Nuclear Engineering and Technology 52 (2020) 1334–1339

Contents lists available at ScienceDirect

Nuclear Engineering and Technology

journal homepage: www.elsevier.com/locate/net

Original Article

Forecasting uranium prices: Some empirical results

Diego J. Pedregal

ETSI Industriales, Universidad de Castilla-La Mancha, Ciudad Real, Spain

ARTICLE INFO

Article history:
Received 5 July 2019
Received in revised form
21 November 2019
Accepted 22 November 2019
Available online 27 November 2019

Keywords:
Uranium price
Unobserved components models
Exponential smoothing
ARIMA
Theta method
Artificial neural networks
Combination of forecasts

ABSTRACT

This paper presents an empirical and comprehensive forecasting analysis of the uranium price. Prices are generally difficult to forecast, and the uranium price is not an exception because it is affected by many external factors, apart from imbalances between demand and supply. Therefore, a systematic analysis of multiple forecasting methods and combinations of them along repeated forecast origins is a way of discerning which method is most suitable. Results suggest that i) some sophisticated methods do not improve upon the Naïve's (horizontal) forecast and ii) Unobserved Components methods are the most powerful, although the gain in accuracy is not big. These two facts together imply that uranium prices are undoubtedly subject to many uncertainties.

1. Introduction

Nuclear power is the largest source of energy with low carbon dioxide emissions after hydropower [1,2]. Its development can facilitate programmes to improve energy security and environmental protection. However, this is not often clearly perceived, due to challenges posed by the high investment costs of power plant construction, unresolved issues of storage and disposal of nuclear waste and general public concern for the safety of nuclear power plants (a disaster such as Fukushima in 2011 remains in the collective memory for a long time) [2].

However, there are other facts that recommend the expansion of nuclear energy, such as the need to reduce greenhouse gas emissions, the rapid growth in energy demand in developing countries and the expected rise in fossil fuel prices. These facts make some uranium unconventional recovery techniques, such as recovery from phosphate rock or seawater, very relevant to the future of the nuclear power production [3,4].

In between the pros and cons is the uranium market deciding the price of uranium which, although representing only up to 3–5% of the total price of nuclear generated electricity [1,2,5] (5–10% according to other sources [6]), remains important for the sustainability of the industry.

The results of market interactions are shown in Fig. 1, that unfolds the evolution of the monthly spot price, specifically the 'Uranium NUEXCO Restricted Price, Nuexco exchange spot' measured in US dollars per pound [7]. It shot up to almost $135/pound in 2007 due to imbalances between uranium supply and demand. Since then, prices have always declined with a short period of growth. The Fukushima disaster in 2011 definitely marked a decline in prices.

There are not many studies about uranium price forecasting, the most recently cited documents are listed in Table 1 (see also references cited in these documents). Both time series and Artificial Intelligence methods focus, by their very nature, on short-term prediction, only a few months and up to a year in advance. Methods based on formal models of how future imbalances between supply and demand will affect prices are often useful in medium and long-term forecasts.

Two shortcomings, from the point of view of the author, are common to the previous studies, namely that the final conclusions are based on forecasts for a single forecasting origin, and they focus on one single method. This is quite unusual in the forecasting literature, where the results are usually provided as an average or any other combination of forecasts estimated across many forecasting origins, and the accuracy of the proposed methods are usually measured against other common alternatives that can serve as benchmarks. The number of examples is immense and would

E-mail address: Diego.Pedregal@uclm.es.
URL: https://blog.uclm.es/diegopedregal/

https://doi.org/10.1016/j.net.2019.11.028

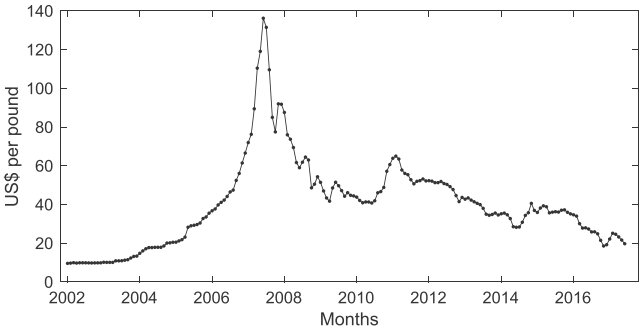

Fig. 1. Monthly uranium price from January 2002 to June 2017.

