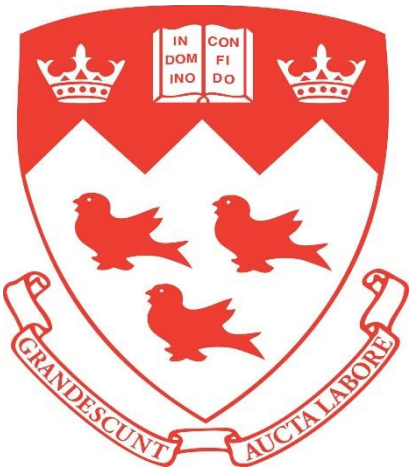

Mineral Commodity Price Forecasting through Time-Series Modeling Techniques and Artificial Neural Networks: A Nickel Case Study

BRIAN ANDONAIRE VILLEGAS

Department of Mining and Materials Engineering

McGill University, Montreal, Quebec

July 2021

A thesis submitted to McGill University in partial fulfillment of the

requirements for the degree of **Masters of Science**

Abstract

The forecasting of mineral commodity prices, yet perhaps the most important factor to the decision-making process of mining companies, still remains an intricate task due to the different elements that drive such prices. To this date, several techniques have been developed with the aim of improving the forecasting accuracy; however, their architecture has also varied in terms of complexity.

This thesis explores the use of Artificial Neural Networks (ANNs) on the forecasting of mineral commodity prices. ANNs have proven to be quite helpful modeling the hidden patterns that many time-series forecasting methods, be they the Autoregressive Integrated Moving Average (ARIMA) and the Generalized Autoregressive Conditionally Heteroscedastic (GARCH) have failed to.

A case study was conducted not only to assess the accuracy of the ANNs on nickel price forecasting but also to determine how this model compares to the ARIMA and the GARCH techniques. This study reveals that, indeed, the forecasted values based on ANNs more accurately captures the real values than those obtained by the other two methods. The same study also shows that, although these three methods are somewhat good at estimating future price values, their accuracy works better for short terms, i.e., from a couple of months to up to two years (or any time unit of interest).

Résume

La prévision des prix des matières premières minérales, pourtant peut-être le facteur le plus important du processus décisionnel des sociétés minières, reste une tâche complexe en raison des différents éléments qui déterminent ces prix. A ce jour, plusieurs techniques ont été développées dans le but d'améliorer la précision des prévisions; cependant, leur architecture a également varié en termes de complexité.

Cette thèse explore l'utilisation des Artificial Neural Networks (ANNs) pour la prévision des prix des matières premières minérales. Les ANNs se sont révélés très utiles pour modéliser les modèles cachés que de nombreuses méthodes de prévision de séries chronologiques, qu'il s'agisse de la AutoRegressive Integrated Moving Average (ARIMA) et de Generalized Autoregressive Conditionally Heteroscedastic (GARCH), ont échoué.

Une étude de cas a été menée non seulement pour évaluer l'exactitude des ANNs sur la prévision du prix du nickel, mais aussi pour déterminer comment ce modèle se compare aux techniques de ARIMA et GARCH. Cette étude révèle que, en effet, les valeurs prévues basées sur les ANNs ressemblent plus exactement aux valeurs réelles que celles obtenues par les deux autres méthodes. La même étude montre également que, bien que ces trois méthodes soient assez bonnes pour estimer les valeurs de prix futures, leur précision fonctionne mieux à court terme, c'est-à-dire de quelques mois à deux ans (ou toute unité de temps d'intérêt).

Acknowledgements

I would like to thank my parents who have always given me their full support, especially as I embarked on this new chapter of my life in a new country.

To the Peruvian Institute of Mining Engineers (IIMP) who are the responsible for making this whole dream of studying in a cutting-edge country come true.

I would also like to thank Professor Kumral for his unstinting support over these two years not only as a supervisor but also as friend.

Content

List of Figures

List of Tables

List of Abbreviations

ACF	Autocorrelation Function
ADF	Augmented Dickey-Fuller test
AIC	Akaike Information Criterion
ANNs	Artificial Neural Networks
AR	AutoRegressive
ARCH	AutoRegressive Conditional Heteroskedasticity
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
CBOT	Chicago Board of Trade
CEI	Chemical Engineering Index
COMEX	Commodity Exchange, Inc.
CPI	Consumer Price Index
CT	Chaos Theory
ETF	Exchange Traded Funds
GARCH	Generalized Autoregressive Conditionally Heteroscedastic
GDP	Gross Domestic Product
LBMA	London Bullion Market Association
LCPI	Logarithmic CPI
LME	London Metal Exchange
LPPI	Logarithmic PPI
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
M&S	Marshall & Swift Index for industrial equipment
MUV	Manufactures Unit Value Index
NR	Nelson Refinery Construction Cost Index
NYMEX	New York Mercantile Exchange
PACF	Partial Autocorrelation Function
PPI	Producer Price Index
R^2	Coefficient of Determination
RMSE	Root Mean Squared Error
SS_E	Sum of Squares
SS_{yy}	Sum of Residuals
WPI	Wholesale Price Index

Chapter 1

Introduction

1.1. Problem statement

Mineral commodity prices are one of the most crucial parameters used in the economic evaluation of a mining project. They have a considerable impact on the company's revenues, heavily influencing the strategic decisions to be made by the company's board. Moreover, since mining companies, in the majority of cases, play almost no role in the determination of the mineral prices, i.e., they are price takers, their plans are deeply dependant on mineral commodity prices.

Regardless of this dependence on mineral commodity price forecasts and the different techniques that have tried to address this issue, mineral price forecasting remains challenging due to the increased volatility in commodity markets. Globalization increased the sensitivities to economic, commercial, natural and political events at every point of the world. The markets of many commodities are international. A great deal of effort has been dedicated to developing and improving the forecasting methods used for the estimation of future mineral prices.

The main two drivers of mineral commodity prices are supply and demand. In fact, many conventional forecasting methods take advantage of historical supply and demand values and, through the use of mathematical techniques, attempt to obtain estimates of mineral prices in the short and long term (Espí Rodríguez and De La Torre, 2013).

The forecasting approach used by mining companies and stockbroker analysts are based on time-series techniques. These methods may be classified as:

a. Simple methods such as exponential smoothing and moving average (MA). In fact, exponential smoothing and MA are smoothing techniques rather than forecasting.
b. More sophisticated techniques such as the Autoregressive Moving Average (ARMA) and Autoregressive Integrated Moving Average (ARIMA).

In addition to these techniques, simulation-based approaches (e.g., geometric Brownian motion and mean reverting techniques) and various machine learning techniques are also used for forecasting purposes.

Time series-based methods assume that past behaviour of a commodity price will continue into the future; however, their main pitfall is that all follow linear models, constraining the future values to be linear functions of past observations (Kriechbaumer et al., 2014). This issue poses a serious problem given that real-time series exhibit a mixture of linear and non-linear patterns.

In order to better model the patterns real-world time series show, several techniques have been proposed, some of which have proven to be more accurate than others at predicting future values. Two of the most widely used non-linear methods are the AutoRegressive Conditional Heteroskedasticity (ARCH) and the Generalized Autoregressive Conditionally Heteroscedastic (GARCH). Theoretically, when the behaviour of the series over time is the only factor taken into account, both generally represent an improvement from the conventional and still widely used ARMA and ARIMA models. Nonetheless, they are no without any drawbacks, as expected from any forecasting technique.

Researchers in the matter of price forecasting have shown that, even when ARIMA and GARCH are the go-to methods for mining companies and traders due to their simplicity, they tend to fall short to machine learning techniques such as Artificial Neural Networks (ANNs) (Chen et al., 2003).

ANNs have shown to more accurately model real-world time series, providing more accurate forecasts (Kristjanpoller and Minutolo, 2015). However, growing research on the factors that affect the capability of forecasting methods, even the most sophisticated one, has yielded interesting results as to what could possibly make one forecasting technique more accurate than the other at predicting future values (Dooley and Lenihan, 2005). This phenomenon will be discussed in later chapters.

In order to properly grasp the accuracy and usefulness of such forecasting methods, a case study needs to be conducted. The selection of the mineral commodity boasted different metals; however, given its critical nature as a raw material and its importance not only as a macroeconomic indicator but also for the industrial sector, *nickel* was chosen.

1.2. Research objectives

- Examine the appropriateness of the most common linear time-series method, i.e., the ARIMA technique.
- Determine the efficiency of the most common non-linear model for the data's volatility: The GARCH model. To that end, a hybrid ARIMA-GARCH model will be used, as it is not possible to forecast the mean values of the data using only the GARCH model.
- Investigate the performance of machine learning method, using ANNs on the forecast of mineral commodity prices. It will be verified the flexibility and non-linear capability of neural networks over the traditional linear and non-linear time-series model.
- Benchmark the outcomes of different methods.

1.3. Original contributions

As discussed in Section 1.1, the main shortcoming of the conventional time-series methods is their linearity assumption behind their models (smoothing techniques, ARIMA). To overcome this particular issue, different models, those based on pure non-linear time series and those based on machine learning techniques, are explored. In fact, machine learning methods such as ANNs have proven to effectively identify the hidden non-linear patterns that the pure non-linear time series models failed to do (Lorente-Leyva et al., 2019), outperforming the conventional time-series forecasting techniques. This research aims to demonstrate the pros and cons of the time-series forecasting techniques and extend more advanced ANNs technique. All techniques to be investigated will be tested through a case study.

1.4.Social and economic impacts

The forecast of mineral commodity prices has a great positive impact on all the stakeholders (i.e., mining corporates, governments, and local communities) of the mining industry. More precisely, the outcome of this research will most likely be beneficial to:

- *Tax makers*, by providing a broader understanding of the future behavior of the mineral commodity prices (Takatoshi and Rose, 2011). Thus, governments can set the most efficient yet reasonable tax rate to maximize benefit for society. Likewise, local communities would have a much clearer idea of the financial/social perks mining project operating in the nearby area would bring so the process to grant social license to operate would be less cumbersome and social risks could be reduced. Finally, mining corporates will have more reliable information on the profitability and sustainability of the operations,

- *Policymakers*, by effectively monitoring mineral prices, which are one of the major factors affecting policy goals (Rudenno, 2012). Consequently, as part of the mining policy, an ideal type of tax regime can be defined so as to balance between the benefit of society and sustainability of mining corporates. That is, the predicted price values can help governments to adopt regulatory actions that ensure a fair taxation system in all the possible price scenarios, allowing neither corruption nor exaggerated demand by the communities when prices rise. Furthermore, the satisfactory forecast of mineral prices facilitates the delineation of governmental mining policies so as to guarantee sufficient future mineral supplies (Radetzki and Wårell, 2016).
- *Investors*, by facilitating their decision-making process regarding their investments. In fact, every mining project is subject to numerous risks associated with the mineral commodity markets, and, therefore, a better and improved prediction of the commodities' prices can lead to better evaluation and optimization of investment opportunities in mining projects.
- *Corporations*, by reducing the level of uncertainty linked to the mineral commodity price. By doing so, the corporations will be able to make an informed decision regarding major strategic decisions: mine portfolio management, the construction of a new mine, or expansion of production (Radetzki and Wårell, 2016). Furthermore, mining companies will also be better equipped to elaborate on a more representative contingency plan that addresses the possible changes in the market conditions.

  A more realistic forecast of the mineral commodity price can also provide a greater understanding of the financial value of a mining company (Schaeffer, 2008). That is, since the mineral commodity price is one of the main drivers in the mineral reserves calculation, which, in turn, are the principal parameter that underpins the liquidity of a mining company, a more accurate price forecasting can help a mining company understand its financial status in the short and long run.

- *Mine planners*, by helping them set more realistic and accurate capacities, developing cut-off grade strategy, and adjusting the corresponding mining method sequence (Rudenno, 2012). Consequently, mine planners are able to better adjust their mine plans, reflected as changes in the production sequencing, the life of mine, the mineral reserves, etc. so as to maximize the mine's profits in both the short, medium, and long term.

1.5. Outline

Chapter 1 states the research problem and identifies both the original contributions to the field and the research's social and economic impacts.

Chapter 2 reviews key statistical concepts and describes the existing methodologies and research in the mineral commodity price forecasting field.

Chapter 3 explains the methodology and the way it is organized,

Chapter 4 analyzes the conventional linear time-series method used in the mining industry: The ARIMA technique,

Chapter 5 examines the currently used non-linear time-series method, GARCH, as well as its joint performance ARIMA-GARCH on the mineral commodity price forecasting.

Chapter 6 evaluates the use of the machine learning technique ANNs in the mineral price forecasting.

Chapter 7 summarizes the key findings/findings of this research and outlines the future work.

Chapter 2

Literature Review

2.1. Commodity markets

The economic sectors can be categorized into the primary, secondary, and tertiary sectors. In this regard, only the primary sector includes raw materials, i.e., metals, agriculture, fishery, and forestry. The outcomes of primary sectors are generally called commodities, which are typically traded in *commodity exchanges* (Radetzki and Wårell, 2016).

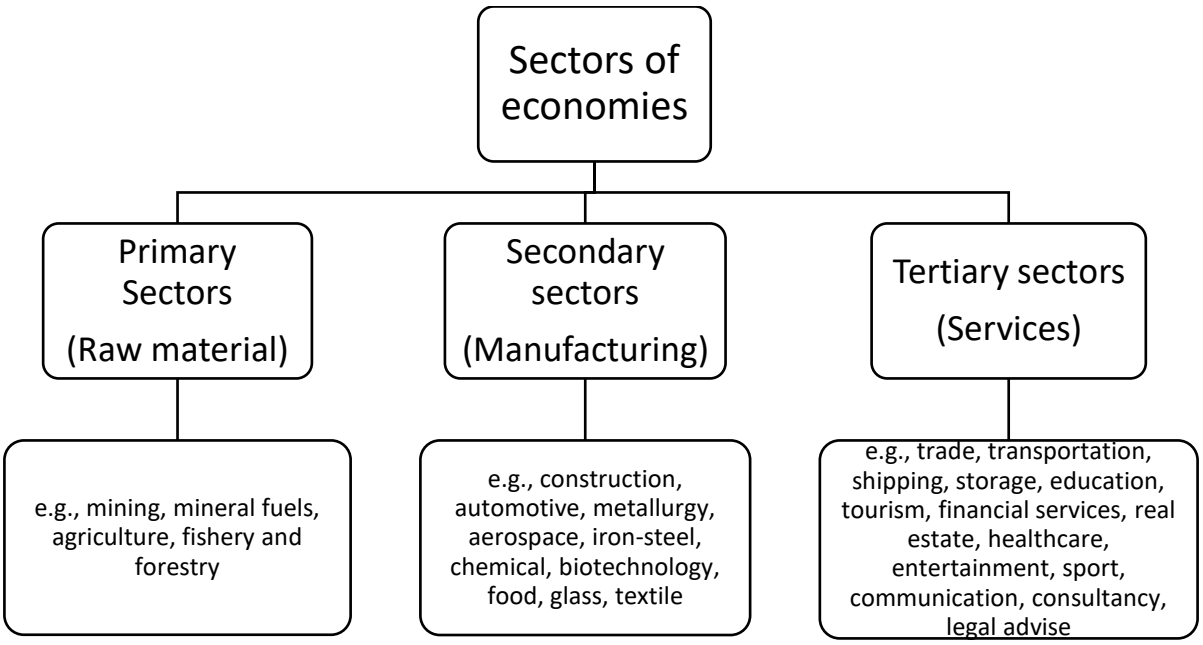

Figure 2.1 - Sectors of economies

(Source: "A Handbook of Primary Commodities in the Global Economy", Chapter 5, Radetzki et al., 2016)

A commodity exchange (also called commodity market) is the market where primary sector products are traded (Rudenno, 2012). Examples of such commodity exchanges are the London Bullion Market (LBMA), London Metal Exchange (LME), the CME Group (Chicago Board of Trade), the New York Mercantile Exchange (NYMEX), and Commodity Exchange, Inc. (COMEX) merged). These exchanges require high product standards (both quality and quantity) so as to guarantee a successful commodity trading between both parties (buyer and seller) (Priolon, 2018). The prices of the commodities are set on a daily basis on the exchanges and companies can buy or sell through the so-called brokers.

These commodity markets are usually driven by two fundamental economic factors: supply and demand (Radetzki and Wårell, 2016). These two elements determine the price that suppliers are inclined to take, and buyers are willing to pay. As a result, any major disruption or changes in both elements would shift the commodity prices.

In many of the commodity exchanges, different commodities are traded; the main tradable commodities can be found in Table 2.1.

Table 2.1 : Tradable commodities

Group	**Examples**
Precious Metals	Gold, Silver, Platinum
Base Metals	Copper, Lead, Molybdenum Nickel, Aluminum
Agro-Based Commodities	Wheat, Corn, Cotton, Oils, Oilseeds
Soft Commodities	Coffee, Cocoa, Sugar
Live-Stock	Live Cattle, Pork Bellies
Energy	Crude Oil, Natural Gas, Gasoline

(Source: “A Handbook of Primary Commodities in the Global Economy”, Chapter 5, Radetzki et al., 2016)

There are different types of markets at which these commodities are traded, such as spot or physical markets, futures markets, auction markets, and commodity markets. Also, given that mining corporates are listed in stock markets, stock markets and commodity price dynamics are closely intertwined (Radetzki and Wårell, 2016).

- The physical market (or the spot market) is a market where commodities are exchanged. That is, the seller supplies the amount of the commodity requested by the buyer, who takes delivery of it. In these markets, the transactional price is also called the spot price. Although, in reality, only a relatively small percentage of the commodity is actually delivered, since transactions are mainly made through future contracts, some exchanges do have storage facilities/warehouses (Rudenno, 2012).
- Auction markets: Sellers introduce the lowest price willing to accept and buyers bid the highest price willing to pay.
- Future markets, as opposed to the stop markets, are the markets where the buyer and seller of the mineral commodity "settle in advance" to a certain amount of the commodity at an agreed-upon price to be delivered at the pre-defined time in the future. Since the price has already been pre-defined and agreed upon by both parties, futures markets can provide a tool for estimating mineral commodity prices. One example of future markets is the NYMEX.
- Commodity markets, as per its name suggests, is the market at which primary sector products are traded.

In general, the exchanges discussed above can serve a number of purposes. One of those is the establishment of competitive and representative commodity prices. In fact, as competitive markets, these exchanges define the commodity prices by weighing in the transactions made by both buyers

and sellers (Takatoshi and Rose, 2011). Consequently, these commodity prices can be taken as benchmarks by outside buyers who wish to purchase commodities at entities outside the exchanges. Furthermore, these prices set by the exchanges can come in handy to not only perform an economic valuation of investment portfolios, but also establish commodity indices as well as exchange traded funds (ETF) (Jones et al., 2019). In addition to the benefits mentioned above, these commodity exchanges can also serve as delivery locations.

An example of commodity markets is the LME, which is one of the world trading centers for mineral commodities, especially non-ferrous metals such as copper and nickel[1], where future buyers and sellers of metal futures and options meet. These commodities can be traded in different forms: an agreed-upon quantity of a metal is bought/sold on a pre-defined date, at a fixed price agreed today, or via an option, in which futures contract are sold/bought at a price (also known as the strike price) agreed today (The London Metal Exchange, 2019). This exchange also offers prospective clients a tool to manage the risks against metal price movements through hedging. That is, by allowing its clients to hedge their commodity prices, the LME helps mine operators to reduce the potential losses that might incur as a result of falling.

This exchange also provides prospective sellers and buyers with a physical market of last resort, via the worldwide network of LME-approved warehouses (The London Metal Exchange, 2019). In fact, LME has approximately 700 warehousing facilities, serving as a spot market with more than 500 brands of metal products listed. That way, through the LME, consumers are able to purchase mineral commodities in times of extreme shortage, whereas producers are also able to sell their metals in times of oversupply, granted that they both meet the criteria set by the exchange.

[1]. The LME is the major world trading center for copper and lead.

Other exchanges of significant relevance are those that belong to the CME, a merger of the NYMEX, and the COMEX. The NYMEX is an exchange where futures and options are traded. The metals it trades are copper, gold, zinc, and silver. This exchange has set, as any other exchange, a set of rules governing the adequate trading, yet its rules are not as strict as those of the LME. In fact, the NYMEX trade is vastly confined to the American market and, taking advantage of the preference the American producers have to deal directly with the buyers.

