, H., Ataee-pour, M., Esfahanipour, A., 2014. Evaluation of the mining projects under economic uncertainties using multidimensional binomial tree. Resour. Pol. 39, 124–133.

Dehghani, H., Bogdanovic, D., 2018. Copper price estimation using bat algorithm. Resour. Pol. 55, 55–61.

Díaz, J., Hansen, E., Cabrera, G., 2020. What do economic variables tell us about commodity volatility? The case of copper. Working paper, University of Chile.

Engel, E., Valdés, R., 2001. Prediciendo el precio del cobre: ¿Más allá del camino aleatorio? No. 100.

Freund, Y., Schapire, R.E., 1996. Experiments with a New Boosting Algorithm. Machine Learning: Proceedings of the Thirteenth International Conference, July 3–6, 1996, Bari Italy. Morgan Kaufman, San Francisco, CA, USA, pp. 148–156.

García, D., Kristjanpoller, W., 2019. An adaptive forecasting approach for copper price volatility through hybrid and non-hybrid models. Appl. Soft Comput. 74, 466–478.

Gong, X., Lin, B., 2018. Structural breaks and volatility forecasting in the copper futures market. J. Futures Mark. 38 (3), 290–339.

Hammoudeh, S., Yuan, Y., 2008. Metal volatility in presence of oil and interest rate shocks. Energy Econ. 30 (2), 606–620.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer.

Ho, T., 1995. Random decision forests. Proc. 3rd Int. Conf. Document Anal. Recogn. 1, 278–282.

James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning with Applications in R. Springer.

Joseph, G., Kundig, K.J.A., 1999. Copper: its Trade. Manufacture, Use and Environmental Status.

Khalifa, A.A., Miao, H., Ramchander, S., 2011. Return distributions and volatility forecasting in metal futures markets: evidence from gold, silver, and copper. J. Futures Mark. 31 (1), 55–80.

Kriechbaumer, T., Angus, A., Parsons, D., Casado, M.R., 2014. An improved wavelet–ARIMA approach for forecasting metal prices. Resour. Pol. 39, 32–41.

Li, G., Li, Y., 2015. Forecasting copper futures volatility under model uncertainty. Resour. Pol. 46, 167–176.

Liu, C., Hu, Z., Li, Y., Liu, S., 2017. Forecasting copper prices by decision tree learning. Resour. Pol. 52, 427–434.

Liu, Y., Yang, C., Huang, K., Gui, W., 2020. Non-ferrous metals price forecasting based on variational mode decomposition and LSTM network. Knowl. Base Syst. 188, 105006.

Lyócsa, Š., Molnár, P., Todorova, N., 2017. Volatility forecasting of non-ferrous metal futures: covariances, covariates or combinations? J. Int. Financ. Mark. Inst. Money 51, 228–247.

Rubaszek, M., Karolak, Z., Kwas, M., 2020. Mean-reversion, non-linearities and the dynamics of industrial metal prices. A forecasting perspective. Resour. Pol. 65, 101538.

Sánchez-Lasheras, F., de Cos Juez, F.J., Sánchez, A.S., Krzemień, A., Fernández, P.R., 2015. Forecasting the COMEX copper spot price by means of neural networks and ARIMA models. Resour. Pol. 45, 37–43.

Schapire, R., 2003. In: Denison, D.D., Hansen, M.H., Holmes, C., Mallick, B., Yu, B. (Eds.), The Boosting Approach to Machine Learning – an Overview. *MSRI Workshop on Nonlinear Estimation and Classification*. Springer, New York.

Shi, L., Chu, L.K., Chen, Y.H., 2011. Dynamic wavelet neural network model for forecasting returns of SHFE copper futures price. Proceedings of the 7th International Conference on Digital Enterprise Technology, DET 2011.

Smith, K.L., Bracker, K., 2003. Forecasting changes in copper futures volatility with GARCH models using an iterated algorithm. Rev. Quant. Finance Account. 20 (3), 245–265.

Todorova, N., 2015. The course of realized volatility in the LME non-ferrous metal market. Econ. Modell. 51, 1–12.

Todorova, N., Worthington, A., Souček, M., 2014. Realized volatility spillovers in the non-ferrous metal futures market. Resour. Pol. 39, 21–31.

Wang, C., Zhang, X., Wang, M., Lim, M.K., Ghadimi, P., 2019. Predictive analytics of the copper spot price by utilizing complex network and artificial neural network techniques. Resour. Pol. 63, 101414.