Table 1
Literature on uranium price forecasting.

Method	References
Time series or statistical models	[8]
Artificial Intelligence	[9–11]
Demand and supply analysis	[1,2,12–14]

roughly include the use of the most typical statistical methods, such as regression, ARIMA, exponential smoothing or Theta methods (see, for example [15–18]). Other relevant methods are those that actually produce a transition between different methods (for example [19–21] and references therein). Finally, Machine Learning techniques have flooded the forecasting area with results that are sometimes strongly questioned by researchers in the field (e.g. Refs. [22–27]).

This paper continues this tradition by using different methods on uranium prices to produce short-term forecasts with the ultimate goal of selecting the best option, if there is an absolute winner. Methods include Naïve, Autoregressive Integrated Moving Average (ARIMA), ExponenTial Smoothing (ETS), Unobserved Components (UC), Theta method, two ANNs and combinations of the above. All of them are systematically estimated for different forecasting origins and the results are summarized with two different global metrics for forecasting horizons ranging from one month to one year.

As far as the author is concerned, this is the first time that such a comprehensive analysis of uranium prices has been carried out in academic literature, both in terms of the number and variety of methods compared and in terms of the number of forecasting origins.

2. Forecasting methods

2.1. Benchmark

The typical method that is very often considered as a benchmark is the so-called Naïve forecast. It consists of projecting the last available data point in the future. It is actually the optimal forecast in the particular case of a Random Walk (RW), a model in which the value of a series at each point in time is the previous one plus a random noise, generally assumed Gaussian with zero mean and constant variance. Taking z_t as a time series measured in time t and a_t as a white noise, a RW obeys the equation $z_t = z_{t-1} + a_t$.

2.2. Autoregressive Integrated Moving Average

Since they were popularized in the early 70's by the first edition of [28], ARIMA models have received a great deal of attention. The formulation of a non seasonal ARIMA (p,d,q) is given in equation (1), where B is the back-shift operator such that $B^l z_t = z_{t-l}$; $\theta_i (i = 1,2,\ldots,q)$ and $\phi_j (j = 1,2,\ldots,p)$ are parameters to estimate; and a_t is a Gaussian white noise with zero mean and constant variance.

$$z_t = \frac{(1 + \theta_1 B + \ldots + \theta_q B^q)}{(1 + \phi_1 B + \ldots + \phi_p B^p)(1 - B)^d} a_t \tag{1}$$

An important issue is the identification procedure to determine appropriate values of p, d and q for a given data set. Typical identification tools, such as simple and partial autocorrelation functions, can generally be explored manually, but the automatic search for the optimal model with the help of information criteria, such as the Akaike Information Criterion or the Bayesian Information Criterion, among others, is a standard nowadays. Two popular approaches to automatic identification are [29,30], which are implemented in well-known software pieces, TRAMO and the forecast package in R, respectively.

2.3. Exponential smoothing

The fundamental principle of exponential smoothing is that forecasts are constructed as a correction of the last observation by a weighted average of past forecast errors that decay in an exponential pattern [30]. This particular pattern reflects the fact that more recent information is more valuable for forecasting the immediate future than older data.

ETS methods were initially used as ad-hoc forecasting tools in the 1960s [31], although they have recently been cast into a much more formal statistical framework, with a procedure for automatic selection of models included, as well as a wider range of models available. The widespread package in R known as forecast is used in the case study below and has become a standard when it comes to ETS methods [30].