LBM is the market at which gold and silver are traded. This trading, which is usually conducted by buyers and sellers directly, is overseen by the Bank of England and is represented by the LBMA. Since all transactions are conducted between both parties to maintain confidentiality, all risks exist only between the two counterparts. In this exchange, the members of the LBM, along with the Bank of England, act as gatekeepers to the market and certify the bullion traded and delivered meets the LBMA standards. Table 2.2 shows the metals and the exchanges they are traded at.

Table 2.2: Main metals and their principal exchanges

Group		**Exchanges**
Base Metals	Copper	LME, COMEX
	Lead	LME
	Zinc	LME
	Nickel	LME
	Molybdenum	LME
	Tin	LME
	Steel	NYMEX
	Aluminum	LME
Precious Metals	Gold	COMEX, LBMA
	Silver	COMEX

2.2. Commodity price mechanisms

In order to understand the intricate economic drivers of the mineral commodity price, it is useful as a starting point to consider the basic economic theory of supply and demand.

Mineral reserves this period = Mineral reserves last period - New mineral consumption last period + Amount of mineral recycled last period + New mineral reserves discovered last period

Figure 2.2- Definition of mineral supply

Theoretically, the price of a mineral commodity settles at the equilibrium point of both the supply and demand curve. As a result, considering a competitive market, this settlement of the price would prevent suppliers and consumers from charging higher price and offering way less money than determined, respectively. In this regard, commodity markets do not behave any different than any competitive market (Priolon, 2018). What really sets commodities markets apart is the group of factors that influences their supply and demand (Zsidisin, 2011).

2.2.1. Supply-driven factors

Mineral commodities can be extracted as main products, co-products, or by-products. The way these commodities are extracted directly affects the mine's output and viability (Radetzki and Wårell, 2016).

Once processed and consumed, some metals can be recovered and recycled. Recycled materials make up what is called "secondary" production and are often referred to as "scrap." Scrap material is usually classified into two categories: new scrap, which is generated in the manufacturing of new goods, and old scrap, which derives from goods that have reached the end

of their useful lives or have become obsolete. In this regard, according to Rudenno (2012), the flow of minerals supplied in a given period needs to consider the relationship between the depletion of a finite mineral stock, current mineral reserves, new mineral supply, recycled materials, and additions to stocks from new discoveries, technological change, or changes in prices.

Unexpected changes in the mineral commodity supply can be attributed to different factors including the opening or closure of mining operation(s), the significance of certain mining operation to the global production as stalled production in a highly production-contributive country due to a company's external or internal factor that can decline the overall world metal production of a certain commodity. In fact, disruptive events such as strikes, accidents in major mines, natural disasters such as floods and fires, failure of important machinery can be detrimental to mineral supply worldwide.

2.2.2 Demand-driven factors

The main driver for the demand of a commodity is associated with its industrial use. This use may change over time due to new technologies or the effects of substitutes or complementary products.

What determines the demand behavior for a mineral at any given point in time are the technological level, the quantity of consumers and their preferences, and the prices of both the commodities and any other substitute products (Takatoshi and Rose, 2011).

The competition between different producers is the determinant of the way goods are produced. In theory, the chosen method would be the one with the lowest cost at any time. That is, if at certain point in time a new method of production with a lower cost than the current one is found, this one will displace the old one.

When the supply of a given metal/mineral increases, the demand will also increase. Consequently, the price of the mineral commodity will rise until the overall production peaks what is wanted from buyers. Similarly, when there is an overproduction of a mineral commodity, buyers will take the lowest price (or last price quoted at an exchange) and this will continue until the excess of material is exhausted and the equilibrium is re-established. Production will decrease and because demand has stayed the same, eventually, the supply and demand situation will equalize, and the price will be back to normal. In a nutshell, the major factors influencing mineral demand include (Radetzki and Wårell, 2016):

- Incomes which are often represented by the Gross Domestic of a country.
- Industrial output
- The past behaviour of the commodity's price
- Prices of substitute and complement materials
- Technology improvements or new technology, which not only help reduce other industries' demand for a mineral commodity; but also affect the competitiveness of a mineral commodity in the market.
- Consumer preferences, which can vary depending on the consumer's income and economic scenario of a region/country.
- Government policies, regulations, activities

Supply, on the other hand, is mainly determined by the production of existing mines, recycling activity, mining production costs, among others.

As part of demand drivers, the consumer behaviour and the economy of lead producers/users of mineral commodities (countries) have to be considered, as well. In fact, China has been playing an important role due to its growing economic power during the last years (Tauber, 2015).

The prices of substitutes of commodities do also affect demand: when the price of such declines, demand for the mineral commodity decreases. In addition to the price-independent drivers, both supply and demand can also be affected by the price itself. That is, price affects both demand and supply as much as they influence price (Tauber, 2015). That is why any approach that fails to account for the inter-dynamics among supply, demand, and price is essentially flawed. Technology changes and the political status of a country, and in general any sudden event, also affect supply and demand. For example, political instability or unfavourable mining political policies can really hinder the financial behaviour of a commodity.

2.2.3. Scarcity of mineral commodities

There is, in fact, a great deal of literature on the scarcity of mineral commodities, i.e., a finite mineral resource stock and fluctuations in their prices. Many scholars suggest that price difference among different commodity prices can be attributed not only to the two main mineral commodity prices' drivers, i.e., demand and supply, but also to the mineral's scarcity, geography and metallurgical properties. According to Rudenno (2012), the scarcity of the minerals, the location of the minerals' ore bodies and their comminution processes, as opposed to the minerals' demand and supply, are the key drivers on the volatility of mineral commodity prices.

Nevertheless, the positive impacts of technological change on mineral supply through more sophisticated mineral exploration seem to have hampered the belief that mineral resource scarcity

plays a role just as important. In fact, recent changes in mining technology have offset the depletion of higher-grade and have allowed companies to exploit lower-grade deposits as they have become now viable from an economic and technical standpoint.

2.3. Commodity price uncertainty and risk management

Price forecasts are used by mining companies to make informed strategic decisions. As a result, they significantly impact any mining company's revenues and profits. They also play a role in the liquidity of a company and, consequently, in its financial planning process.

For instance, a decrease in commodity prices could:

- reduce producers' sales revenue, leading to not only a re-evaluation of the current business strategy but also to a possible decline of the company's total value.
- affect the viability of certain mine zones that were economically viable under the original prices. As a result, the production levels of the company's mines, as well as the whole mine planning and scheduling, could be changed to prevent potential losses.
- Potentially alter the company's current inventory management solutions, and, therefore, the company's earnings due to a decline in the inventory directly impact the company's dividends.

Whereas a rise in commodity prices could:

- increase the company's sales revenues and, thus, its financial value.
- encourage competition among producers as new entrants seek to take advantage of the higher prices.
- reduce profitability for consumers (of such commodities).

The forecasting of commodity prices is a difficult task. Although over the years substantial progress has been made and many different techniques from many different branches of economics, finance, and computer science have been developed, the increased volatility of commodity prices has made forecasting more burdensome than ever. Table 2.3 summarizes the effects of price variations on inventory, sales, purchasing, and revenue.

Table 2.3: Impacts of price movements

Price movement	**Impact**			
	Inventory	**Sales**	**Purchasing**	**Revenue**
Decline in commodity price	Leading to higher cost of inventory (constraint in cash flow)	Reduced sales volume and impact negatively profitability	Increase in purchasing power	Reducing earnings
Rise in commodity price	Leading to lower cost of inventory (increase in cash flow)	Increased sales values due to higher price	Decrease in purchasing power	Increasing earnings

In order to understand the challenges in mineral commodity forecasting, it is important to understand the nuances of commodity markets. According to Tauber (2015), the response of a commodity price behaviour as a result of changes in supply and demand varies across different commodities. He found that a decrease in supply would yield a much larger price increase in copper than in gold, with a percentage increase twice as greater. His analysis also indicated that copper prices are much more sensitive (than those of gold) to changes in mining costs (via mining cost inflation) and China's GDP. These price differences among different commodities can be explained by the variations of supply and demand curves of said commodities. That is, the location of the equilibrium point of the supply and demand curves in the supply-and-demand diagrams would indicate the sensible commodity prices are prone to.

Mineral commodity prices have somehow the same drivers for all commodities, but that the effect occurs in different degrees. Therefore, it is vital to identify those that affect them the

most. For example, the most powerful drivers of gold prices are the investor demand and political stability of the countries. However, copper and iron ore prices are more sensitive to Asian markets. (Arslanalp et al., 2016). That is why a convincing grasp of the commodity drivers would allow the investors to properly identify and construct different case scenarios that account for how the dynamic supply-and-demand evolves. That is, how supply- and demand-side players would respond to changes in market prices and its inevitable risks.

Currently, there are two ways a company usually deals with price uncertainty: it can be treated as a 'procuring risk' or as a 'tradable risk.' To mitigate said risk, the former focuses on improving supply chain management, whereas the latter puts more emphasis on the financial risk and hedging of the commodities (Deloitte, 2017).

Corporates dealing with procuring risk is used to concentrate all their efforts in a risk transfer and avoidance framework (Deloitte, 2017). However, nowadays, although it might seem that all risk is being passed on to consumers, corporations are increasingly embracing new practices aimed to find a balance in the supply chain so as to offset most of the consumer-associated risk. In fact, many companies have begun to develop management measures that involve cost control in the whole supply chain as well as negotiations with suppliers for the settlement of a fixed price (Deloitte, 2017).

Companies exposed to tradable risk have as common practice hedging the mineral commodity price risk. As explained in previous sections, hedging is the process of mitigating the risk of price fluctuations in the physical market by settling in advance the price for the same commodity in the futures market (Deloitte, 2017). There are two main reasons for a company to hedge the commodity price risk:

- To lock in a future price which is attractive

- To secure a price fixed against disruptions that might affect the price values

The decision to eliminate the price risk entirely or just a part of it is solely made by the company, which bases its choice on the amount of price risk it is prone to accept.

2.4. Deflator for mineral commodity prices

As mentioned, the historical price data were the most important component to forecast the commodity prices. These data are usually given as the current (nominal) values. Before applying a forecasting technique, the effect of the inflation should be eliminated. In other words, the current values must be converted to constant (real) values. This conversion is achieved by using the deflators or the price indices. The selection of an appropriate index is a difficult task.

A deflator is a statistic that allows price data to be measured at different points of time in terms of a base period, which can be any period of time.

As mentioned, deflators are usually price indices that are used to report historical inflation. An index relates the price of a particular year to the price of a year of reference, referred to as the base year of the price index series, which is given an index of 100. There exist many price indices; however, some of the more common include:

- CPI: Consumer Price Index
- PPI: Producer Price Index
- WPI: Wholesale Price Index
- CEI: Chemical Engineering Index
- M&S: Marshall & Swift Index for industrial equipment
- NR: Nelson Refinery Construction Cost Index
- MUV: Manufactures Unit Value Index

From those, the CPI is based on a fixed basket of goods and services purchased by average urban households, while the PPI program measures the average change over time in the selling prices received by domestic producers for their output. On the other hand, the MUV is used as a proxy for the price of developing country imports of manufactures in the US dollar terms.

Any commodity price-related work involves the unquestionable use of a deflator (CPI, PPI, MUV, etc.) as a prior tool to homogenize a data set (Cuddington and Jerrett, 2008). Usually, all the prices are expressed as a statistic of the actual values at a time given, in what is called nominal value. As a result, a comparison among prices at different points in time would be inaccurate.

No deflator can make the claim that it is universally "most relevant." Ultimately, the deflator selection depends on what relative prices one is most interested in. For example, from the perspective of a financial investor considering investments in commodities as an asset class, the CPI might prove to be more suitable as the investor is interested in knowing how their prices move over time relative to the CPI (Cuddington and Jerrett, 2008). On the other hand, if one is looking for the price of metal inputs relative to output prices, then one would want to select the particular outputs of interest. Here the PPI might be viewed as more relevant because it includes capital as well as consumption goods and excludes distribution costs and indirect taxes (Heap, 2005).

2.5. Mineral commodity price forecasting

Mineral commodity price forecasting is a complex task that has been addressed by different techniques; however, the ones that have been widely used, due to its understandability and simple structure, are the time-series methods.

Time series forecasting uses past observations of a variable to develop a model describing the underlying relationship between said observations. The model is then used to extrapolate the time series into the future (Zang, 2003).

There are, in fact, several different time series models used in commodity price forecasting. Traditional statistical models include moving average, exponential smoothing, ARMA, ARIMA, (ARCH), and GARCH.

The ARMA model combines Autoregressive (AR) and Moving Average (MA) models. It assumes that a sequence of random variables depends upon time and that the series mean and variance do not change over time (stationary). This assumption hampers the utility of these methods for mineral commodity price forecasting as commodity prices tend to have a non-stationary behavior. In response to this lack of application, the ARIMA model, the generalized version of the ARMA model, was developed. This model uses a linear function to describe past behavior and random errors. ARIMA models assume that lagged random changes of the dependent variables are represented by an algebraic equation with fixed coefficients estimated from past data (Cortez, 2018). Despite the fact that the ARIMA model offers a more realistic forecasting approach, it is not able to handle the volatility and nonlinearity that are present in the data series (Yaziz et al., 2013).

Both the ARCH and the GARCH models have been largely used to measure volatility in pricing (Rudenno, 2012). The perks of such models are that they account for the heteroscedasticity phenomenon (variances of the error are not constant), allowing to consider volatility in clusters as a feature of the model. Table 2.4 summarizes time series models used for forecasting along with pros and cons.

Table 2.4: Summary of commonly used time series models for forecasting

Type of model		Pros	Cons
Time series	MA and Exponential Smoothing	Simple and easy to implement	Fail to consider the diverse components of a time series
	ARMA	Takes into consideration past market behavior	The random variables rarely fluctuate in a regular pattern
	ARIMA	Flexibility to represent different time series, their statistical properties and the use of the Box–Jenkins methodology	Fails to account for the non-linear part of the series
	ARCH	Considers the heteroskedasticity of series, allowing to account for volatility clustering	Not useful for a point estimation. Better result when combined with ARMA
	GARCH	Inclusion of previously squared volatilities (variances) in the calculation of the future returns' variance	Provides an improvement form the ARCH model by including past variances

Recently, ANNs have been extensively studied in time series forecasting. The main advantage of this method is its flexible non-linear modeling capability (Zhang, 2003). That is why ANNs seem to be an alternative to traditional linear methods. In fact, compared to time-series methods, machine learning techniques have demonstrated better performance for forecasting mineral commodity prices over time (Mombeini and Yazdani-Chamzini, 2015).

ANNs have proved to add significant advantages in the forecasting “process”. One important advantage of the ANNs models (over both time-series linear and non-linear methods) is

that no prior assumption of the model form is required, being this model determined by the characteristics of the data, instead (Wongsathan and Seedadan, 2016).

In this thesis, we will thoroughly review, analyze, and compare the three more important forecasting methods used by the bulk of literature (ARIMA, GARCH, and ANNs), aiming to determine which one of those is more accurate and under what conditions one of them would prevail over the others.

Chapter 3

Methodology

3.1 Introduction

Three different forecasting models are studied, analyzed, and compared. The first two methods fall under the category of time-series forecasting techniques and the other can be identified as a machine learning technique. These methods are ARIMA, which is fully evaluated in Chapter 4; GARCH, which is thoroughly developed in Chapter 5; and ANNs, which is reviewed in Chapter 6. The performances of the proposed techniques above are tested through a case study using historical nickel prices. The price of nickel was adjusted for inflation, and its outliers were identified and corrected. The nickel prices are forecasted for the next 24 months. The outcomes of each technique are then compared.

The forecasts for each forecasting method are contrasted against the real values of nickel prices. The metrics to analyze each estimated value will be the prediction error, which is the difference between actual and forecasted values.

Since this metric does not say much about the prediction quality itself, each one of the metric values will be compared against each other to determine the most accurate model. The comparison across different models will be analyzed in different installments of 12, 18, and 24 months, aiming to find out which model is more suitable at different forecasting stages, and the reason behind this occurrence are discussed.

3.2. Data

The data includes 456 observations of monthly nickel prices (in USD per ton) from January 1980 to December 2017. The monthly nickel prices (in USD per ton) from January 2018 to December 2019 were used to validate the accuracy of the forecasted values. In order to conduct a proper test, all these values (from 1980 to 2019) were adjusted for inflation using the CPI as a deflator, which is "preferred" on the basis of the literature recommendations and the availability of the indexes' observations over time (Irwin and Sanders, 2011). Also, due to the lack of some indexes' availability for the time span of main data, the deflator selection was narrowed down to two choices: the CPI and the PPI. If this selection were to be made solely from the investors' perspective, the CPI could be viewed as more suitable; and, since one of the main goals of metal price forecasting is the decision to lean towards the use of the CPI (Irwin and Sanders, 2011). However, for the sake of affirmation, the trends of both indexes were compared graphically to better understand the possible differences of choosing one over the other.

Both the CPI and PPI values were taken from the website of the U.S. Bureau of Labor Statistics (with values updated as of November 2020). Figure 3.1 shows the log values of the CPI and the PPI from 1980 to 2019. As can be seen from this figure, the logarithmic values of both indexes resemble one another to some extent, with the CPI risen a little more rapidly than the PPI, though they both follow the same trend and fluctuations. Consequently, the trend of the real nickel prices would be somewhat less positive than that of the nickel price values deflated by the PPI, yielding more conservative dataset values. This detail, along with the fact that the selection of one index over the other would lead to little different real nickel prices, suggests that, in this case, the CPI is the most suitable deflator for nickel price forecasting.

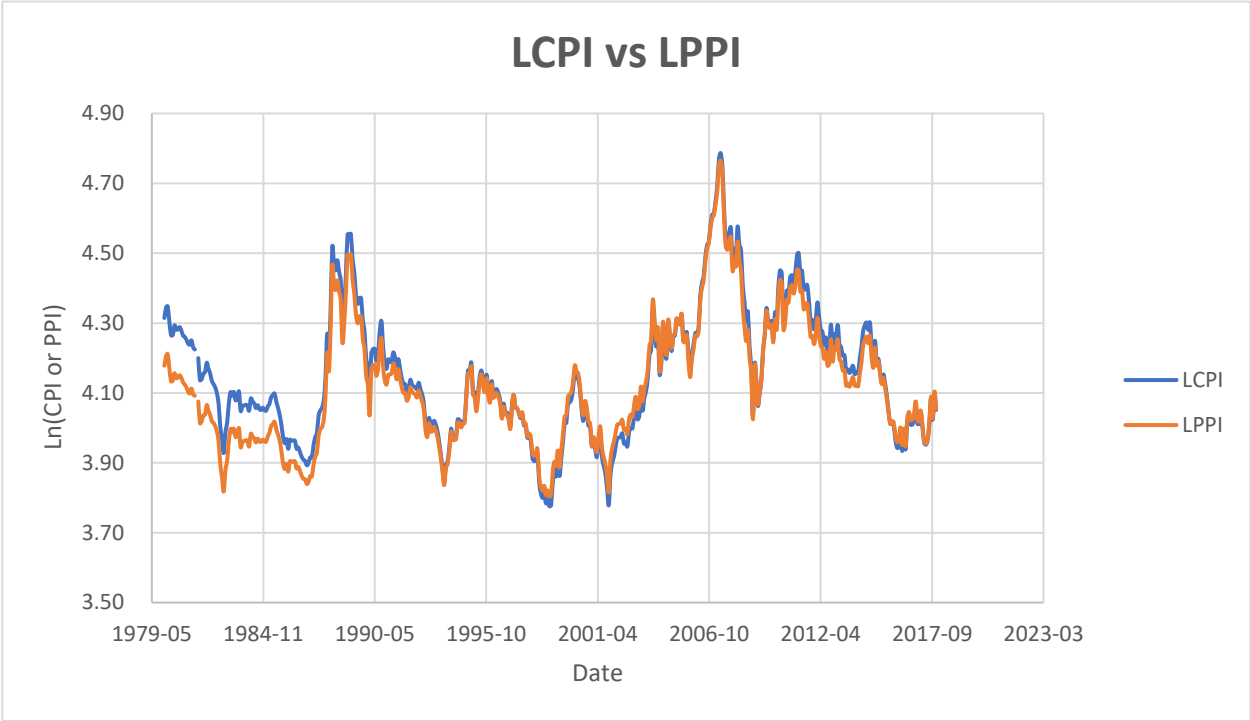

Figure 3.1: Comparison between logarithmic CPI values and logarithmic PPI

3.3. Performance evaluation

The main problem is how much deviations from actual or real values are acceptable. Theoretically , the forecasting techniques estimate future values; so, unless we establish some kind of benchmark that allows us to compare the real values and the forecasted values, there is no way to assess the quality of the forecasting model (Omar et al., 2016).