As an example, equation (2) shows a common possibility for non-seasonal time series with an additive level and a slope.

$$\begin{aligned} &\text{Level}: && T_t = \alpha z_t + (1-\alpha)(T_{t-1} + b_{t-1}) \\ &\text{Slope}: && b_t = \beta(T_t - T_{t-1}) + (1-\beta) b_{t-1} \\ &\text{Forecast}: && \widehat{z}_{t+1} = T_t + b_t \end{aligned} \tag{2}$$

2.4. Unobserved Components

Although UC models are often considered with seasonal components, they can also be specified without one, as would be the case with the uranium price; see equation (3) and [32–35].

$$\begin{aligned} z_t &= T_t + I_t \\ \begin{bmatrix} T_{t+1} \\ T^*_{t+1} \end{bmatrix} &= \begin{bmatrix} 1 & 1 \\ 0 & \alpha \end{bmatrix} \begin{bmatrix} T_t \\ T^*_t \end{bmatrix} + \begin{bmatrix} \eta_t \\ \eta^*_t \end{bmatrix} \end{aligned} \tag{3}$$

T_t and I_t stand for the trend and irregular components, respectively; T^*_t is referred to as the trend 'slope', $0 < \alpha \leq 1$, η_t and η^*_t are independent white noise sequences with variances σ^2_η and $\sigma^2_{\eta^*}$, respectively. Playing with constraints on the parameters different trend types emerge, like a RW ($\alpha = 0; \sigma^2_{\eta^*} = 0; T^*_1 = 0$); RW with drift ($\alpha = 1; \sigma^2_{\eta^*} = 0; T^*_1 \neq 0$); local linear trend (LLT, $\alpha = 1$); smoothed trend ($\alpha = 1; \sigma_{\eta^2} = 0$) and damped trend (DT, $0 < \alpha < 1$). The common assumption about the irregular component is to consider it a Gaussian random variable.

In this paper, the selection of the best combination of trend and irregular is done by minimizing the Akaike Information Criterion,

shown in equation (4), where L^* is the likelihood value at the optimum and k the number of parameters in the model, i.e., all variances involved, α in equation (3) and each non-stationary state in the model [35].

$$AIC = -2ln(L^*) + 2k \tag{4}$$

The key to success is to choose a sufficiently large set of UC models to look for, so that the approach can model as many different time series as possible. The set of models to look for are 7, and are all possible combinations of trends (none, RW, LLT, DT) and irregular (none or random Gaussian noise) with the exception of the 'no trend' and 'no irregular' model.

Note that some of the models above are not considered, strictly speaking, UC models in the literature. Two examples: i) a RW is identified with an RW trend with no irregular; and ii) an LLT with no irregular and $\sigma^2_{\eta^*} = 0$ and $T^*_1 \neq 0$ is a RW with drift. In this way, the set of models considered is somewhat broader than the UC models.

As far as the author is concerned, this is the first time that such a procedure has been used for UC models.

2.5. Theta method

Theta is an ad-hoc method that was the winner of the M3 forecasting competition [17,36]. It is then an additional benchmark against which any other method can be compared. Although this method showed great robustness when applied to many time series, it does not mean that it is the best when applied to a single time series.

The model exploits the idea of modifying the local curvatures of a time series. This is achieved by applying a coefficient (called θ) to the second differences in the time series. The series deflates as *theta* approaches zero. In fact, for $\theta = 0$ the output is a linear trend and for $\theta = 1$ the output is the same series. However, these are not necessarily extreme values, since in reality the complete method is basically based on the combination of the theta lines forecasts with $\theta = 0$ and $\theta = 2$ [36].

2.6. Artificial neural networks

Neural networks are models that are constructed as a composition of layers of interconnected units (neurons) that allow fairly flexible ways of specifying models in many areas of science. They have been used extensively in prediction over the last thirty years with variable success. The results were not conclusive yet, although a new wave of neural networks is arriving hand in hand with deep and recurrent neural networks (mainly Long Short Term Memory nets). However, it is too early to see if they are really useful for forecasting and, in addition, deep learning techniques make sense when it is necessary to simultaneously forecast many time series, which is not the case for this paper.