After the nominal nickel prices are converted to the real nickel prices, the real monthly nickel price values from January 2018 to December 2019 are used as reference benchmarks to assess the performance of the forecasted price. This procedure are applied to the different case studies using the three different forecasting techniques —ARIMA, GARCH, and ANNs.

In order to evaluate the performance of these forecasting methods, the criterion to be utilized will be the prediction error — the disparity between the forecasted values and the real

values. In fact, there are different error measures that research on fields such as engineering, computer science, statistics, and econometrics have adopted, making a consensus on the best performance evaluators a problematic task. However, it is possible to establish some kind of pattern in the machine-learning-based forecasting literature, allowing us to identify the most commonly used performance metrics rather than striving for the best metric. As it turns out, in 2018, Botchkarev (2018) carried out one of the most comprehensive studies regarding the error measures in machine learning regression, forecasting, and prognosis. Botchkarev (2018) reviewed over 100 peer-reviewed papers published from 1983 to 2017 and found that the most commonly used performance metrics are: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Coefficient of Determination (R^2). Botchkarev (2018) also found that although RMSE, MAE and MAPE are the metrics, most researchers rely on, the prevalence of one of them over the other has changed drastically over the 25-year timeline: in the 80s the dominance of RMSE was clear, whereas in the 90s MAE took over as the leading metric to assess machine-learning forecasting models and remained in that position until the beginning of the 2000s when MAPE and R^2 moved up to the first place. Furthermore, Botchkarev's research also sheds light on an important matter concerning forecasting model comparison: the use of different performance metrics to compare distinct forecasting models, aiming to pick the one that yields the most accurate forecasts. It turns out that the error measures previously mentioned cannot be used to compare different models unless they have the same units as well as belong to the same measure category. RMSE, MAE, and MAPE belong to a category called "primary metrics", whilst R^2 belongs to a category called "composite metrics" (Botchkarev, 2018). Primary metrics typically involve the calculation of point distance and the aggregation of point distance results across the data set values. This type of metrics also served as a foundation for the

construction of other metrics such as composite metrics and hybrid metrics. For instance, the coefficient of determination R^2, which can be categorized as a composite metric, involves the combination of the sum of squares (SS_E) and the sum of residuals (SS_{yy}), as seen in the formulas below.

Based on the extensive body of literature, the following performance measures were chosen: RMSE, MAE, MAPE, and R^2.

Let $\hat{y}_i$ be the expected or predicted values, y_i the observed or real values, $\bar{y}$ the average of the set of values of the series, and n the number of values of the data set; each one of the error measures can be expressed as:

$$\text{(i)} \qquad RMSE = \sqrt{\frac{\sum_{i=1}^{n}(\hat{y}_i - y_i)^2}{n}}$$

$$\text{(ii)} \qquad MAE = \frac{\sum_{i=1}^{n}|\hat{y}_i - y_i|}{n}$$

$$\text{(iii)} \qquad MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{\hat{y}_i - y_i}{y_i}\right|$$

$$\text{(iv)} \qquad R^2 = 1 - \frac{SS_E}{SS_{yy}} = 1 - \frac{\sum_{i=1}^{n}(\hat{y}_i - y_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}$$

The RMSE calculates the square root of the mean squared error, yielding a more representative measure since the RMSE is measured in the same units as the data set values as opposed to the mean squared error. This particular feature also allows the RMSE to be compared with different primary metrics with the same units as the data itself.

Botchkarev (2018) also ascertained that the RMSE might pose a hidden threat because when working with a large set of data, this error measure gives quite disproportionate weight to

very large pint distance among the data set values, leading to what is known as occasional large error. That is why it is widely suggested the data be adjusted for outliers before any kind of modelling.

MAE, as can be inferred from its formula, has the same units as the data's, and it is usually smaller than the RMSE since the latter is divided by the square root of *n* as opposed to just *n* (which is the denominator of the MAE). It is also believed that the MAE is a more precise measurement since it avoids the need for squaring the point distance in the calculation.

MAPE is perhaps the easiest-to-understand metric among those described above since its ability to express the error in terms of percentage is appealing to a greater group of researchers.

R^2 shows how much of the variability in dependent variable can be explained by the independent variable(s). R^2 is a positive number that can only take values between zero and one. A value for R^2 close to one shows a good fit for the forecasting model and a value close to zero presents a poor fit.

It is also worth stressing that the performance metrics of the same groups can only be compared between models whose errors are measured in the *same units*. That is, if one model's errors are in absolute units while another's are in logged units, their error measures cannot be directly compared. In such cases, it is advisable to convert the errors of both models into comparable units before computing the various measures. In this particular thesis, both the RMSE, MAE and MARE will be jointly used to compare forecasting models, and only the R^2 of each model will be compared among the others R^2s.

3.4. Modelling

For the cases studies in this thesis, the software R version 3.6.1 for Windows was used. All computational works was realized by a computer Acer, Intel(R) Core(TM) i5-8265U CPU @ 1.60GHz, with 8.00 GB RAM.

Chapter 4

Forecasting Through Auto-Regressive Integrated Moving Average (ARIMA) Models

4.1. Introduction

The commodity price is key to assessing both the economic and technical feasibility of a mining project; therefore, it plays a major role in the decision making of the whole hierarchy ladder of a mining company from upper manager to mine planners (Espí and De La Torre, 2013).

One common yet somewhat outmoded practice in the mining industry when assessing viability of a mining project is to simply overlook the future price fluctuations and set one fixed value over the entire operating life of the project. Though the risks of this approach can be, to some extend, hindered via sensibility analysis on the key variables determining project feasibility, this traditional method seems to ignore past observation patterns that might be helpful to describe the future behavior of certain metal price, doing this traditional practice not so ideal especially when dealing with real-world price uncertainty.

Mineral commodity forecasting has always been one of the most difficult tasks in many different areas such as economics, statistics, engineering, etc. undertake, often with no definite answer to the best forecasting model. In fact, the wide range of forecasting models includes some of the following: (1) those based on mathematical and statistical calculations such as time-series methods (mainly studied and covered by quantitative finance) in what is called objective forecast,

and (2) those based on qualitative techniques such as experts' opinions and supply and demand analysis in what is known as subjective forecast (Rudenno, 2012). Although some progress has been made over the past few decades, the ultimate choice of the most suitable forecasting method comes down to the following criteria: sample size, stationary, seasonality, length of the forecast (short term or long term), among others (Rudenno, 2012).

For a long time, researchers, forecasters and alike relied on simple forecasting models, such as exponential smoothing and moving averages to estimate future prices and, though such methods are still widely used to quickly obtain a trend curve that might serve as a guide to how the future values will behave over a given time frame, their accuracy has been shown to be outperformed by more sophisticated methods such ARMA and ARIMA models (Brockwell and Davis, 2016). Both models, the latter being built on the former, belong to a larger forecasting model category called time-series forecasting models. This type of forecasting model uses past observations of a certain variable to describe the relationship and hidden patterns among those observations and extrapolates the future values of the said variable (Guidolin and Pedi, 2018). Due to their simplicity and easy implementation, the ARMA and ARIMA models have become widely accepted and even nowadays are the "go-to" models for stock and commodity price forecasting. ARMA is built on the assumption that the observations vary over time in a regular pattern (Ru and Ren, 2012). That is to say that, over time, the observation values fluctuate around a constant mean and wiggles in a regular pattern, i.e., stationary (Nau, 2019). ARIMA model, which is an extension of ARMA, describes the behavior of the past observations, or rather the lagged values of those observations, and the random errors using a linear function. That is, the coefficients of the linear function are calculated based on the behavior of past data observations. The advantage of the ARIMA models over the ARMA models is the removal of the stationary constraint, allowing for treatment and

modeling of more real-world data, which most of the time does not resemble a stationary behavior (Paolella, 2019).

In this chapter, we aim to design an ARIMA model to predict the future price values of nickel in the short run. In fact, in the bulk of literature regarding stock price forecasting, time-series forecasting models have proven to be more accurate when the time span chosen ranged from a couple of months to up to a few years (Ru and Ren, 2012).

This chapter is structured as follows: Section 4.2. discuses the literature review regarding the use of ARIMA models in different settings, such as stock prices, metal prices, and agricultural commodity prices. Section 4.3 explains the basic concepts and frameworks of ARIMA models and how the key parameters in the model architecture can be found. Section 4.4 develops an ARIMA model for nickel price forecasting for the short run. This section analyzes and determines the best ARIMA model that fits the data given, using the model parameters to predict the future values of nickel price. It also discusses the results of the case study and analyzes what the results might mean in the context of the ARIMA model accuracy to estimate forecasts. This section summarizes the main conclusion and takeaways of the case study, as well.

4.2. Literature review

A time series is a sequence of data points taken successively over time (Box et al., 2016). The intervals of time at which those observations are taken are equally spaced (Chatfield, 2010). That is, when collecting data for a certain variable, the time interval must remain constant and could take up any value such as months, days, hours, etc.

In order to better understand a time series and a forecasting technique, it is imperative to be familiar with its different components. According to Chatfield (2010), all time series can be

decomposed into the following elements: (i) a *seasonal component or variation*, which is the type of variation that arises when observing a similar pattern over the same particular time of a year. (ii) The *trend* component is present when, over a prolonged period of time, the time series displays either upward or downward behavior. This statement does not mean that the time series cannot exhibit some fluctuations within the observed period of time, but rather an overall decline or growth. (iii) The *cyclical* component, which usually ranges for longer periods from several years to several decades, but not seasonal variations. (iv) The other components, which can be *regular or irregular fluctuation* or variations the time series might display, include all remaining elements after the seasonal and trend components are removed.

Generally speaking, forecasting time series methods use past observations of a variable (univariate forecasting) or set of variables multivariate forecasting) to predict their future values. In the particular case of this research, the focus will be univariate forecasting of mineral commodity prices. These time series forecasting methods include ARIMA, ARIMA, ARCH, and GARCH, among others.

The ARIMA models are perhaps the most valuable and common element in a forecaster's toolkit. It is both one of the most widely used forecasting models as well as the model that serves as a ground for more advanced and technically demanding forecasting methods. The foundation of the ARIMA models can be attributed to the early work of Box and Jenkins — the Box-Jenkins model.

Box and Jenkins cemented the basics for time series control, analysis and forecasting (Chatfield, 2010), but with the inclusion of Reinsel and others they developed a more thorough revision and analysis of ARIMA models that included a new outlier detection process, a revised test for unit roots and an extended estimation procedure (Box et al., 2016).

To this date, there are a few papers that have solely focused on the use of ARIMA for mineral commodity price forecasting, with the majority of them being an assessment of the different time series forecasting methods (Matyjaszek et al., 2019). This finding can be attributed to the fact that ARIMA models are considered to be of the simplest estimation models there are and, to some extent, not worthy of study alone. Nonetheless, in this research, a whole section is devoted to the ARIMA-based forecasting model because it is the author's belief that in order to make a better comparison among different forecasting models, a detailed explanation of each of the models is needed.

4.3. The ARIMA model

4.3.1. ARIMA model components

In the ARIMA-based forecasting model, the future values of a variable are a linear combination or a linear function of random errors and past observations or the lagged values of those. That is, an ARIMA model can be explained by a set of three different parameters (p, d, q). Each one of these parameters is related to a component of the ARIMA model (Lütkepohl, 2005).

- The auto-regressive component (p) refers to the use of past values in the regression. The parameter "p" specifies the number of terms in the model that describe the dependency among successive observations and can be thought of as the memory of the model regarding past data observations. For example, a model with $(p = 0)$ means that there is no relationship between adjacent observations, whereas a model with two auto-regressive terms $(p = 2)$ is one in which an observation depends on (is predicted by) two previous observations. This model's regression equation could be expressed as:

$$y_t = c + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + e_t \quad (4.1)$$

Where y_t shows the variable y measured at time t, y_{t-1} shows the variable y measured at time $t-1$, y_{t-2} shows the variable y measured at time $t-2$, α_1 and α_2 are the numerical coefficients calculated from simulations, c is a constant, and e_t is the random error.

- The Moving average component (q) represents the error of the model as a combination of previous random shocks or error terms e_t. In a nutshell, the term "q" denotes the number of errors to be included in the model. That is, a moving average process with $q = 0$ does not have moving average components, whilst a moving average process with $q = 2$ could be expressed as:

$$y_t = \theta_1 e_{t-1} + \theta_2 e_{t-2} + e_t \quad (4.2)$$

Where y_t represents the variable y measured at time t, e_t represents the random shock at time period t, e_{t-1} represents the random shock at time $t-1$, and e_{t-2} represents the random shock at time $t-2$, θ_1 and θ_2 are numerical coefficients.

- The Trend component (d) refers to the number of differences needed among pairs of observations to make a non-stationary series to be stationary. That is, we need to find a series that vary around a constant mean level, neither decreasing nor increasing systematically over time, with constant variance. For example,

$$If\ d = 0, Y_t = y_t$$

$$If\ d = 1, Y_t = y_t - y_{t-1}, and\ so\ on$$

Where Y_t is the differenced series, y_t represents the variable y measured at time t, y_{t-1} represents the variable y measured at time $t-1$, and d is the number of differences needed to make a process stationary.

Therefore, a general equation for an ARIMA process, which is the mix of both moving averages and autoregressive components can be written as:

$$y_t = c + \alpha_1 X_{t-1} + \ldots + \alpha_p X_{t-p} - \theta_1 e_{t-1} - \theta_q e_{t-q} \quad (4.3)$$

4.3.2. Stationarity

A time-series model is considered to be stationary when the process remains in what Box et al. (2016) defined as "statistical equilibrium." That is, the time-series data points vary over a fixed constant mean and with constant variance. This particular property is what sets ARMA and ARIMA models apart. As it turns out, one of the main requirements of ARMA models is for the time series to be stationary; whereas ARIMA provides a more flexible option as it "relaxes" this requirement and allows for the modeling of any type of data, both stationary and non-stationary ones (Paolella, 2019). In fact, ARIMA models have been proven to be of quite an importance as in the real world of business, engineering, mathematics, finance and economics the data presents itself with no defined patterns and is better represented by non-stationary processes (Dooley and Lenihan, 2005).

Checking whether a time series has the property of stationary can include different procedures ranging from simple inspection of the series graph to more sophisticated computational tests. The simplest method to check for stationary is to inspect the graph of a time series, draw an imaginary line that represents the means and assess whether the different upward and downward

fluctuations vary around it, offsetting one another over the period of time selected (Brockwell and Davis, 2016). However, this method is not completely reliable as some time series might appear to be non-stationary when, in reality, they are, and vice versa (Berlinger et al., 2015). The other procedure, on the other hand, involves the so-called Augmented Dickey-Fuller (ADF) test. The ADF test has the null hypothesis that the time series is non-stationary. Mathematically:

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} \sum_{i=1}^{h} \delta_t y_{t-p} + e_t \tag{4.4}$$

Where α and β are numerical coefficients (different from those of equations above) of the time trend and p is the lag order of the autoregressive process. The ADF test using a *t*-test checks the following hypothesis.

$$H_o: \gamma = 0 \text{ (non-stationary)}$$

$$H_A: < 0 \text{ (stationary)}$$

The intuition behind the ADF test is that testing $\gamma = 0$ is equivalent to y_t following a unit root process. In fact, the equation the ADF test uses can provide different insights if a careful look is taken at each one of their components. If both $\alpha = 0$ and $\beta = 0$, we can say that y_t follows a random walk process with drift since both the constant and time trend are equal to zero. On the other hand, if $\beta = 0$, y_t would follow a random walk process with drift. The analysis of the different values of α and β is important since, though no change needs to be made in the setup of the hypothesis, the values of the t statistics the rejections zones are under changes as both coefficients take up on different values (a combination of zeros).

4.3.3. Differencing

One way to make a non-stationary time series stationary is to compute the differences between consecutive observations. This is known as differencing.

Differencing can help stabilize the mean of a time series by removing changes in the level of a time series, and therefore eliminating (or reducing) trend and seasonality. Sometimes, transformations such as logarithms can help to stabilize the variance of a time series and need to be made before differencing the time series.

Another way to identify the number of differences a time series needs to be stationary is to look at the Autocorrelation Function (ACF) plot. For a stationary time series, the ACF will drop to zero relatively quickly, while the ACF of non-stationary data decreases slowly (Hyndman and Athanasopoulos, 2018). This particular statement will be thoroughly covered in the next section.

4.3.4. Autocorrelation and Partial Autocorrelation Functions (ACF/PACF)

The rule of thumb for identifying the essential parameters for any ARIMA model lies in the proper analysis of the ACF and Partial Autocorrelation Function (PACF). In order to understand both functions, it is important to be familiar with a concept that, at first glance, it might seem self-explanatory. In fact, both the autocorrelation and partial autocorrelation functions are computed on a series of sequential lags. A *lag* can be defined as the time periods between two different points of a time series. For instance, lag 1 is between Y_t and Y_{t-1}, lag 2 is between Y_t and Y_{t-2}, and, in general, lag k is between Y_t and Y_{t-k}.

4.3.4.1. Autocorrelation function (ACF)

The autocorrelation of Y_t at lag k, which gives correlations between Y_t and Y_{t-k}, can be defined by

$$\rho_k = \frac{cov(Y_t, Y_{t-k})}{\sqrt{var(Y_t).var(Y_{t-k})}} \tag{4.5}$$

Since the process is defined for stationary time series, the value of the varies does not change across time and, therefore, ρ_k formula can de simplified to

$$\rho_k = \frac{cov(Y_t, Y_{t-k})}{var(Y_t)} \tag{4.6}$$

The plot of ρ_k against lag k is what is called the ACF. Furthermore, the equation above can be farther simplified to

$$\rho_k = \frac{\frac{1}{N-k}\sum_{t=1}^{N-k}(Y_t-\bar{Y})(Y_{t-k}-\bar{Y})}{\frac{1}{N-1}\sum_{t=1}^{N}(Y_t-\bar{Y})^2} \tag{4.7}$$

Where N is the number of observations on the whole series, k is the lag, $\bar{Y}$ is the mean of the whole series,

$$\bar{Y} = \frac{1}{N}\sum_{t=1}^{N} Y_t \tag{4.8}$$

and the denominator is the variance of the whole series, as well.

4.3.4.2. Partial autocorrelation function (PACF)

The partial autocorrelation is technically a conditional correlation. That is, it is the correlation between two variables when the value of another one or a set of them is considered to

be known (assumption). In fact, the partial autocorrelation function can be thought of in the context of linear regression with Y as a dependent variable and x_1, x_2, x_3 as independent variables. The partial autocorrelation between Y and x_3 would be equal to the correlation considering how both Y and x_3 are related to x_1 and x_2. That is, in the context of a linear regression the partial autocorrelation function can be expressed as

$$\frac{covariance(Y, x_3|x_1, x_2)}{\sqrt{variance(Y|x_1, x_2).variance(x_3|x_1, x_2)}}$$

This equation can be translated in terms of a time series. Let Y_t be the values of the time series observation at time t, and k is the lag. The partial autocorrelation between x_t and x_{t-k} can be defined as the regression between x_t and x_{t-k}, conditional on x_{t-1}, x_{t-2}, …, x_{t-k+1} (set of variables that comes between x_t and x_{t-k}). Therefore:

- The partial autocorrelation at lag 1 is equal to the autocorrelation function at lag 1.
- The partial autocorrelation function at lag 2 can be expressed as

$$\frac{Covariance(x_t, x_{t-2}|x_{t-1})}{\sqrt{Variance(x_t|x_{t-1}).Variance(x_{t-2}|x_{t-1})}}$$

(the values of the variances are equal provided that the time series is stationary)

- The partial autocorrelation function at lag 3 is defined by

$$\frac{Covariance(x_t, x_{t-3}|x_{t-1}, x_{t-2})}{\sqrt{Variance(x_t|x_{t-1}, x_{t-2}).Variance(x_{t-3}|x_{t-1}, x_{t-2})}}$$

And so on for any lag.

4.4. Case study: Nickel price forecasting

4.4.1. The best set of parameters (p, q, d)

The selection of the parameter d is perhaps the easiest one to make. As pointed out above, this parameter can be determined either visually or through computing methods. Since the former has already been explained, only the latter will be described. The software R provides the *diff* package. The only input needed is the time series and the command yields the number of differences needed to make the series stationary (Tsay, 2012).

The ACF is a tool in identifying the order of an $MA(q)$ process because it is expected to "cut off" after lag k. However, the ACF is not as useful in the identification of the order of an $AR(p)$ process since it includes a mixture of exponential decay and sinusoid expressions, which cannot really tell what the order for the autoregressive part is. Hence, while ACF could help indicate that the process might have an AR structure, it fails to provide further information about the order of such structure (Hyndman and Athanasopoulos, 2018). For that, PACF of the time series has proven to be more useful.

For an $AR(p)$ model, the lag k at which the PACF graphs cuts offs is a good predictor of the value of p. Theoretically, PACF shuts off past the order of the model. That is, the partial autocorrelations are equal to 0 beyond the order of the model corresponding to the autoregressive part of the ARIMA process.

Nonetheless, the methodology explained above means little if the time series is not treated in the context of stationary. That is, since non-stationary time series have an ACF that remains significant for half a dozen or more lags, rather than quickly declining to zero, it is vital the time series be differenced until stationary is achieved. The values of d correspond to the number of

differences needed to achieve said stationary. Once differenced, the stationary property of a time series need to be checked until both visually and statistically it meets the requirements of stationary.

Table 4.1: Summary of methodology for non-seasonal ARIMA parameter's identification

Process	**ACF**	**PACF**
$\boldsymbol{AR(p)}$	X	Helps to identify the value of p
$\boldsymbol{MA(q)}$	Helps to identify the value of q	X

4.4.2. Future prices

As mentioned before, the data consists of 456 observations of monthly nickel prices (in USD per ton) from January 1980 to December 2017. The monthly nickel prices (in USD per ton) from January 2018 to December 2019 were used to validate the accuracy of the forecasted values. In order to conduct a proper test, all these values (from 1980 to 2019) were adjusted for inflation using the Consumer Product Index (CPI) as a deflator.

The data's noisy fluctuations were smoothed using the moving average technique (though any particular smoothing technique is feasible, and none represents a 'better' fit). This moving average smoother of order 'm' averages the nearest 'm' periods of each observation. As neighboring observations of a time series are likely to be similar in value, averaging eliminates some of the randomness in the data, leaving a smooth trend-cycle component. The mathematical expression of a moving average can be expressed as:

$$T_t = \frac{1}{m}\sum_{j=-k}^{k} y_{j+k} \tag{4.9}$$

where

$$k = \frac{m-1}{2}$$

This moving average is not the same as the one described at the beginning of this thesis, as this smoothing technique is related to the actual values of our variable rather than the errors. Theoretically, the higher the order of the moving average, the smoother the original series becomes (Tsay, 2012). However, the order value chosen should neither smooth the series too much so as to 'arrest' the series' trend nor help very little to smooth the data. Based on repeated attempts, a moving average smoother of order eight was selected.

Then, since a time series is built on three essential components (seasonality, trend and random), the series was deconstructed to better understand its behavior and prepared a foundation for building a forecasting model. To that end, an additive deconstruction was carried out as a simple inspection revealed that, though the series did not follow any specific seasonality-trend pattern, the seasonal component did not generally increase as the trend went up (except for a few outlier values).

Figure 4.1 shows the different components of a time series. The time series contains a quite constant seasonal component that might have trouble in the fitting of the model, especially since the basic models are usually designed for non-seasonal models. The trend of the time series also shows that the series is not stationary; it will need to undergo any kind of treatment to make it stationary. When treating the seasonal component, one has to be careful not to forget it since that issue might lead to naïve forecasts. However, it is not always the case that the seasonal component bears importance to the forecast of a time series (Hyndman and Athanasopoulos, 2018).

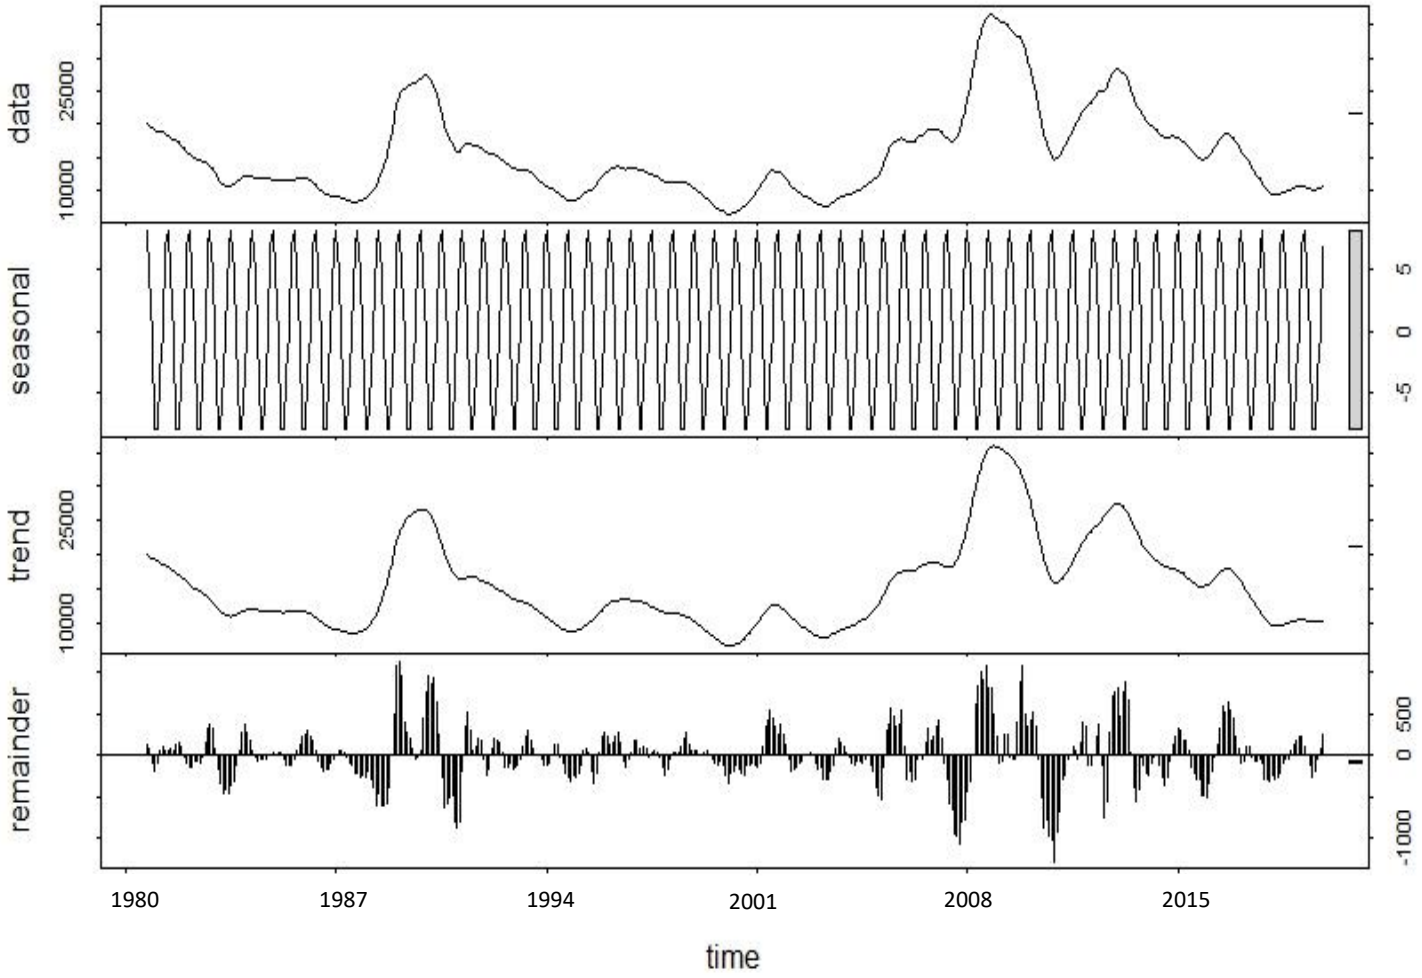

Figure 4.1: Seasonal, trend and random components of the time series

Since the first inspection of the time series graph reveals that the data might not be stationary, the Augmented Dickey-Fuller test was conducted. The values of the said test were:

Augmented Dickey-Fuller Test

Dickey-Fuller = -2.4792, Lag order = 7, p-value = 0.3752

alternative hypothesis: stationary

Since the p-value was greater than 0.05, the null hypothesis of non-stationary was rejected, in favor of the alternative hypothesis. That is, the ADF test confirms that the time series is actually stationary. This means that the times series need to be differenced to become non-stationary. To do so, the command *diff* was run using the software R. This command yields the number of differences the time series needs to get rid of the stationarity. After running the command *diff*, it was determined that the series only needed one difference to become stationary. To prove this, a new ADF test was conducted to see whether the p-value was less than 0.05.

Augmented Dickey-Fuller Test

Dickey-Fuller = -5.2069, Lag order = 7, p-value = 0.01

alternative hypothesis: stationary

Since now the p-value was less than 0.05, the null hypothesis was rejected in favor of the alternative hypothesis, confirming that now the differenced series is stationary.

After a time series was stationarized by differencing, the AR or MA was determined. A quick inspection of the ACF graph revealed that the correlation function cut-offs at lag 8-9, hinting

the possibility that q could take on the value 8 or 9. Although it might seem like PACF is a bit volatile, it actually shows some important insights that can come in handy to determine the p order of AR. The inspection also showed that the PACF cut off at lag 2 or 3. It can be argued that some values at some lags are well above the 95% significance boundaries (blue dotted lines), but that can be attributed to the fluctuation of the data as those values do not exceed the boundaries by that much and seem to have and overall constant trend).

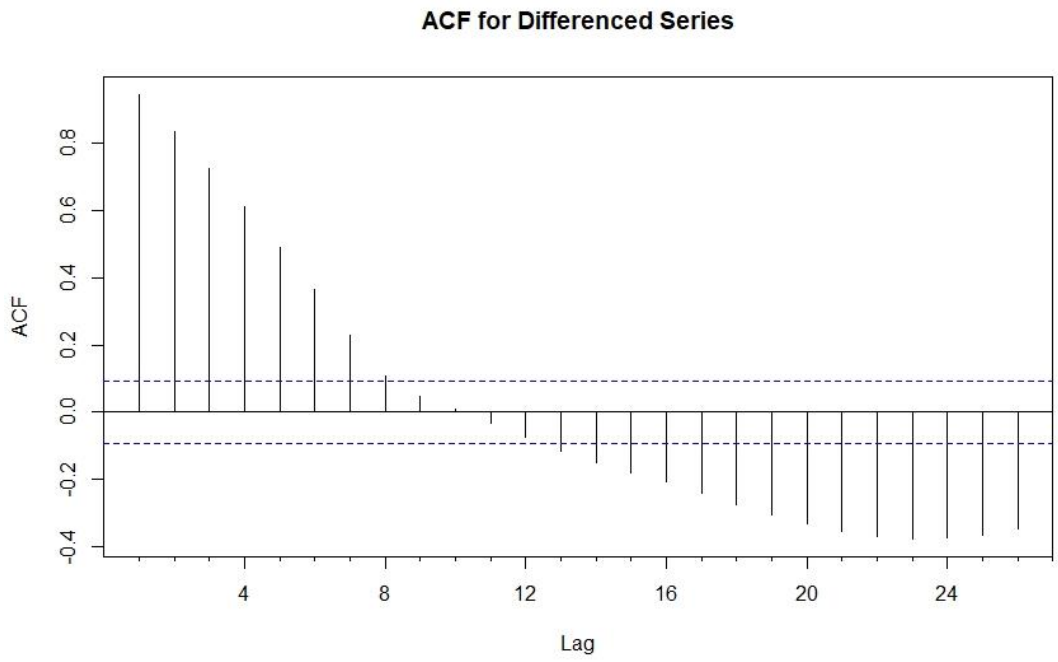

Figure 4.2: Autocorrelation function at different lags for the time series

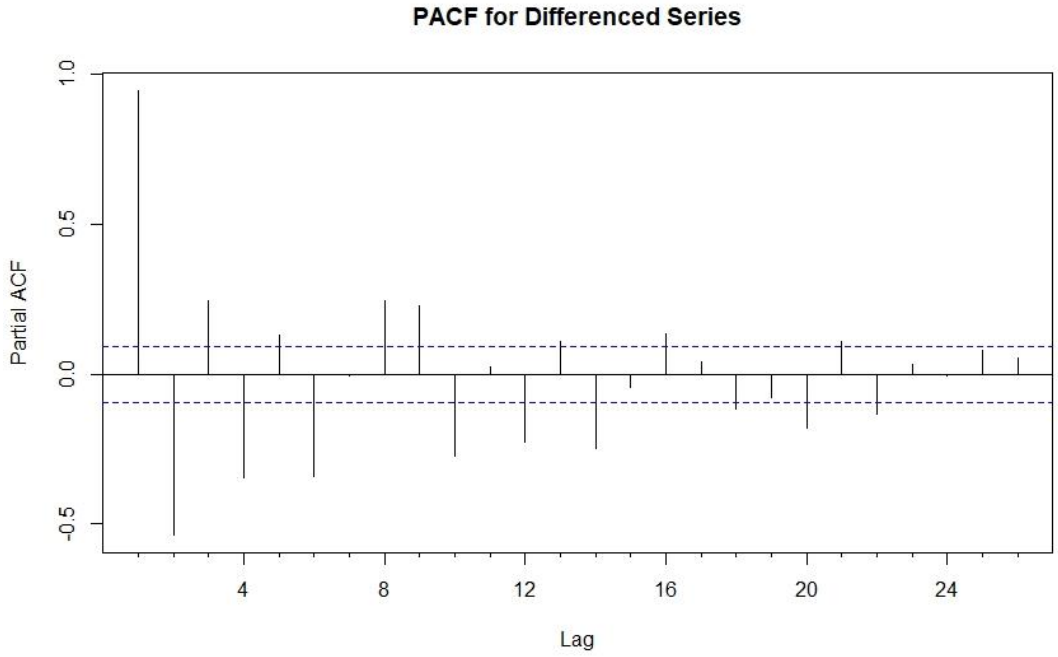

Figure 4.3: Partial Autocorrelation function at different lags for the time series

Both Figure 4.2 and Figure 4.3 show the autocorrelation and partial-autocorrelation function values of the series at different lags. Both figures allow us to identify the values of the parameters for the model. Note that both figures indicate a ballpark of the parameters and such need to be tested in order to find out what the best set of parameters is.

The different possible values of p and q mentioned above provided a series of combinations, with each one of them yielding different ARIMA models. The Arima models tested were (3,1,2), (3,1,8), (3,1,9) and (2,1,2). Also, the Akaike criterion (AIC) was used as a comparable measure of the model's suitability. The model with the lowest AIC values is the most suitable (Hyndman and Athanasopoulos, 2018).

Table 4.2 shows the values of the AIC for the different set of parameters selected. These values will help us determine what the best set of parameters really is.

Table 4.2: Values of Akaike coefficients for different set of parameters

ARIMA	AIC
(3,1,2)	5,643.53
(3,1,8)	5,446.01
(3,1,9)	5,447.73
(2,1,2)	**5,440.08**

Also, the assumption of no correlated errors was tested. A simple test of the residuals revealed that indeed there was no autocorrelation among the errors.

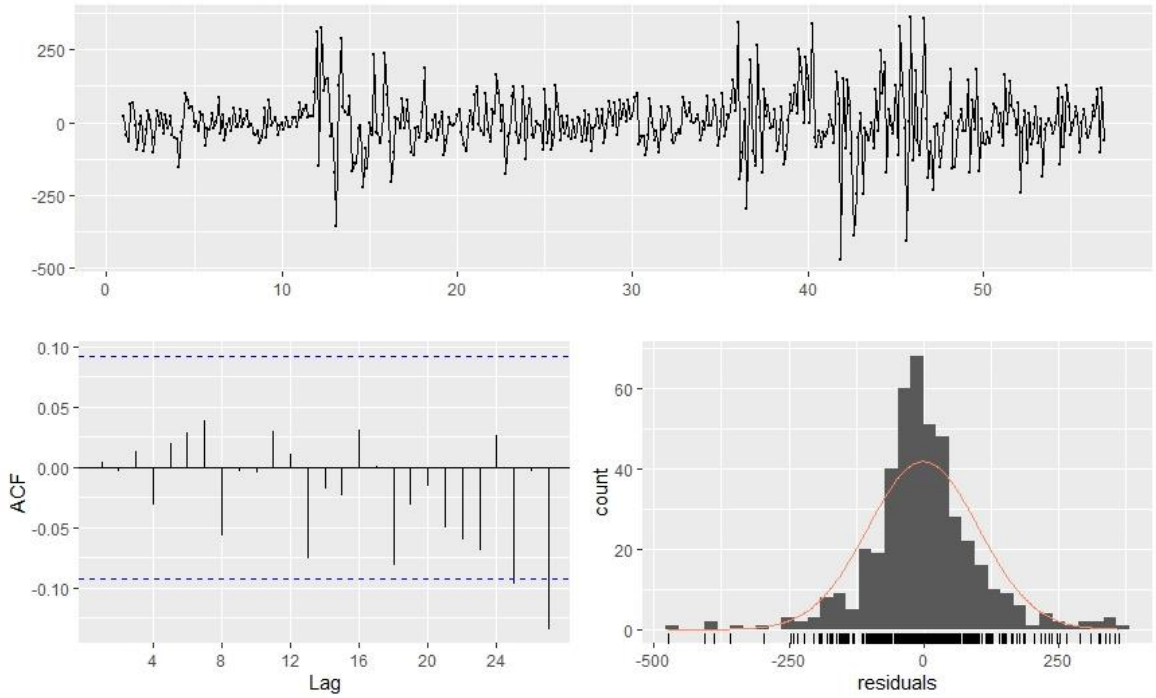

Figure 4.4: ACF and PACF of the model residuals for ARIMA (2,1,2)

Figure 4.4 shows the statistical properties of the residuals of the series once the model was fitted. For the model to be statistically right, the residuals need to behave like white noise (Hyndman and Athanasopoulos, 2018). That is, their autocorrelation needs to be as close as possible to zero. In order to check this assumption, a simple inspection of the spikes of the ACF could be made. If 95% of the spikes lie within the blue-dotted lines, then no correlation could be expected and the residuals can be considered as white noise. This is clearly the case for the graph above.

It is also worth pointing out that when testing the set of parameters mentioned above, the seasonally adjusted data was re-seasonalized as to avoid results that weren't consistent with what might be expected based on the past behaviour of the time series, such as obtaining horizontal or flat lines as forecast. That is, when running the command *forecast ()* in the software R, which calculates the forecasts for a time series given a particular model, once the specifications made was to include the seasonal component in the forecasting model since by default this command tend to seasonally adjust the time series to a non-seasonal model. This proved quite a useful technique since when the seasonal component was not included, not only the behaviour of the time series did not make sense (horizontal line), but also the values of the AIC obtained were much lower that those yielded when using a non-seasonal ARIMA model.

Finally, the best model was fitted to the time series, and the nickel prices for the next 24 months were calculated.