Two architectures have been tested in the data, namely Multilayer Perceptrons (MLP) and Extreme Learning Machines (ELM). Both have a single hidden layer and the number of neurons is automatically estimated avoiding overfitting by following the algorithms in Refs. [37,38].

2.7. Combination of methods

There is an extensive literature reporting that the combination of forecasting methods significantly improves accuracy, see for example [17,37]. Very often simple combinations are reported as the most powerful, such as mean and median, see Ref. [39] and references therein. These possibilities will be studied in the case of the uranium price.

3. Results and discussion

In order to compare the performance of all the forecasting methods listed in the previous section, an experiment consisting on repeated forecasts has been launched with an initial forecasting origin set in December 2012 used to fit the designated methods. For each method, forecasts are produced one year ahead and the forecasting origin is then moved forward one month. Forecasts are then made again for all methods. This process is repeated until the end of the sample is reached. In total, 43 rounds of 12-month forecasts were carried out.

This comprehensive assessment is completed with the help of two error metrics, namely the symmetrical absolute mean percentage error (sMAPE) and the scaled absolute mean error (MASE), see equations (5) and(6) and [17,40].

$$\text{sMAPE}_h = h^{-1} \sum_{i=1}^{h} \frac{2|z_{T+i} - \widehat{z}_{T+i}|}{|z_{T+i}| + |\widehat{z}_{T+i}|} \times 100 \tag{5}$$

$$\text{MASE}_h = h^{-1} \sum_{i=1}^{h} \frac{|z_{T+i} - \widehat{z}_{T+i}|}{(n-1)^{-1} \sum_{r=2}^{n} |z_r - z_{r-1}|} \tag{6}$$

In equations (5) and (6) z_t and $\widehat{z}_t$ are the actual and forecast values at the time t, respectively; T is the forecasting origin; h is the forecasting horizon (from 1 to 12 in this paper); and n is the number of observations in the fitting sample. Both metrics have become the most important in recent forecasting literature, because they have proven to be very useful in many applications and are free of some drawbacks [17,40]. The sMAPE metric avoids distortions of the standard criterion of non-symmetric MAPE and problems for values close to zero. The MASE metric compares the out-of-sample performance of the model with the in-sample performance of a simple RW (actually our Naïve method).

Another issue taken into account in this work is whether a Box-Cox transformation adds any value in terms of prediction due to its power to stabilize the variance of the series [41]. The transformation is based on the selection of a λ parameter, with the limit case of $\lambda = 0$ being reduced to a natural logarithmic transformation according to equation (7). Although very often the logarithm transformation is chosen arbitrarily ($\lambda = 0$), rigorous methods are available, such as maximum likelihood or minimization of the coefficient of variation of the transformed series [41,42]. The latter is used in this document because it is model-independent and can be properly estimated prior to model identification.

$$z(\lambda) = \begin{cases} \dfrac{z^{\lambda} - 1}{\lambda}, & \text{if } \lambda \neq 0 \\ \log(z), & \text{if } \lambda = 0 \end{cases} \tag{7}$$

Several packages have been used to calculate model forecasts: R is used as the base platform for all calculations; ARIMA, ETS and Theta are estimated by the functions auto. arima, ets and theta, respectively, of the package forecast [30]; ANNs are estimated with the nnfor package [38]. UC models deserve special mention, as this is the first time that an automatic model selection procedure based on information criteria has been used in a UC context. This software is developed in C++ mounted in R with the help of RcppArmadillo [43].

Ratios of mean sMAPE errors across all forecasting origins considered for each method in the previous section with respect to mean sMAPE of the Naïve method are reported in Table 2 and

Table 2
Ratios of mean sMAPE measure across all forecast origins for each method over the Naïve mean sMAPE. Bold letters indicate the best model in each row.

Without Box-Cox transformation								
h	ARIMA	AR(1)	ETS	UC	Theta	MLP	ELM	COMB
1	1.048	0.989	0.995	**0.943**	1.086	1.105	1.035	0.952
2	1.077	0.989	1.019	0.963	1.105	1.145	1.053	**0.961**
3	1.066	0.985	1.026	**0.964**	1.102	1.142	1.061	0.967
4	1.049	0.981	1.016	0.960	1.094	1.135	1.072	**0.952**
5	1.041	0.978	1.005	0.962	1.086	1.140	1.083	**0.953**
6	1.039	0.976	0.999	0.963	1.079	1.132	1.093	**0.953**
7	1.036	0.974	0.996	0.964	1.077	1.134	1.104	**0.953**
8	1.031	0.973	0.998	0.964	1.077	1.142	1.111	**0.949**
9	1.028	0.972	1.002	0.964	1.077	1.158	1.115	**0.947**
10	1.027	0.971	1.016	0.964	1.075	1.166	1.118	**0.948**
11	1.024	0.970	1.035	0.964	1.073	1.177	1.121	**0.955**
12	1.021	0.969	1.056	**0.964**	1.071	1.180	1.121	0.967
With Box-Cox transformation								
h	ARIMA	AR(1)	ETS	UC	Theta	MLP	ELM	COMB
1	0.976	0.990	1.000	**0.918**	1.120	1.049	1.055	0.951
2	0.996	0.984	1.029	**0.953**	1.131	1.128	1.097	0.982
3	0.997	0.983	1.043	**0.954**	1.149	1.171	1.131	0.988
4	0.987	0.978	1.036	**0.940**	1.146	1.198	1.157	0.976
5	0.978	0.975	1.029	**0.933**	1.140	1.206	1.180	0.968
6	0.976	0.973	1.014	**0.931**	1.136	1.206	1.201	0.961
7	0.972	0.969	1.009	**0.930**	1.139	1.213	1.226	0.960
8	0.968	0.967	1.004	**0.937**	1.144	1.232	1.246	0.961
9	0.963	0.967	0.996	**0.949**	1.146	1.250	1.260	0.962
10	0.963	0.967	0.991	**0.955**	1.148	1.280	1.274	0.961
11	0.971	0.966	0.988	0.960	1.150	1.314	1.286	**0.959**
12	0.982	0.966	0.986	0.959	1.148	1.349	1.291	**0.957**

shown graphically in Fig. 2 for forecasting horizons ranging from 1 to 12 months. A column is added for an AR(1) with intercept, which is the best version of all ARIMA models included in Ref. [8]. In addition, several combinations of forecasts were tested, and the table presents only the best of them all, namely the mean of ETS and UC (under the heading COMB). Values below 1 in Table 2 indicate that a particular method is more accurate than the Naïve (the smaller the better), while values above 1 imply that the Naïve is better. Equivalent MASE results are not shown because they were mostly consistent with sMAPE and are available from the author upon request. Bold letters indicate the best model for each forecast horizon (in rows).

Some relevant observations follow:

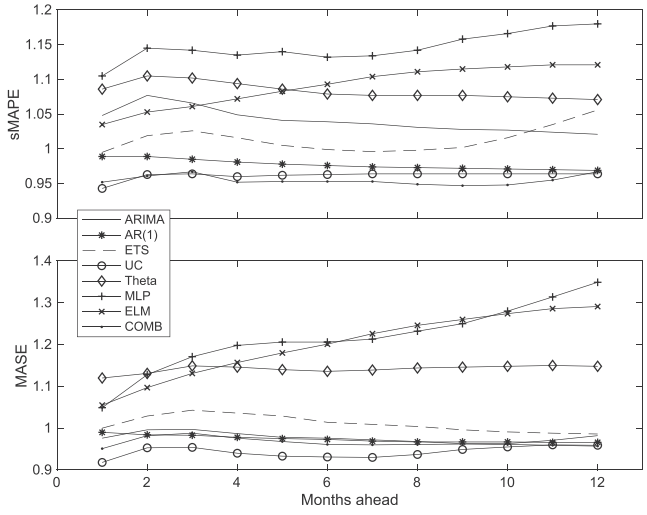

Fig. 2. Ratio of mean error metrics of all forecasting methods to Naïve. sMAPE on top panel and MASE on bottom panel.