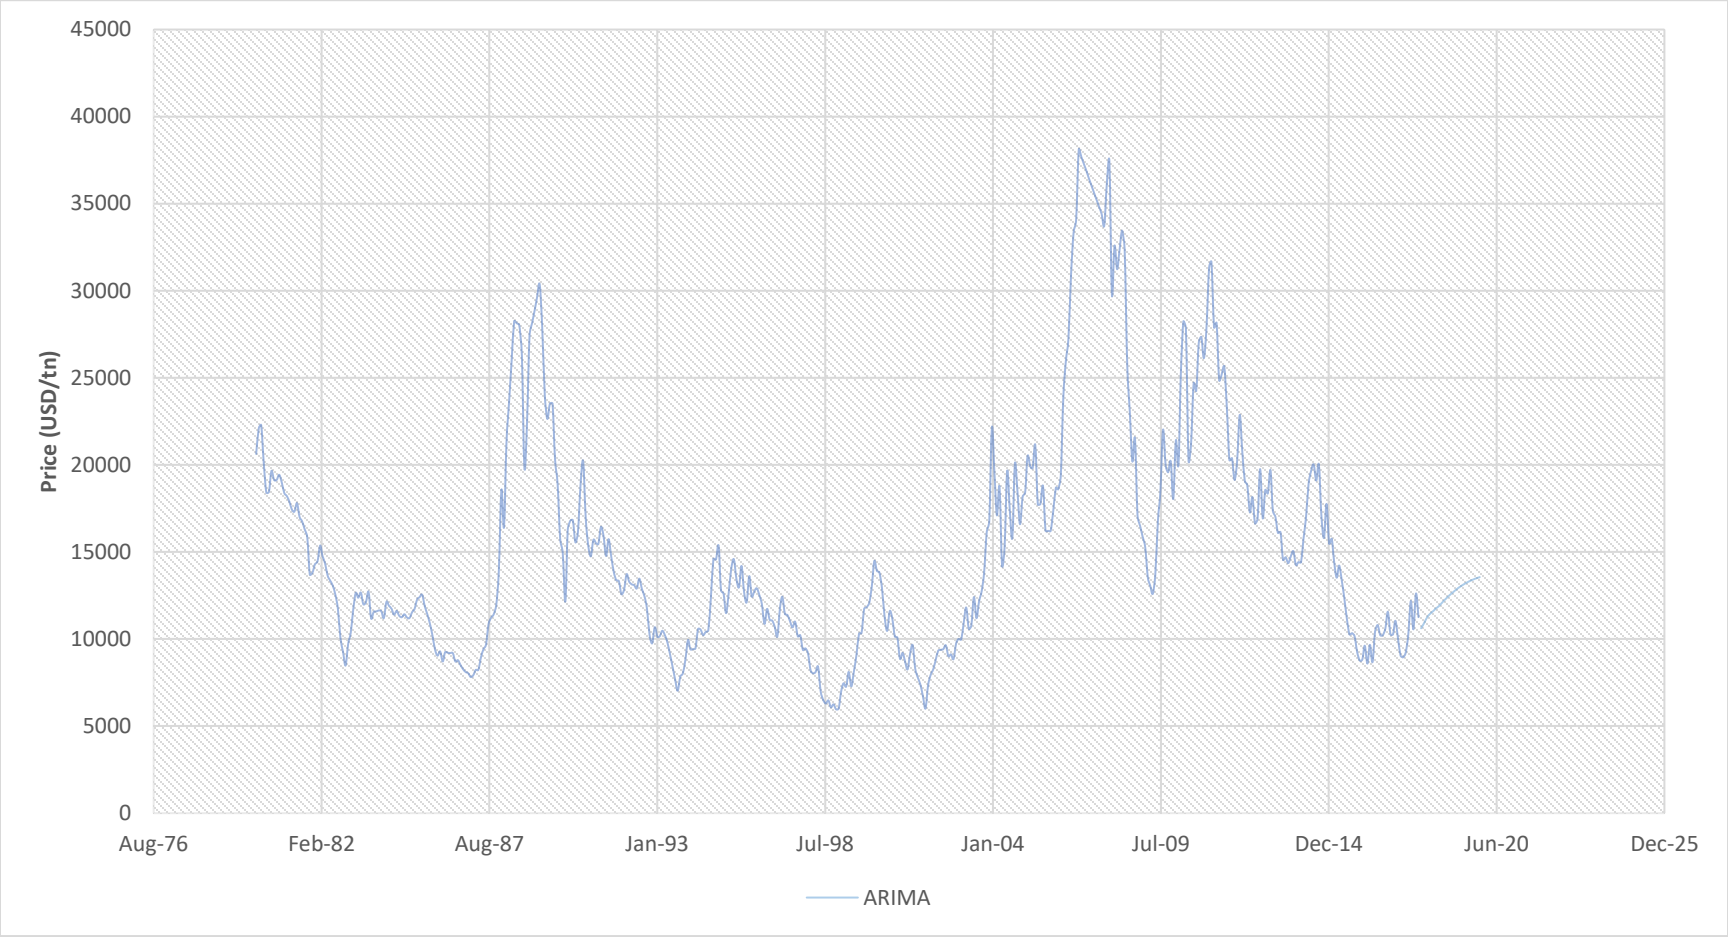

Figure 4.5: Forecasting of nickel prices for the next two years using ARIMA

Figure 4.5 shows a comparison of the real values of the nickel price (blue line) from January 2018 to December 2019. Since this thesis was written in 2020, the nickel price data can be taken advantage of in order to precisely determine how accurate the model is in predicting nickel prices. To have a better idea of the behaviour of the forecasted values, a line with said values was also graphed (blue line).

Table 4.3. shows the values of the forecasted nickel prices from January 2018 to December 2019.

Table 4.3: Summary of future values of Nickel prices using an ARIMA model

Date	Estimate (USD/tn)	Date	Estimate (USD/tn)
Jan-18	10,615.78	**Jan-19**	12,625.59
Feb-18	10,874.57	**Feb-19**	12,749.48
Mar-18	11,136.58	**Mar-19**	12,865.15
Apr-18	11,366.00	**Apr-19**	12,972.58
May-18	11,502.20	**May-19**	13,071.84
Jun-18	11,635.02	**Jun-19**	13,163.06
Jul-18	11,772.11	**Jul-19**	13,246.44
Aug-18	11,889.04	**Aug-19**	13,322.18
Sep-18	12,050.84	**Sep-19**	13,390.57
Oct-18	12,205.84	**Oct-19**	13,451.89
Nov-18	12,353.55	**Nov-19**	13,506.47
Dec-18	12,493.56	**Dec-19**	13,554.65

Chapter 5

Forecasting Through Generalized Autoregressive Conditionally Heteroscedastic (GARCH) Models

5.1. Introduction

Unlike both ARMA and ARIMA forecasting techniques, which model the mean of any time series, ARCH process and its more extended and comprehensive version, GARCH, change the main focus to the series' volatility (Bauwens et al., 2012). Time series' volatility forecasts are vital to a myriad of financial instruments such as risk measure and hedging (Viens et al., 2012).

Many researchers have suggested that the applications of ARMA and ARIMA models for forecasting are limited because the former assumes that the data show constant volatility (stationary), whereas the latter not only impose linearity constrains to the forecasting model but also are not able to handle volatility (Brockwell and Davis, 2016). To overcome these issues, both the ARCH and GARCH processes were developed to reflect the time-varying volatility most real-world time series show (Viens et al., 2012).

Financial data, especially stock returns and commodity prices, are heteroskedastic — the magnitude of their error terms changes over time, with intermittent periods of high and low values (Hyndman and Athanasopoulos, 2018). The natural yet vexed question that arises when forecasting such data is the accuracy of the forecasting model. In order to determine how accurate a forecasting model is, a parameter that plays a key role is the variance of the error terms of the time series

(Knight and Satchell, 2007). In fact, the square roots of the error terms' variance are the ones that provide a measure to assess the time series' volatility over time. Furthermore, despite what its name suggests, GARCH models not only allow for the modeling of the data error variance but also the mean of the process (Dritsaki, 2018). In some papers this forecasting technique is called "the hybrid models ARMA-GARCH"; however, in this chapter and forward, for the sake of uniformity and in completion with several computer packages such as R and Phyton, this process will be called GARCH.

In this chapter, we aim to design a GARCH model to estimate the future price values of nickel in both the short run and long term, considering the definition of the long run as the estimation of forecast one year or more ahead and short term as forecasts looking out less than one year in the future (Bernard et al., 2006). The reason behind this categorization is to test the hypothesis that as the forecast horizon increases, the efficiency of the different forecasting methods considered in this research changes, establishing different groups of time-series forecasting techniques that work better depending on the forecasting term (Rudenno, 2012). Furthermore, since the ARIMA-based forecasting results have already been determined, a comparison between the forecast using an ARIMA model and a GARCH model will be conducted. A series of error measures will be calculated to numerically assess the accuracy of both models as well as to determine which model represents a better fit for the nickel price data given.

This chapter is structured as follows: Section 5.2. discuses the literature review regarding the use of GARCH models for forecasting purposes in different settings, such as oil prices, metal prices and electricity prices. Section 5.3 explains the basic concepts and frameworks of ARCH and GARCH models, as well as their differences. This section also addresses how the key parameters in the GARCH models can be found. Section 5.4 describes the GARCH model that

best fits the nickel price data given and uses such a model to estimate the future value of nickel prices. Section 5.4 also discusses the results of the case study and analyzes the efficiency of both the ARMA model and GARCH model forecasting processes. The main takeaways from the analysis are provided at the end of this section.

5.2. Literature review

In 1982, Nobel prize winner Robert F. Engle introduced the ARCH models aimed to tackle economic time series with time-varying volatility (Engle, 2001). A few years later, Bollerslev (1986) introduced the more general class of the ARCH process, commonly known as the GARCH process (Hyndman and Athanasopoulos, 2018). Before the introduction of both methods, there were no forecasting methods for the modeling the variance of the residuals, as most tended to focus on the estimation of the mean. The closest empirical method to the ARCH and GARCH models back then entailed the calculation of the standard deviation using a fixed amount of the latest observations of the time series (Engle, 2001). At that time, this type of estimation made a lot of sense due to the belief that most recent observations played a greater role in the determination of the variance of tomorrow's residual and that the earliest observations' coefficients were non-significant in comparison to the latest's, assigning a zero coefficient automatically to the variance of older past observations (Engle, 2001). Both the ARCH and GARCH processes shifted that constrain by allowing the weight of the variance of each of one of the residuals to be determined by the data instead (Engle, 2001).

Given the importance of the ARCH/GARCH models to the financial markets, most of the ARCH/GARCH-related research done, yet limited, has been conducted using price returns instead of the prices themselves. This finding follows the idea that the price returns, or returns of any

financial data for that matter, exhibit much more attractive statistical properties, especially when it comes to normalization and stationary properties (Mills, 2015). However, even among return indicators, there has been a strong debate as to which, namely the simple returns or the logarithm returns of a financial variable, is the better fit when it comes to dealing with financial data. There is a great bulk of finance-related papers making the case for logarithm returns over simple returns. The consensus is that financial data modelling would be better off when using logarithm returns due to their time-additivity and complexity properties (Hudson and Gregoriou, 2015). When dealing with returns, there is a widely-used term called compounding returns, which is the cumulative return of the series over time. The problem with compounding return is that, though returns themselves are normally distributed, their product is not (Hudson and Gregoriou, 2015). Therefore, the use of logarithms provides not only a normally-distributed foundation onto which financial data can be modelled but also a much easier architecture, since only the logarithm of the initial and last period variables are used in the compounding returns' calculation. Given the unanimous accord towards the use of the log returns (of prices), these values will be the input for the case study to be described lines below.

The assumption under which the ARCH model was developed is that the variance of future returns is a function of the variance of the residuals of previous observations (Engle, 2001). The main difference between the ARCH and GARCH models is the inclusion of previously squared volatilities (variances) in the calculation of the future returns' variance (Bauwens et al., 2012).

The body of literature on commodity price forecasting using GARCH processes is, to some extent, scarce. A handful of them uses gold prices, oil prices, and energy prices as case studies for GARCH-based forecasting. Many applications consider a hybrid method that includes the combination of ARMA and GARCH models so as to account for both the linear and non-linear

features most real-world time series exhibit. For example, Yaziz et al. (2013) considered a hybrid model ARIMA-GARCH to forecast daily gold prices. Using a sample of daily gold prices taken every five days from November 26th, 2005 and January 18th, 2006, they used a two-phase process that involves 1) the modeling of the linear data of the time series through the use of an ARIMA process, and 2) the modeling of the residuals of the first phase (or non-linear data) using a GARCH model. When comparing the hybrid model ARIMA-GARCH against the single forecasting models ARIMA and GARCH, they found, using MAE and MSE as error measures, that the hybrid model represents a potential better forecasting model in terms of accuracy. This finding, according to them, confirmed the research-based idea that the combination of different models aimed to tackle any time series' different features might outperform the use of single forecasting models (Newbold and Granger, 1974) (Zhang, 2003). Dritsaki (2018) used the same approach, though this time the variable is not oil prices themselves, but rather the return of oil prices since the effectiveness of GARCH models have been proven and widely studied in the context of financial data such as the return of stock prices (Ruppert, 2011), let alone the almost unanimous consensus that they use of log returns provide a better statistical foundation in the context of normalization and stationary properties.

In terms of the selection of the GARCH model for forecasting commodity prices, there is a large body of literature suggesting that, in fact, GARCH process does actually provide better performance over other parametric and non-parametric time series forecasting methods. For example, Hou and Suardi (2012) compared GARCH models and other non-parametric methods to both model and estimated the volatility of oil price return and found that a GARCH forecast for data volatility does yield superior performance. This finding was verified by Agnolucci (2009), whose research showed that GARCH models provide significant forecasts for the volatility of

crude oil futures prices. The research on the superiority of GARCH models over some other time-series forecasting methods is not only limited to the volatility of oil price returns; the studies such as those of Wang and Wu (2012), which focused on the forecast energy market volatility, make a case for GARCH forecasting models, with the univariate GARCH process being outperformed by the multivariate GRACH model. Nonetheless, for the sake of simplicity and based on previous research done in the area of metal price forecasting, this chapter, and in general the whole thesis, will focus on univariate forecasting models, as the data selected is composed of one single variable — the mineral commodity price.

5.3. The GARCH model approach

For quite a long time, ARMA and ARIMA processes were the most applied models for financial forecasting. However, as more research was conducted in the quest for finding a forecasting model that can most reliably predict the future values of the series, linear models started losing momentum in favor of a more stylized model that could explain some of the series' features that, otherwise, might have been overlooked (Lama, 2015; Tully and Lucey, 2007). Properties such as volatility clustering — the tendency of large changes in the main (and unique) variable to be followed by other larges changes, forming a sort of cluster along with the series of persistent volatility (Engle, 2012) — cannot be captured by the tradition ARIMA models.

Let P_t be the price of a financial asset at time t, $Z_t = \log P_t - \log P_{t-1}$ the log return of the said price at time t, and h_t is the conditional variance. An ARIMA process will consider h_t to be constant over time (and, therefore, independent of t) (Bernard et al., 2006). Nonetheless, real-world time series rarely show stationary behavior over time. Therefore, the foundation of ARCH and GARCH models aims to incorporate the effects of such properties into the model design. That

is, the volatility, measured by the conditional variance h_t of the series, will depend on the past values of Z_t^2

$$h_t = a_0 + \sum_{i=1}^{p} a_i . Z_{t-i}^2 \quad (5.1)$$

Where $Z_t = \sqrt{h_t} . \varepsilon_t$ with ε_t being independent and identically distributed as well as normally distributed with $N(0,1)$. Also, p is some integer and i is an integer that can take on the values 1 to p, a_0 is a constant greater than zero ($a_0 > 0$), and the coefficients a_i can be equal or greater than 0 ($a_i \geq 0$). The value of p depends on the number of past observations that the models used in the calculation of the conditional variance (Bauwens et a., 2012).

The GARCH model, on the other hand, postulate a more comprehensive model by including past volatilities into equation 5.1.

$$h_t = a_0 + \sum_{i=1}^{p} a_i . Z_{t-i}^2 + \sum_{i=1}^{q} b_i . h_{t-1} \quad (5.2)$$

Where i is an integer that can take on the values 0 to q, the constant $a_0 > 0$, and the coefficients $a_i \geq 0$, $b_i \geq 0$

5.4. Case study: Nickel price forecasting

The previously described dataset in Sections 3.2 and 4.4.2 was also used for this case study. The inflation was also adjusted using the CPI.

Since the majority of research using GARCH models for forecasting purposes has been conducted on financial time series and due to their time-additivity and normalization-related properties, it is vital to wok with the log returns of the original nickel prices (Bauwens et al., 2012;

Knight and Satchell, 2007). Let P_t be the nickel price at time t and Z_t be the log returns of the nickel prices at time t; then this value can be calculated using the following equation:

$$Z_t = \log P_t - \log P_{t-1} \tag{5.3}$$

5.4.1. The best set of parameters (m, s)

There are many procedures to determine the best set of parameters (m, s) for the GARCH process in price forecasting, neither representing an absolute better approach over the other. In fact, the choice of the best set of parameters for the GARCH models has proved to be a difficult task and, therefore, the model GARCH (1,1) has long been the gold 'standard' when it comes to GARCH modelling of financial data due to its simplicity (Chatfield, 2010). One of the most recommended and widely used tools is the identification of both the autocorrelation function (ACF) and partial autocorrelation function (PACF) plots of the series Z_t (log returns of the nickel prices) and their squares ${Z_t}^2$. For instance, if Z_t appears to be white noise — the series do not show any autocorrelation (or it is close to zero) — shown by the lag spikes in the autocorrelation function plot (ACF) lying within the 95% bounds, and ${Z_t}^2$ appears to be $AR(1)$, then an $ARCH(1)$ model for the variance is suggested (Engle, 2001). If the PACF of the ${Z_t}^2$ suggests $AR(m)$, then $ARCH(m)$ might work. GARCH models may be suggested by an $ARMA$ type look to the ACF and PACF of ${Z_t}^2$ (Tsay, 2012). In practice, this rule of thumb will not always work as correctly, and a little experimentation with various ARCH and GARCH structures might need to be tested.

Another common methodology is the use of software to estimate the best set of parameters. Luckily, the software R allows us to rank a different set of GARCH parameters (m, s) based on the values of the AIC. The procedure would be the following:

- The optimal set of ARMA orders would be based on the values of p and q for the ideal $ARMA(p,q)$ for the time series, which were already determined in Chapter 4. These values would be combined with values $(1,1)$ and $(2,2)$, just in case the optimal values do not coincide with these values.
- Usually, the orders of the GARCH model take on the values of 1 and 2 or any combination of these numbers.
- Each combination of $ARMA(p,q)$ and $GARCH(m,s)$ should be tested and ranked according to its AIC. The ARMA-GARCH model with the lowest AIC should be picked.
- The best set of parameters of $ARMA(p,q)$ and $GARCH(m,s)$ will be fitted to the time series.

From Chapter 4, we determined that the ideal order values for the AMRA model were $(2,2)$. In order to determine the lowest AIC for ARMA-GARCH model the set of orders tested (p,q) (m,s) was the following: $(1,1)$ $(1,1)$, $(1,1)$ $(2,2)$, $(2,2)$ $(1,1)$, and $(2,2)$ $(2,2)$.

Table 5.1. shows a summary of the AIC for a different set of parameters for the ARMA-GARCH model. As can be seen from this table, the set of orders $ARMA(2,2) - GARCH(1,1)$ provides the lowest AIC. Though these values coincide with what most literature suggests, it should not be granted that the best set of parameters for ARMA-GARCH would coincide with the best ARMA model for the time series. It is always good practice to analyze a different array of orders. These values will then be used to forecast the nickel prices 24 months ahead. The results from this forecast will be shown in the next subsection.

Table 5.1: AIC for a different set of parameters for ARMA-GARCH

Model	Parameters	Akaike
ARMA	(1,1)	17.063
GARCH	(1,1)	
ARMA	(1,1)	17.109
GARCH	(2,2)	
ARMA	**(2,2)**	**17.037**
GARCH	**(1,1)**	
ARMA	(2,2)	17.082
GARCH	(2,2)	

5.4.2. Future prices

Using the optimal set of parameters, the nickel prices for the next 24 months starting January 2018 was forecasted. Unlike the ARIMA process, the input variable for the GARCH model was the log of the return of nickel prices. Therefore, to calculate the nickel prices back, the reverse of equation 5.4 would be used.

Let P_t be the nickel price at time t and $Z_t = \log P_t - \log P_{t-1}$ the log returns. If Z_{t+1} is the forecasting log return one unit of time ahead (same as time interval in between two consecutive observations of the time series), then the formula for converting the forecasting log return to the original price is

$$P_{t+1} = 10^{Z_{t+1} + \log P_t} \tag{5.4}$$

Figure 5.1 shows the log-returns of the original time series as well as the forecasted log returns (Z_{t+1}) 24 months ahead of starting January 2018. These values were converted (back) to the original nickel prices using Equation 5.4. For example, for January 2018, the value of the nickel price is:

$$P_{January-2018} = 10^{Z_{January-2018} + \log P_{December-2017}} \quad (5.5)$$

Since the values of Z_{t+1} are already known for the next 24 months, the value of $P_{January-2018}$ can be determined. This value is then used to calculate the nickel price for February 2018.

$$P_{February-2018} = 10^{Z_{February-2018} + \log P_{January-2018}} \quad (5.6)$$

Analogically, the rest of the values can be calculated following the same procedure.

The results from such conversion are shown in Table 5.2. Table 5.2 summarizes the converted values of the forecasted log returns for the next 24 months from January 2018 to December 2019.

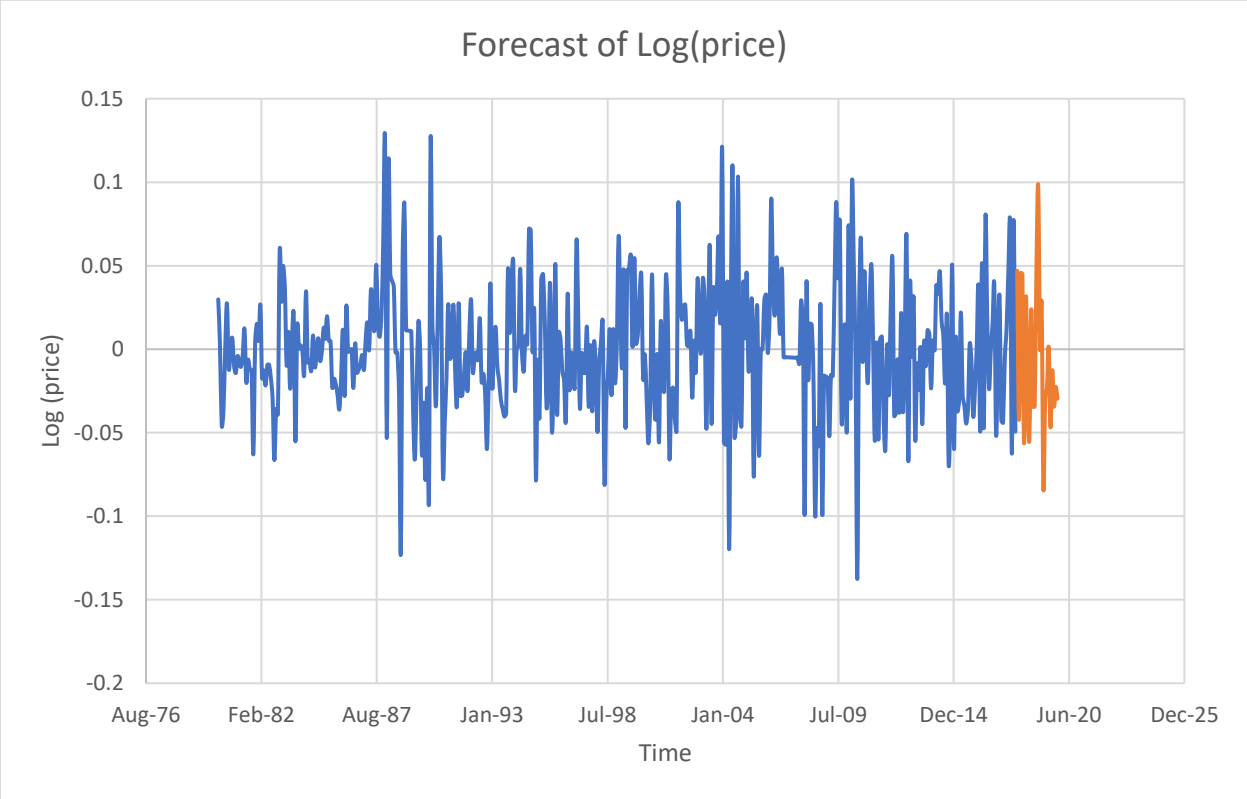

Figure 5.1: Forecasted log-returns 24 months ahead starting January 2018

Table 5.2: Summary of future values of Nickel prices using an GARCH model

Date	Estimate ($/t)	Date	Estimate ($/t)	Date	Estimate ($/t)
Jan-18	12,542.894	**Sep-18**	11,567.80	**May-19**	11,365.03
Feb-18	11,378.765	**Oct-18**	10,690.40	**Jun-19**	10,773.84
Mar-18	12,634.961	**Nov-18**	9,880.60	**Jul-19**	10,801.73
Apr-18	14,012.372	**Dec-18**	11,040.53	**Aug-19**	9,697.19
May-18	12,307.016	**Jan-19**	13,830.08	**Sep-19**	9,420.43
Jun-18	13,237.602	**Feb-19**	13,860.16	**Oct-19**	8,705.79
Jul-18	12,392.029	**Mar-19**	14,758.49	**Nov-19**	8,262.94
Aug-18	10,947.678	**Apr-19**	12,183.24	**Dec-19**	7,718.88

Chapter 6

Forecasting Through Artificial Neural Networks (ANN) Models

6.1. Introduction

One of the most critical issues many mining companies face when evaluating a mining project or determining the financial liquidity of current operations is the price prediction of the mineral(s) they exploit or plan to. In fact, due to the sensibility of mineral prices to external market forces, its importance in the decision-making process in the mining industry is of major concern.

Despite the drawbacks that time-series forecasting models such as ARIMA and GARCH have, mining companies and traders still rely on them for mineral price forecasting. Though these models help have a clearer idea of the future path of the mineral prices, recent research has proved that, in the short run, both ARIMA and GARCH forecasting models tend to be outperformed by current machine learning techniques (Labys, 2006)

In the last decade, ANNs have been extensively studied as an alternative to popular linear and non-linear time series forecasting models. In fact, due to the networks' non-linear modeling capability, it is believed that ANNs are able to better resemble the complex real world with a high degree of accuracy (Chen et al., 2003).

The main idea behind an artificial neural network architecture is based on how the different networks of neurons in the human brain work when processing information. Their ability to work

parallelly and interconnected with the aim of delivering certain function(s) is what makes this particular method a strong forecasting technique (Zhang, 2003).

In this research, an ANN model will be developed and trained to forecast the future prices of nickel in the short term. Their forecasts will be validated against not only the real or observed values but also the forecasts of other models, aiming to assess the quality of this forecasting model over the most common and popular time-series forecasting methods.

6.2. The Artificial Neural Network (ANN) approach

An ANN model can be thought to work the same way the intricate network of neurons in the human brain does when processing information. That is, neural networks allow to establish different complicated relationships across and among the set of values of any data. In fact, artificial neural networks are machine-learning forecasting techniques that allow rather complex non-linear relationships between the forecast(s) and its predictors. There is where their advantage over time-series linear techniques lies (Zhang, 2003).

The architecture of neural network resembles a network of neurons distributed in different layers; its architecture revolves around one single element, also called a *single neuron*, which is the base unit of a neural network. In the simplest design of an artificial neural network, also known as *single layer perceptron*, there are only two layers: the *input layer*, where the external information is received, and the *output* or *forecast layer*, where the problem solution is obtained (Panella et al., 2012). The neurons in the input layer are also known as *inputs* or *predictors*, whereas those in the output layer are also called *outputs* or *forecasts*. In this type of set-up a single neuron receives input and computes an output. Each input comes with an associated weight (w_j), which is assigned on the basis of its relative importance to other inputs (Panella et al., 2012). In

the output layer, a weighted linear combination function is applied to transform the input(s) into a final output.

The set-up from Figure 6.1 takes the input values X_1, X_2 and X_3, and assigns them weights w_1, w_2 and w_3. These values along with the pre-established input 1 with weight b, also known as bias, are jointed together simulating a weighted linear combination. As can also be seen in Figure 6.1., the only member in the linear combination of the inputs that does not depend on the input values themselves is the bias. This term is extremely important in the linear combination set-up since it provides every neuron with a trainable constant value (Panella et al., 2012).

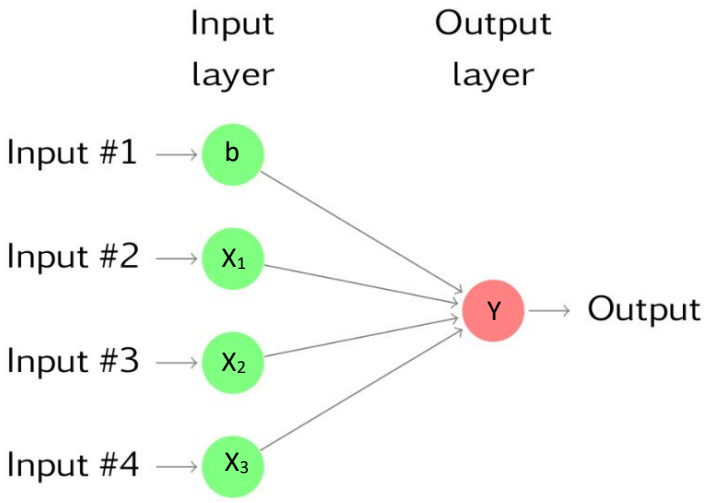

Figure 6.1. Architecture of a single neuron with three input neurons and a bias

Unlike the simplest cases of an artificial neural network, in which its architecture can be thought as that of a linear regression, the most complex cases contain both an intermediate layer known as hidden layer and hidden neurons (Mombein and Yazdani-Chamzini, 2015). Since the double-layer neural network or single neuron works exactly as either a linear or multilinear linear regression depending on the number of inputs, the nonlinearity feature mentioned at the very

beginning of this chapter is not yet accomplished. It is, in fact, the *hidden layer* the one that accounts for the non-linear advantage of artificial neural networks and makes them a much more representative and accurate forecasting technique (Mombein and Yazdani-Chamzini, 2015).

This multilayer network , also known as *multilayer feed-forward network,* has an interesting architecture in which each one of the neurons of the input layer acts as the inputs to the *hidden layer* and undergoes a *weighted linear transformation* (Kristjanpoller and Minutolo, 2015). These outputs become the inputs to the final layer, and the final output is the result of a non-linear modification of the hidden layer outputs. For instance, the input into a hidden neuron "j" are combined to give:

$$Z_j = a_j + \sum_{i=1}^{n} w_j Y_j \tag{6.1}$$

Where a_j is the bias factor and Z_j is the linear combination of all the input neurons (Y_j) and their respective weights (w_j) to the hidden layer.

The parameter a_j and w_j are actually learned from the data after a series of simulations undertook by the network using different starting points randomly selected (Tsay, 2012). That is, firstly, both values are taken randomly, and are constantly updated using the observed data values and the predicted ones, which are added to the observed data as soon as they are estimated. These coefficients are usually restricted to the so-called decay parameter, which is set to be equal to 0.1. The network is usually trained several times using different random starting points, and the results are averaged (Tsay, 2012).

In this multilayer network, unlike other networks such as the recurrent neural networks, the information moves only forward and there is no connection between the neurons of different layers. That is, once the input values from a layer, say the input layer, move forward to the next layer, they do not turn back to create any kind of cycle (Kristjanpoller and Minutolo, 2015).

Figure 6.2. illustrates the design of a simple neural network with an intermediate or hidden layer. Each one of the input dots follows a path forward that always yields one last value.

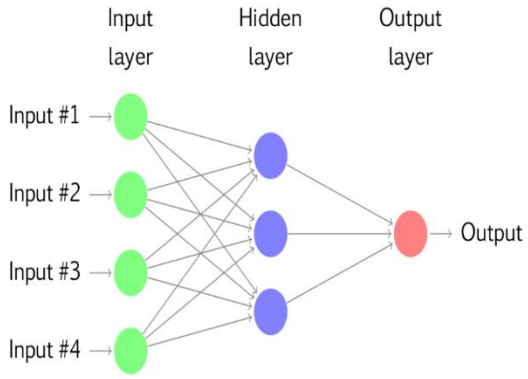

Figure 6.2. Architecture of a multilayer network with a hidden layer consisted of 3 hidden neurons, 4 input neurons and an output layer

The outputs of the hidden layer, now the inputs of the output layer, are transformed to a final result using a non-linear function called the *activation function***.** This activation function is the sole responsible for providing the non-linearity feature to the model (Wongsathan and Seedadan, 2016). Given that real-world time series are non-linear, the activation function plays a crucial role in the setup.

All activation functions, no matter what the formula they take, take a single input value and transforms it into a real number (Panella et al., 2012). There are several activation functions; the most common ones are:

- Sigmoid: this function takes a value and, via Equation 6.2, transforms it into a positive real number between 0 and 1. Figure 6.3 shows that the range for the sigmoid function is between 0 and 1, no matter what the value in the domain is.

$$\sigma(x) = 1 / (1 + \exp(-x)) \tag{6.2}$$

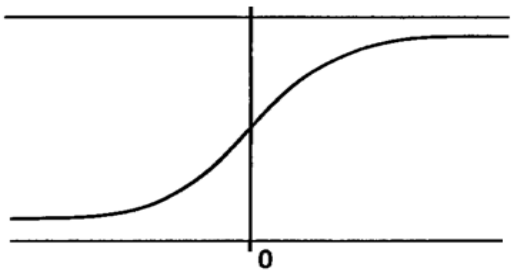

Figure 6.3: Sigmoid function

- Tanh: this function takes a value and, through Equation 6.3, transforms it into a real number ranging between -1 and 1, inclusive. Figure 6.4 shows that the tanh function can take a maximum value of 1 and a minimum value of -1 or any value in between that range.

$$\tanh(x) = (\exp(x) - \exp(-x))/ (\exp(x) + \exp(-x)) \tag{6.3}$$

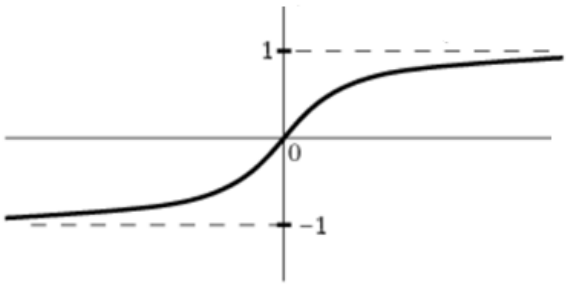

Figure 6.4: Tanh function

- ReLU (Rectified Linear Unit): This function takes a value and compares it with zero (as can be seen from Figure 6.4), yielding the maximum between both numbers. That is, it replaces the negative values (or zero) with zero and the positive values with the same positive values. Figure 6.5 shows that the ReLU takes on zero for all negative values, whereas it resembles the identity function for positive input values.

$$f(x) = \max(0, x) \tag{6.4}$$

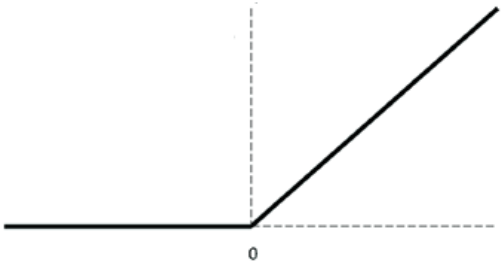

Figure 6.5: ReLU function

6.3. Case study: Nickel price forecasting

6.3.1. Artificial Neural Network (ANN) forecasting model

In the ANN architecture for time series, the past values or rather their lagged values are computed as inputs and the future value is the output (Khashei and Bijari, 2010). That is, if y_t is the observation at time t, then

$$y_{t+1} = f(y_t, y_{t-1}, y_{t-2}, \dots, y_{t-n}) \tag{6.5}$$

y is the values of the forecasted value one unit of time ahead, be it a month, year, day, etc., and n indicates the number of past observations the network needs to conduct proper training and

subsequent re-trainings. In fact, the network follows the same training procedure described in the previous section and, once the first forecasted values is calculated, it is immediately added to the training patterns the network will use for the second-step-ahead estimated value (Khashei and Bijari, 2010). This process is iterative and continues for as long as number of time units ahead the networks is provided with. That is, if we follow the same example cited above, the forecasted values two-step ahead, three-step ahead, and so on will be:

$$y_{t+2} = f(y_{t+1}, y_t, y_{t-1}, \cdots, y_{t-n+1}),$$

$$y_{t+3} = f(y_{t+2}, y_{t+1}, y_t, \cdots, y_{t-n+2}),$$

$$y_{t+4} = f(y_{t+3}, y_{t+2}, y_{t+1}, \cdots, y_{t-n+3}),$$

$$\vdots$$

$$y_{t+k} = f(y_{t+k-1}, y_{t+k-2}, y_{t+k-3}, \cdots, y_{t+1}, y_t, \cdots, y_{t-n+k-1})$$

The generalized formula that the neural network will follow is:

$$Y_t = w_0 + \sum_{i=1}^{q} w_j . g(w_{0,j} + \sum_{i=1}^{q} w_{i,j} . Y_{t-i}) + \varepsilon_t \qquad (6.6)$$

where Y_{t-i} is the subset of the lag values of the time series to be used by the network, $w_{i,j}$ are the arc weight coefficients feeding the hidden layer, $g(x)$ is the activation function, and w_j are the coefficients of the hidden layer outputs aiming to feed the forecast layer.

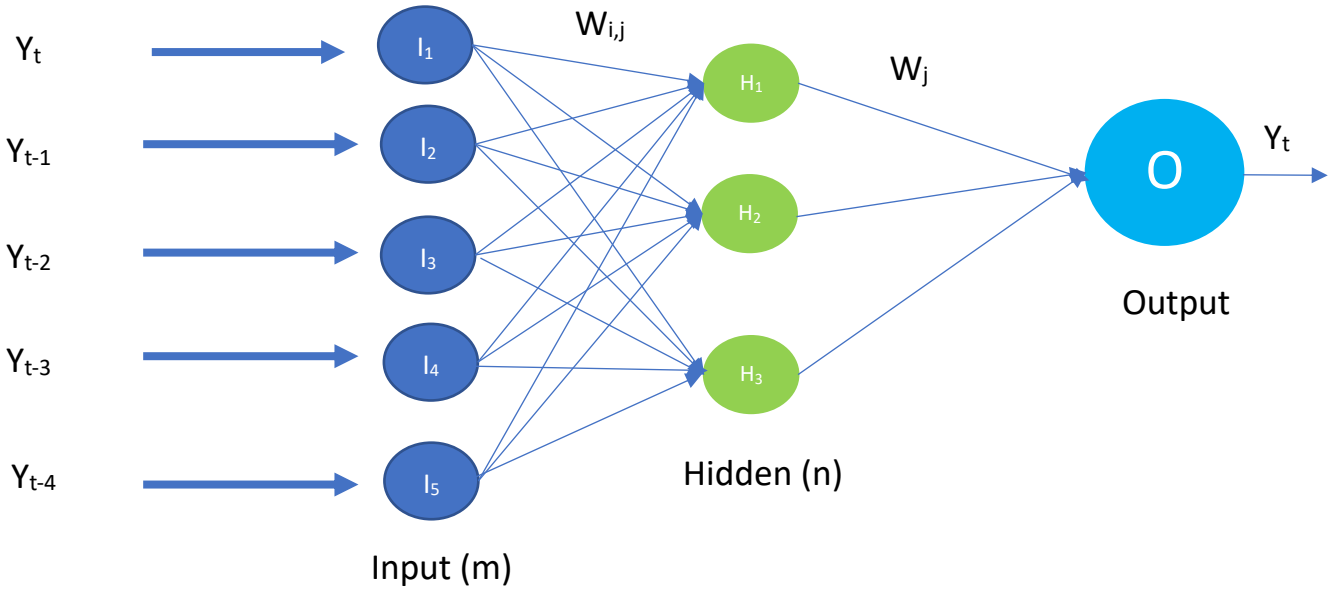

Figure 6.6: Architecture of a multilayer feed-forward neural network with 5 neurons and 3 neurons in the input and hidden layers, respectably

For the purposes of this study, the sigmoid function will be used to transform the hidden layer outputs to the output layer.

$$g(x) = \frac{1}{1+e^{-x}} \tag{6.7}$$

Another way to look at neural networks and one that actually helps to understand what needs to be determined in order to estimate forecasts is parameter-based notation. Since an artificial neural network is equivalent to a non-linear autoregressive model for time series forecasting, its typical notation is *NNAR(p, P, k)* (Berlinger et al., 2015). The parameter p indicates the number of lagged inputs (from the data set) that the neural network used to re-train itself, whereas k indicates the number of nodes in the hidden layer and P responds to the number of differences needed for the time series to achieve stationary. For instance, NNAR (9,1,5) would indicate a neural network that uses the last nine observations of the data for modelling.

Although many different approaches exist in order to find the optimal architecture of an ANN, such as the pruning algorithm, polynomial time algorithm, the canonical decomposition techniques, etc., these methods are usually quite complex and difficult to implement (Zhang et al., 1998). Furthermore, none of these methods can guarantee the optimal solution for all real forecasting problems. To date, there is no simple clear-cut method for the determination of these parameters; however, there are a couple of research-based methods that have been proven to be most effective across most forecasting literature.

For the purpose of this study, the software R will be used to avoid tricky and complex calculations; however, the yielded optimal parameters will be contrasted against the rule of thumb laid out by Hyndman and Athanasopoulos (2018). Both authors proposed that the optimal number of lags for the network are the same optimal number of lags (according to the AIC) for a linear AR(p) model. This value can be therefore obtained by inspecting the partial autocorrelation diagram of the time series.