- There are always better methods than Naïve, which implies that there is some structure in the data that can be modeled by some of the estimated methods. However, all values are above 0.9 and many near 1 and above 1. This is clear proof that uranium prices are quite difficult to predict.
- The *lambda* parameter of the Box-Cox transformation varies between −0.53 and −0.3, indicating that the default value of 0 that is often used in practice is inappropriate for uranium prices. Transformation is useful to induce greater precision in those methods that are better (with the exception of COMB). However, in the worst performing methods the Box-Cox transformation does not add any value (Theta, MLP, ELM and ETS).
- The best option in general is the UC using the Box-Cox transformation, which produced improvements between 4.5 and 8.2% over the Naïve, depending on the forecasting horizon. This is the model that would have produced the best results if it had been used over time in the experiment.
- The second best is COMB with Box-Cox transformation, although it is the first when the Box-Cox transformation is not used. In fact, Table 2 reflects the ambiguity of the forecast combinations sometimes reported in the literature [39], since in the upper part of the table the COMB option is often better than the two models individually (UC and ETS), while in the second part, one of the methods used in the combination is the best (UC).
- The results of the ETS method are not good at all. The ETS is better than the Naïve only in a few cases and by a small margin. The same can be said of ANNs (both MLP and ELM), poor results can be attributed to the shortness of time series, as it is well known that artificial intelligence methods are data hungry.
- ARIMA methods ranks in the middle, since it turned better than Naïve, ETS, Theta, MLP and ELM methods when Box-Cox transformation is taken into account.
- The AR(1) with intercept proposed in Ref. [8] for quarterly data is better than ARIMA for automatically identified monthly data

[30]. This may imply some problems in the ARIMA identification procedure. This AR(1) model becomes a RW with drift (i.e., $z_t = z_{t-1} + a_t + c$), because the AR parameter always becomes smaller than − 0.99, implying that a unit root is dominant.

- Theta method is especially bad. These results contrast with the excellent results that this method achieved in contexts where many time series have to be forecast [17]. This result may seem surprising at first, because in the end the Theta method is theoretically similar to some of the other methods considered in this paper. In fact, it can be seen as a moving average of exponential weighting with constant or as an RW with drift and additional noise. But the problem is that the Theta method is actually a combination of two 'theta lines' with predefined fixed parameters θ (0, and 2). This rigidity is clearly not appropriate for the series in this paper, although it might be suitable for other series or in contexts with many time series [17].
- Fig. 2 adds some further insights about the observations made so far, mainly because it visually emphasizes the different magnitude of the errors. In general, both error metrics produce the same classification of methods with the exception of the ARIMA model, which is penalized by the sMAPE. Depending on the accuracy of the forecasts, two separate groups of methods can be conceived, namely ANN with Theta on the one hand and the rest on the other.

4. Conclusions

The aim of this research is to find the best strategy for predicting uranium prices, based on an experiment that compares nine forecasting methods and takes into account the stabilizing transformation of the Box-Cox variance. The conclusions are drawn from an exhaustive experiment in which the origin of the forecasts is systematically changed and all models are rigorously tested. The methods used include some benchmarks, statistical methods and artificial intelligence methods.

Not all methods outperform the simplest forecasts obtained with a RW (often called Naïve). As a matter of fact, only ARIMA, UC and the average of UC and ETS perform better. The best general method is UC and this should be the method to use in case the forecasting error wanted to be minimized.

These conclusions apply locally to this time series of uranium prices. In no way can a general conclusion be drawn about the performance of the forecasts of any of the methods involved in other contexts or in other time series. In fact, some of the worst methods worked very well in forecasting competitions, where many different time series have to be forecast.

All the methods considered are univariate in nature, which limits their potential for predicting turning points, such as the drastic one observed in 2007, or for incorporating the effects of uranium unconventional recovery techniques, such as recovery from seawater or phosphate rock. These effects could be implemented with extensions of the previous models to include input drivers, provided that such useful drivers exist and sufficiently accurate models can be identified from the data.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the Vicerrectorado de Investigación y Política Científica from UCLM through the research group fund program (PREDILAB; DOCM 24/01/2019 [2019/536]) and by the European Regional Development Fund. The author would also like to thank two anonymous referees and the editor for their valuable comments.

References

[1] D. Cole, The Global Uranium Market, Reserve Bank of Australia, 2015. Reserve Bank of Australia.
[2] D. Kryzia, L. Gawlik, Forecasting the price of uranium based on the costs of uranium deposits exploitation, Miner. Resour. Manag. 32 (2016) 93–110.
[3] N. Haneklaus, Y. Sun, B. Roland, B. Lottermoser, E. Schnug, To extract, or not to extract uranium from phosphate rock, that is the question, Environ. Sci. Technol. 51 (2) (2017) 753–754. URL 10.1021/acs.est.6b05506.
[4] B. Parker, Z. Zhang, L. Rao, J. Arnold, An overview and recent progress in the chemistry of uranium extraction from seawater, Dalton Trans. 3 (2018) 1–6. URL 10.1039/c7dt04058j.
[5] w. Mays, Limitations to progress in developing uranium resources, in: Proceedings of the 30th Internatio -nal Symposium on Uranium and Nuclear Energy, World Nuclear Association, 2005, pp. 39–60.
[6] I.A.E. Agency, Uram-2018: ebb and flow — the economics of uranium mining, accessed: 2019-06-20, https://www.iaea.org/newscenter/news/uram-2018-ebb-and-flow-the-economics-of-uranium-mining.
[7] International Monetary Fund, Primary commodity prices, accessed: 2019-06-20, https://www.imf.org/en/Research/commodity-prices.
[8] S. Kim, W. Ko, H. Nam, C. Kim, Y. Chung, S. Bang, Statistical model for forecasting uranium prices to estimate the nuclear fuel cycle cost, Nucl. Eng. Technol. 49 (5) (2017) 1063–1070.
[9] Q. Yan, S. Wang, B. Li, Forecasting uranium resource price prediction by extreme learning machine with empirical mode decomposition and phase space reconstruction, Discrete Dynam Nat. Soc. (2014) 1–10, 2014.
[10] J. Chen1, Y. Zhao, Q. Song, Z. Zhou1, S. Yang, Exploration and mining evaluation wywtem and price prediction of uranium resources, Min. Miner. Depos. 12 (2018) 85–94.
[11] J. Chen, Y. Zhao, Q. Song, Z. Zhou, S. Yang, Exploration and mining evaluation system and price prediction of uranium resources, Min. Miner. Depos. 12 (2018) 85–94.
[12] S. Kahouli, Re-examining uranium supply and demand: new insights, Energy Policy 39 (2011) 358–376.
[13] M. Dittmar, The end of cheap uranium, Sci. Total Environ. 461–462 (2013) 792–798.
[14] A. Monnet, S. Gabriel, J. Percebois, Long-term availability of global uranium resources, Resour. Policy 53 (2017) 394–407.
[15] V. Gómez, A. Maravall, Automatic modeling methods for univariate series, in: A Course in Time Series, John Wiley & Sons, Inc., 2001, pp. 171–201.
[16] R. Hyndman, A.B. Koehler, J.K. Ord, R.D. Snyder, Forecasting with Exponential Smoothing: the State Space Approach, Springer Science & Business Media, 2008.
[17] S. Makridakis, M. Hibon, The m3-competition: results, conclusions and implications, Int. J. Forecast. 16 (4) (2000) 451–476.
[18] M.A. Villegas, D.J. Pedregal, Automatic selection of unobserved components models for supply chain forecasting, Int. J. Forecast. 35 (1) (2019) 157–169 (special Section: Supply Chain Forecasting).
[19] K.S. Chan, H. Tong, On estimating thresholds in autoregressive models, J. Time Ser. Anal. 7 (3) (1986) 179–190.
[20] D. van Dijk, T. Teräsvirta, P.H. Franses, Smooth transition autoregressive models - a survey of recent developments, Econom. Rev. 21 (1) (2002) 1–47.
[21] G. Sbrana, A. Silvestrini, Random switching exponential smoothing: a new estimation approach, Int. J. Prod. Econ. 211 (2019) 211–220.
[22] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer Publishing, New York, 2009.
[23] S. Haykin, Neural Networks and Learning Machines, Prentice Hall, New Jersey, 2008.
[24] G. Zhang, Times series forecasting using a hybrid arima and neural network model, Neurocomputing 50 (2003) 159–175.
[25] H. Zheng, J. Yuan, L. Chen, Short-term load forecasting using emd-lstm neural networks with a xgboost algorithm for feature importance evaluation, Energies 10 (2017) 1–20.
[26] S. Makridakis, E. Spiliotis, V. Assimakopoulos, Statistical and machine learning forecasting methods: concerns and ways forward, PLoS One 13 (3) (2018) 1–26, https://doi.org/10.1371/journal.pone.0194889. URL 10.1371/journal.pone.0194889.
[27] S. Makridakis, E. Spiliotis, V. Assimakopoulos, The m4 competition: results, findings, conclusion and way forward, Int. J. Forecast. 34 (2018) 802–808.
[28] G.E.P. Box, G.M. Jenkins, G.C. Reinsel, G.M. Ljung, Time Series Analysis: Forecasting and Control, fifth ed., John Wiley & Sons, 2015.
[29] V. Gómez, A. Maravall, Automatic modeling methods for univariate series, in: A Course in Time Series, John Wiley & Sons, Inc., 2001, pp. 171–201.
[30] R.J. Hyndman, Y. Khandakar, Automatic time series forecasting: the forecast package for R, J. Stat. Softw. 3 (27) (2008) 1–22.
[31] R.G. Brown, Statistical Forecasting for Inventory Control, McGraw-Hill, 1959.