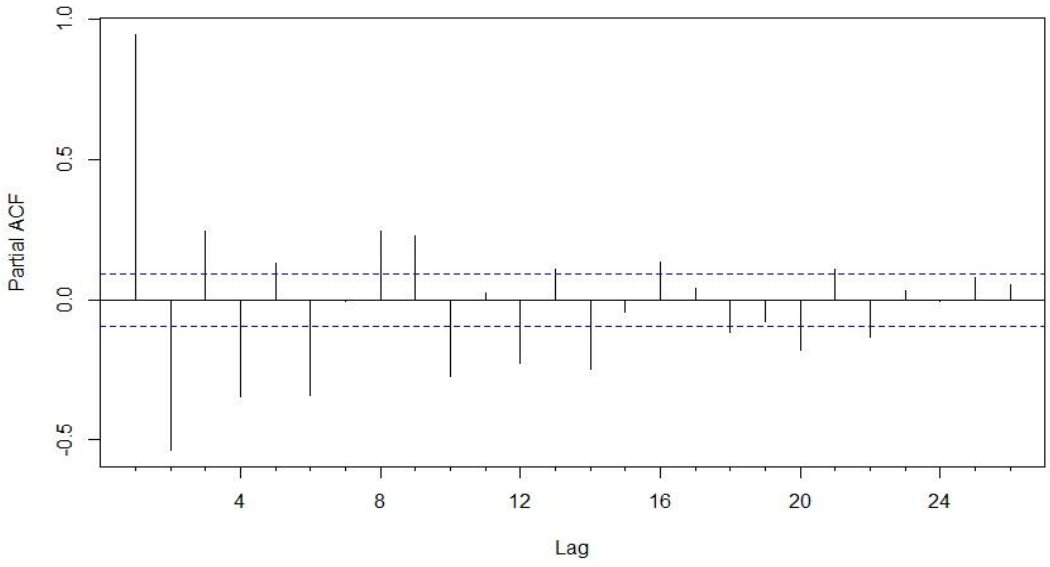

Figure 6.7: PACF of the time series for nickel prices.

Figure 6.7 shows the partial-autocorrelation function at different lags of the time series. After analyzing the partial autocorrelation correlogram of the time series, we can clearly see that the graph cuts off at lag 13, indicating that the optimal value of p for the neural network would be 13.

Furthermore, Hyndman and Athanasopoulos (2018) ascertained that the default value of P is set to be 1 and the values of k can be calculated with the following formula:

$$k = (p + P + 1)/2$$

(rounded down to the nearest integer)

The value of P (equal to 1) actually coincides with what the Dickey-Fuller test (stationary test) yielded (Chapter 04).

If we plug the value of p into the formula to calculate k, we obtain (rounding down to the nearest integer):

$$k = (13 + 1 + 1)/2$$

$$k = 7$$

Therefore, the optimal set of parameters for the neural network would be (13,1,7).

In R (software), the function $nnetar()$ helps determine the optimal set of parameters. Coincidentally, the results for the set of parameters of the networks is the same as that determined by the rule of thumb stated above.

6.3.2. Future prices

As previously stated, the optimal set of parameters for the nickel price time series is (13,1,7). Nickel prices for the next 24 months will be forecasted. However, in order to develop a more comprehensive knowledge of the future values of nickel price, a series of 1,000 simulations will be performed. That way, the distribution of the future path will be better understood and more defined. The overall distribution of the forecasted nickel price will be the average of the set of simulations.

This simulation process also helps determine the extreme future sample paths the forecasted nickel prices might take, establishing what is known as prediction intervals. These intervals describe the range of the set of paths the future nickel prices might fall into.

In Figure 6.8, it can be seen that the simulation of eight possible sample paths (out of the 1,000 simulations performed) for the nickel price forecasts for the next 24 months. These simulations were randomly selected just to show how their path behaves over time. Clearly, it can also be seen that ANNs are able to model cyclicity that the ARIMA model in Chapter 4 was unable to. It can also capture the asymmetry that models such as ARIMA (Chapter 4) and GARCH (Chapter 5) failed to.

Figure 6.9, on the other hand, shows the average of the 1,000 simulations performed. Although the software R allows you to conduct as many simulations as wanted to train the neural network, Hyndman and Athanasopoulos (2018) stated that running simulations a few hundred or thousand times is enough to have a clear understanding of what the future path of the forecasted values will be. In fact, carrying out more simulations might only slow down the calculation process and provide no significant changes. This remark can be clearly seen in Figure 6.10.

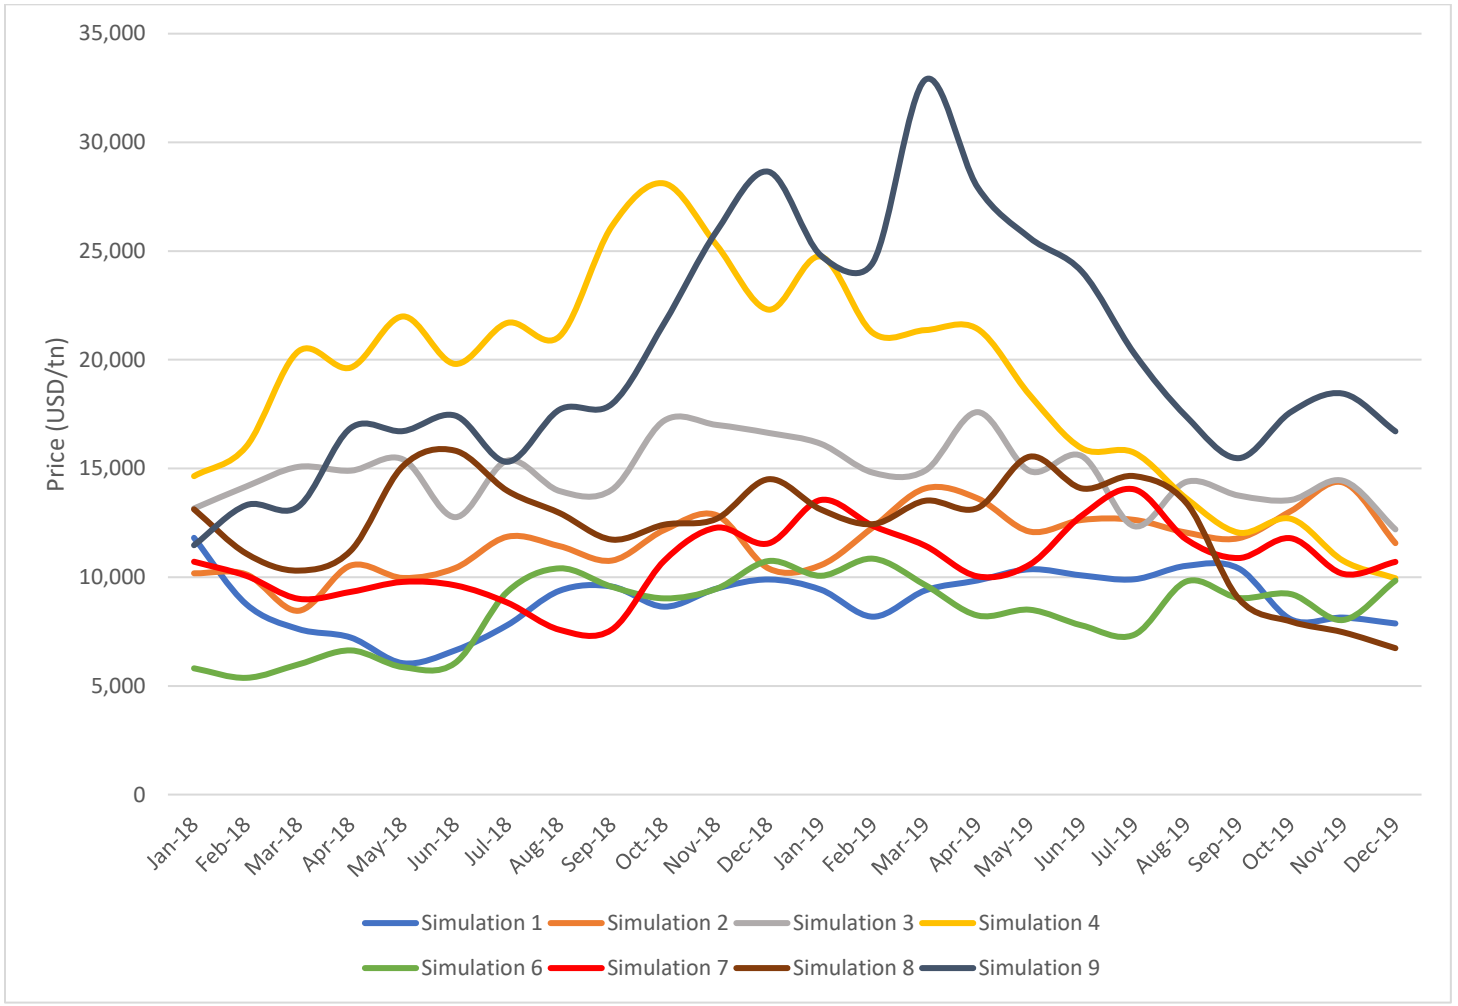

Figure 6.8: Simulations of the future sample paths for forecasted nickel prices for the next 24 months

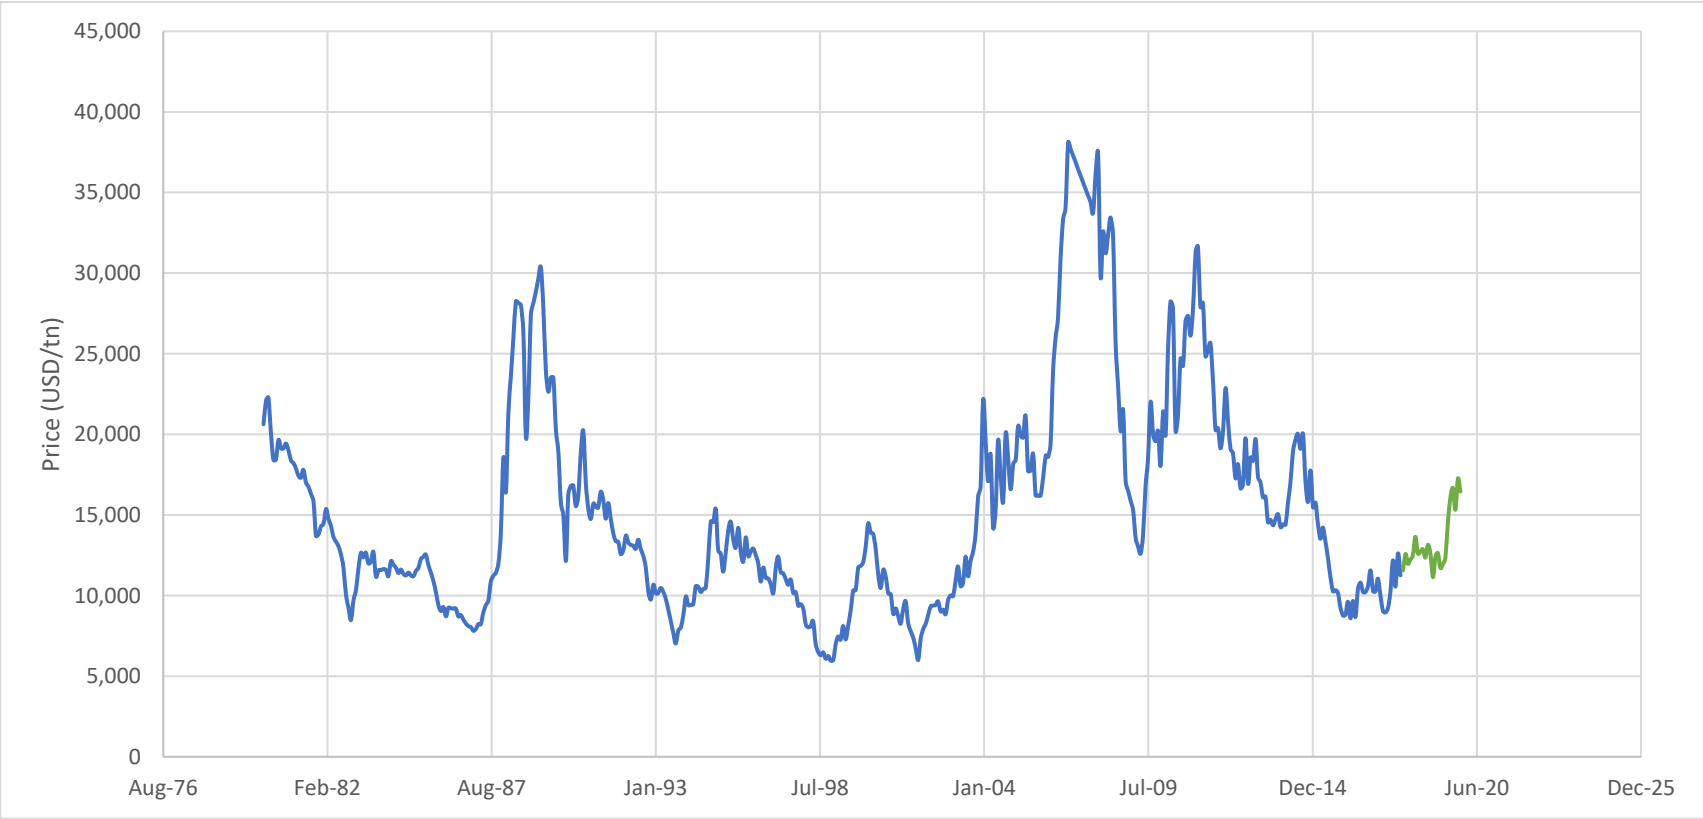

Figure 6.9: Average of the 1000 simulations for forecasted nickel prices for the next 24 months (green line).

Figure 6.10: Average nickel price future paths simulated iteratively (100, 500, 900 & 1,000 simulations)

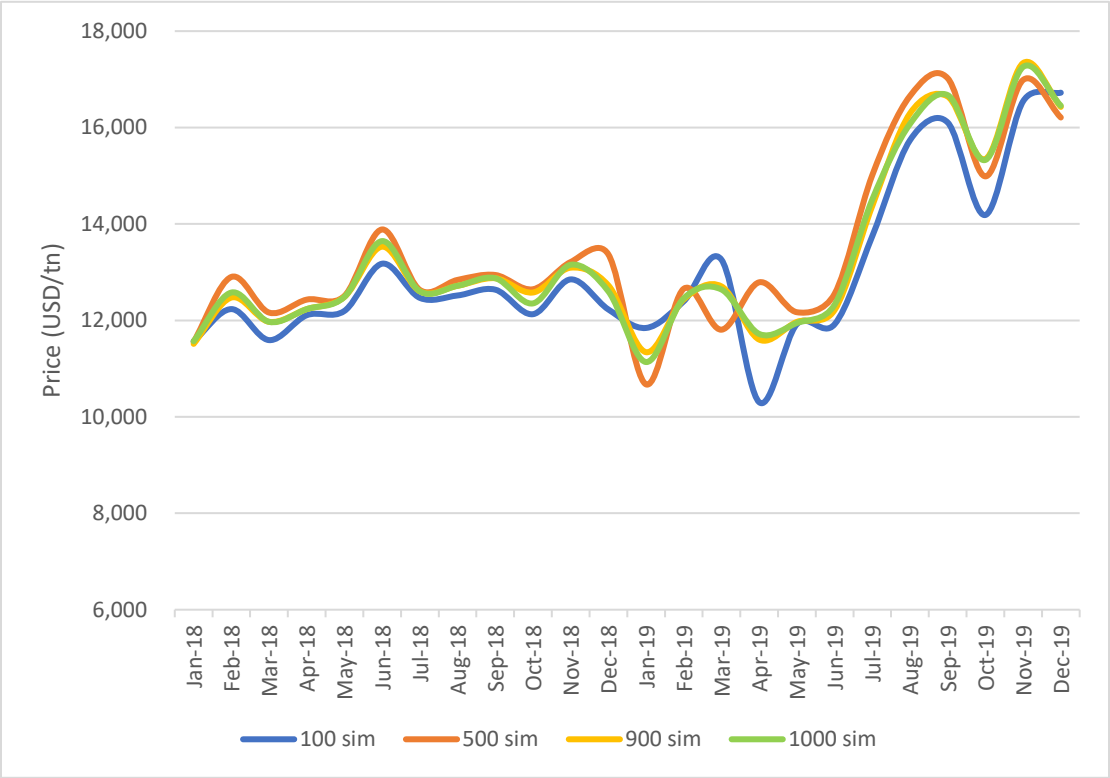

Figure 6.10 shows that there is almost no difference between the future path for the forecasted values ran 900 and 1,000 times. That is why the chosen threshold was 1,000 simulations for the sake of accuracy and speed.

Table 6.1. provides the values the forecasted nickel prices using Artificial Neural Networks (ANN) training from January 2018 and December 2019.

Table 6.1: Summary of future values of Nickel prices using an ANN model

Date	Estimate (USD/tn)	Date	Estimate (USD/tn)	Date	Estimate (USD/tn)
Jan-18	11,453.64	**Sep-18**	12,787.38	**May-19**	10,442.38
Feb-18	11,839.57	**Oct-18**	12,370.79	**Jun-19**	10,122.14
Mar-18	11,785.11	**Nov-18**	12,366.88	**Jul-19**	9,780.74
Apr-18	12,505.11	**Dec-18**	12,024.75	**Aug-19**	9,393.29
May-18	12,083.58	**Jan-19**	11,703.19	**Sep-19**	9,078.55
Jun-18	12,554.33	**Feb-19**	11,519.98	**Oct-19**	8,729.99
Jul-18	12,685.72	**Mar-19**	11,046.09	**Nov-19**	8,423.09
Aug-18	12,544.30	**Apr-19**	10,840.76	**Dec-19**	8,135.52

Figure 6.11, in turn, provides a picture of what the most likely future paths for the forecasting of nickel prices in the next 2 years will look like. It shows the two 'extreme' values the simulations (for the forecasted nickel price) can take as well as the average of the values of the 1,000 simulations. The other simulations lie within those two 'extremes.'

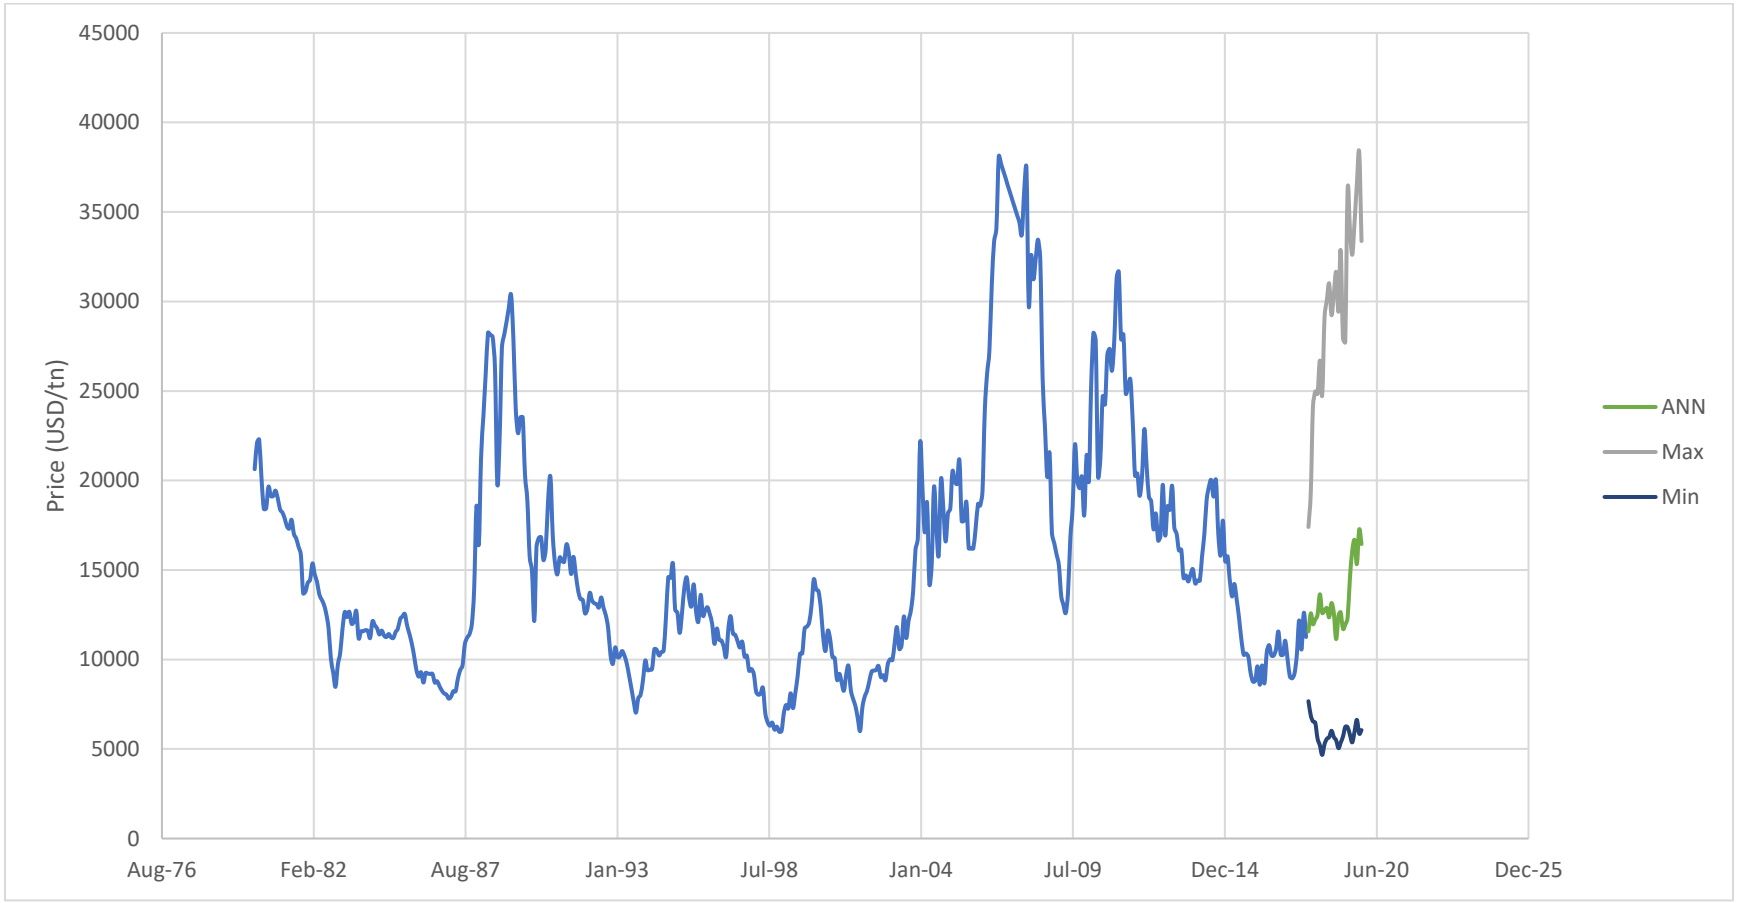

Figure 6.11: Prediction intervals (maximum and minimum) for forecasted nickel prices for the next 24 months

Chapter 7

Discussion and Conclusions

7.1. Performance assessment of the different models

The observed or real monthly nickel prices from January 2018 to December 2019 are used as reference benchmarks to assess the performance of a) the ARIMA model over time, b) GARCH forecasting model over the ARIMA process, and c) the ANN model in comparison to the ARIMA and GARCH models. In order to grasp how well those models behave over one another over time, a comparison at different points in time will be carried out.