[32] A.C. Harvey, Forecasting, Structural Time Series Models and the Kalman Filter, Cambridge university press, 1989.
[33] P.C. Young, D.J. Pedregal, W. Tych, Dynamic harmonic regression, J. Forecast. 18 (6) (1999) 369–394.
[34] C.J. Taylor, D.J. Pedregal, P.C. Young, W. Tych, Environmental time series analysis and forecasting with the captain toolbox, Environ. Model. Softw 22 (6) (2007) 797–814.
[35] J. Durbin, S.J. Koopman, Time Series Analysis by State Space Methods, second ed., Oxford University Press, 2012.
[36] V. Assimakopoulos, K. Nikolopoulos, The theta model: a decomposition approach to forecasting, Int. J. Forecast. 16 (4) (2000) 521–530.
[37] N. Kourentzes, D. Barrow, S. Crone, Neural network ensemble operators for time series forecasting, Expert Syst. Appl. 41 (9) (2014) 4235–4244.
[38] N. Kourentzes, Nnfor: time series forecasting with neural networks, r package version 0.9.6, URL, https://CRAN.R-project.org/package=nnfor, 2019.
[39] D.K. Barrow, N. Kourentzes, Distributions of forecasting errors of forecast combinations: implications for inventory management, Int. J. Prod. Econ. 177 (2016) 24–33.
[40] R. Hyndman, A. Koehler, Another look at measures of forecast accuracy, Int. J. Forecast. 22 (4) (2006) 679–688.
[41] G.E.P. Box, D.R. Cox, An analysis of transformations, J. R. Stat. Soc. Ser. B 26 (2) (1964) 211–252.
[42] V.M. Guerrero, Time series analysis supported by power transformations, J. Forecast. 12 (1) (1993) 37–48.
[43] D. Eddelbuettel, C. Sanderson, Rcpparmadillo: accelerating r with high-performance c++ linear algebra, Comput. Stat. Data Anal. 71 (2014) 1054–1063, https://doi.org/10.1016/j.csda.2013.02.005. URL.