The criteria to analyze the performance of the forecasting models will be the use of different performance measures: RMSE, MAE, and MAPE.

Table 7.1 summarizes all the error measures for different forecasting timeframes for the ARIMA, GARCH, and ANN processes. The decision to include values of performance measures at different time periods was based on the idea that both the ARIMA, GARCH, and ANN forecasting models are more suitable for short-term forecasting. That is, it is possible the accuracy of those models changes over time (either increases or decreases or both) and that the relationship between the two might change as well, implying that at point of time the accuracy of a model may be outperformed by the other that used to be a more accurate model.

Figure 7.1 through 7.3 shows a graphical comparison between forecasts 24 months ahead using ARIMA, GARCH, and ANN processes.

Figures 7.1, 7.2, and 7.3 show the trend of the error measures MAE, MAPE, and RMSE over different time frames (6, 12, 18, and 24 months). Both lines in each one of the graphs will be compared against an imaginary line located right on the horizontal axis (y = 0). That is, the more the line trend deviates from the horizontal axis, the less accurate the forecasting model is. This procedure will help determine an empirical cut-off at which the relationship between different forecasting models changes.

Figure 7.4 shows that the ARIMA model predicts an overall increase of the nickel price for the next 24 months, whereas the GARCH model estimates a dramatic decline of the price after month 15 (March 2019). The forecasts using the ANN model, on the other hand, based on a quick glance, do seem to be able to resemble the overall behavior of the real prices' time series (both in terms of the overall trend and fluctuations).

7.1.1. Performance assessment of the ARIMA model

As can be seen in Figure 7.4, the predicted values using an ARIMA model resembles, for the most part, the overall trend of the real nickel price values. Nonetheless, the ARIMA-based forecast is unable to predict the upward and downward fluctuations the real values show.

The forecast also shows the linearity problem many researchers point out regarding ARIMA forecasting. The forecasted values go up across the timeframe given, providing not-so-real nickel price values. However, if a careful look were taken at the predicted values after month 18, it can be shown that, due to ARIMA models' linearity, the model is able to estimate, though not precisely and only in terms of the trend, the boom in prices due to the Indonesian ban on nickel exports.

From Figure 7.1, 7.2, and 7.3, it can be seen that in months 6 and 12 both the error measure lines for the ARIMA models deviate more from the horizontal axis. Due to the linearity feature of the ARIMA models, it is expected that during the first 6 to 12 months, the performance of the model is somewhat larger, especially if the time series have shown signs of extreme volatility (ups and downs). This seems to be the case for the ARIMA model presented in this thesis. Over time, and due to the volatility of mineral commodity prices, the model's forecasting accuracy tends to stabilize since many of the larger forecasting values tend to be offset by the drops in mineral commodity prices as part of their natural price cycle.

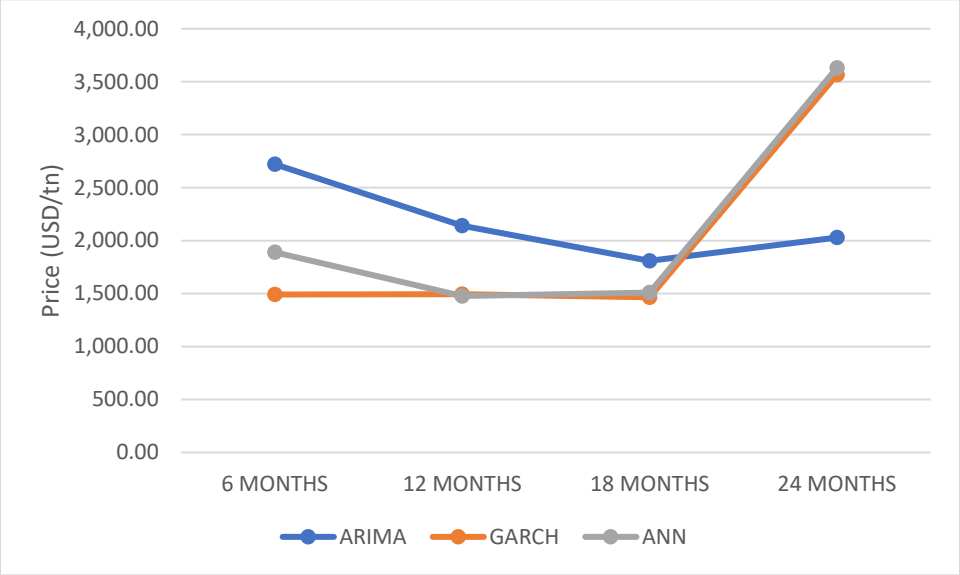

Figure 7.1: Performance of ARIMA, GARCH and ANN forecasting models over time measured via RMSE. Price of nickel (USD/tn) vs time (months)

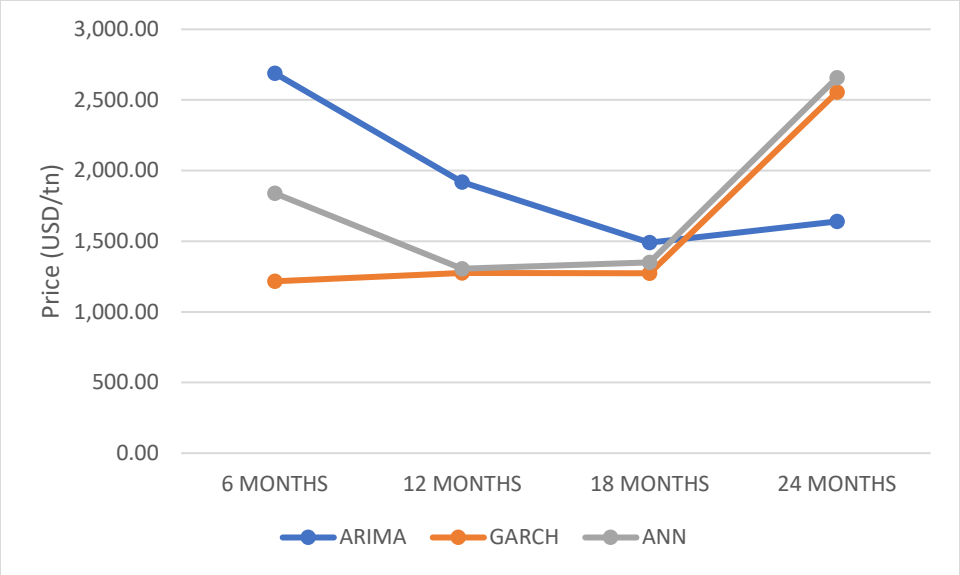

Figure 7.2: Performance of ARIMA, GARCH and ANN forecasting models over time measured via MAE. Price of nickel (USD/tn) vs time (months)

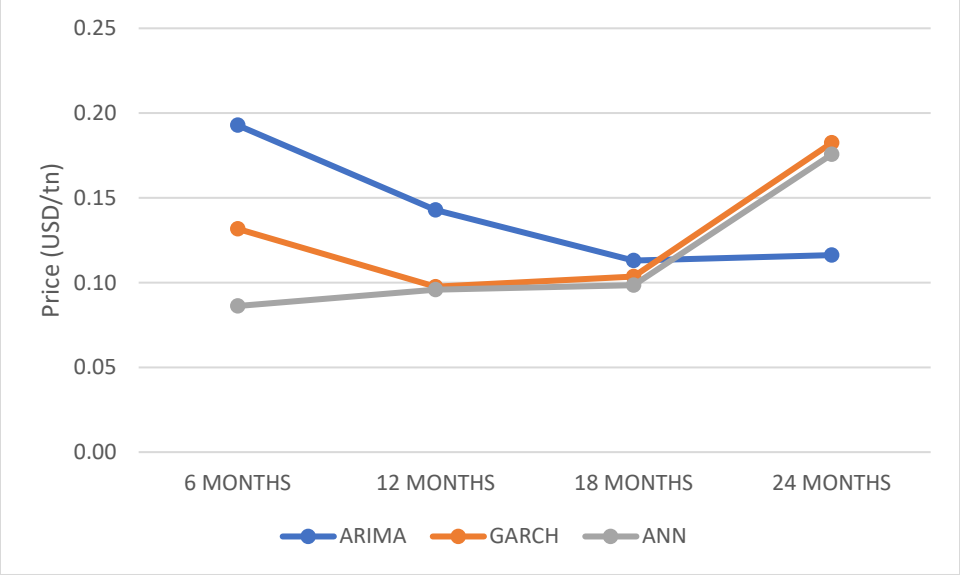

Figure 7.3: Performance of ARIMA, GARCH and ANN forecasting models over time measured via MAPE. Price of nickel (USD/tn) vs time (months)

These figures also show that, overall, forecasts using ARIMA models yield better results over the whole 24-month horizon, when compared to the other models. However, when the trend of each of the error measures is compared at different times, the results vary: the cut-off for optimal precision for the ARIMA model seems to be 18 months. Nonetheless, this value should be taken with a grain of salt and only be understood in the context of what it represents in relation to the other models' cut-off. The accuracy of ARIMA-based forecasting is bell-shaped. From months 6 to 18, the accuracy of the forecasts using ARIMA models increases (comes close to the horizontal axis) and then, decreases between forecasting month 18 and month 24, suggesting that ARIMA forecasting models work best in the short run (as explained before, the definition of short run can seem to be broad; nonetheless this thesis is considering it as any value that is less than 24 months). Drawing a quantitative answer of what constitutes an optimal cut-off would be subjective as many studies might need to be conducted in order to determine the said threshold.

Table 7.1: Error measures for different models and timeframes (N, in months)

Model	**ARIMA**	**GARCH**	**ANNs**
N (months)	6	6	6
RMSE	2,720.85	1,490.51	1,890.07
MAE	2,688.89	1,216.41	1,840.36
MAPE	0.1928	0.1316	0.0862
N (months)	12	12	12
RMSE	2,141.60	1,493.91	1,477.04
MAE	1,918.19	1,274.63	1,305.59
MAPE	0.1429	0.0976	0.0958
N (months)	18	18	18
RMSE	1,809.29	1,464.87	1,509.60
MAE	1,490.03	1,273.37	1,351.13
MAPE	0.1130	0.1036	0.0985
N (months)	24	24	24
RMSE	2,029.68	3,566.34	3,630.33
MAE	1,639.49	2,554.71	2,657.44
MAPE	0.1163	0.1824	0.1756

Based on Table 7.1, it can be seen that for almost all the scenarios (different timeframes) the accuracy of the forecasts using the ARIMA model is the worse among those of the other two forecasting models. It is only when considering a time span of 2 years for the forecasted values that the accuracy of the model improves. This behavior can be fairly explained by the statements made in the paragraphs above.

7.1.2. Performance assessment of the GARCH model

As can be seen from Figure 7.4, the predicted values using a GARCH model resemble the overall trend of the real nickel price values for the first fifteen months (January 2018 to March 2019). Nonetheless, the forecasts using the GARCH process are unable to replicate the extreme increase jump the real values show from June 2019 to September 2019. This behavior contrasts with the future path of the ARIMA-based forecast, which surprisingly is able to predict, though not precisely, the trend the nickel prices' path takes from May 2019 forward. Although the GARCH forecasting behavior might seem odd since the GARCH process models the mean (ARMA) and volatility of the time series, the downward trend can be attributed to the high volatility of the forecasted series show.

From Figures 7.1 to 7.3, it can be seen that, unlike the ARIMA forecasting technique, the shape of the accuracy curve of the GRACH model is not definitive. What is clear, however, is that the accuracy of this model does tend to be quite stable up until month 18. This finding is clearly stated in Table 7.1. In that we can see that the forecasts using this model tend to have, in almost all case, the second lowest error measures, losing only to the ANN forecasting method and indicating that the GARCH forecasting technique represents a much better approach than the ARIMA model.

In addition, the accuracy of GARCH forecasting models seems to decrease, though slightly until month 18, as time passes. This outcome could indicate and confirm that the GARCH forecasting technique, just like the ARIMA model, works best in the short run and looses its efficiency when the forecasting time span is too wide.

If the comparison to the ARIMA forecasting method were to be made for the complete forecasting horizon (24 months), the accuracy of the GARCH model would pale in comparison to that of the linear model. However, this finding should be considered carefully since the surge in prices that the ARIMA forecasting method is able to suggest can be attributed to a unique external factor that is not reflected in the past behavior of the time series.

It can be argued that, even though the GARCH forecasting method failed to precisely portrait the trend of the time series after month 18, the sudden and dramatic decline in the forecasted prices from May 2019 on might seem like a bit of an exaggeration . Nevertheless, in this special case, this finding seems to make sense somehow (from a quantitative perspective) since the time series of nickel prices exhibit great periods of high volatility, a factor that is taken into account in the calculation of the forecasted values using this model.

7.1.3. Performance assessment of the ANN model

As can be seen from Figure 7.4, the ANN model is the one that seems to yield better results in the nickel price forecasting. Both the trend and fluctuations of the time series of the real nickel prices seem to be accurately estimated by this method.

Table 7.1 confirms this finding. For all the forecasting timeframes up until 18 months, the error measures of the forecasting method is one of the lowest among all, only rivalling the ones from the GARCH model. For instance, based on Figures 7.1 through 7.3, if the RMSE and the MAPE

were to be chosen as the determinants of which model is most suitable, the ANN one will win by a slight margin over the GARCH method. Conversely, the opposite would happen when considering the MAE as the main error measure. However, when the error measures for the 24-month forecast are compared, the results vary: among the three forecasting models this thesis reviewed and analyzed, the artificial neural network forecasting model is no longer the most suitable one, whereas the ARIMA forecasting model is the most appropriate. This occurrence might be attributed to the fact that ANNs are better suitable for shorter forecasting timeframes (than those of the other two models). In fact, as the ANN forecast moves forward, more and more past observations are dropped, relying instead on forecast to predict further future values. Therefore, the longer the forecasting horizon, the less reliable and accurate the ANN model.

The results explained above (sections 7.1.1 to 7.1.3) seem to agree with some of the findings in the body of literature in mineral commodity price forecasting. That is, forecasting methods, be they ARIMA, GARCH, or ANN, tend to be more accurate in the short run (somewhat between a few months and up to two years). This can be seen in Figures 7.1, 7.2 and 7.3, being the ARIMA-based forecast somewhat of an outlier due to external factors affecting the price in 2019 (Indonesian ban).

Furthermore, as expected, the results show that the GARCH forecasting model does represent an improvement over the traditional linear forecasting models such as ARIMA.

Figure 7.4: Summary of different forecasting models for nickel price. The timeframe of the forecasts is 24 months

7.2. Conclusions and Future Work

Mineral commodity prices have been the most important factor to consider when assessing the economic viability of a mining project or any strategic decision for that matter. And yet, the estimation of future mineral commodity prices remains ever challenging and a matter of simplicity and accuracy. In fact, conventional forecasting methods, though as simple as they can get, have been proven not to be able to resemble the behaviour of real-world time series. They do come in handy when a raw estimation of future prices is needed; however, their forecasted values should be taken lightly and be used just to ascertain the price's trend as opposed to the amount it will increase or decrease by.

This thesis not only assessed the accuracy of two of the most common linear and non-linear forecasting methods, ARIMA and GARCH, respectively, but also outlined the use of a machine learning method that has proven to model accurately the hybrid patterns (linear and non-linear) that real-world time series exhibit, i.e., the ANN forecasting technique.

This work showed that the most widely used forecasting method, i.e., ARIMA, though easy to use, falls short in terms of the accuracy when estimating future values to other more sophisticated forecasting models, be they the GARCH or the ANN processes. This inability to portrait the behaviour of the real future prices of a mineral commodity can be attributed to the constrain that working with a linear model imposes to the forecasted values.

The GARCH forecasting technique, on the other hand, despite being built on the same idea of using past observations to determine future prices, exhibits a much more accurate representation of the real nickel prices. Due to the alleviation of the linear constraints of the ARIMA model, the GARCH method was able to predict the overall trend, cyclicality, and fluctuations of the time series representing the nickel prices fairly well.

Nowadays, computer science has provided yet another tool in the wide range of forecasting techniques. They have been proven to be quite capable of identifying the hidden patterns (both linear and non-linear) that other traditional forecasting methods fail to. Through constant training of the model over and over for a number of times, machine learning techniques are able to represent the behaviour of real-world time series more accurately.

The ANN forecasting method, in this thesis, has shown a remarkable improvement over both the ARIMA and the GARCH techniques in terms of accuracy over time. Nonetheless, it seems that the ANN models tend to work better when the span of the forecast is shorter and up to 18 months since their accuracy starts decreasing past this threshold. In fact, this occurrence makes sense due to the fact that as part of the training of the ANN model, lesser past observations are considered in the training of the next value.

The work done in this thesis has also revealed what was mentioned lines above. Overall, all the forecasting methods explained in Chapter 4 through 6 show better performance at predicted mineral commodity prices in the short run. The time span for the forecast is still up for debate; however, it can range from a few months to a year (or a year and a half). Indeed, the performance in the error measures for all the forecasting techniques tends to deviate a lot more past the threshold (only for this case study) of 18 months, indicating that long-term forecasting using said techniques might lead to wrong interpretations.

Future work on this topic could involve the integration of the ARIMA model into the ANN framework, creating a hybrid model. This setup would require ‘separating’ the linear and non-linear components of the time series, allowing each one of the components to be approached by the correct technique (Zhang, 2003). That is, for the time series originated from a linear process the ARIMA technique could be used; whereas for that generated from a non-linear process, the

ANN framework should be applied. This methodology should yield a forecasting accuracy greater than those obtained using each one of the methods alone, as per it is suggested by the literature.

Other famous machine learning methods used in the forecasting literature are the Support Vector Machines (SVM) and Fuzzy Systems (FS). Based on experience minimization and statistical learning theory, the SVM technique seems to offer better both learning capacity & performance and generalization in comparison to traditional forecasting models; however, it is limited to small sample size and might not be able to yield accurate results when working with large data sets (Zhang et al., 2015). The FS method, on the other hand, is not able to learn and only provides approximate value(s) rather than precise one; nonetheless, its ability to work with incomplete date makes it a popular forecasting tool (Tapia Cortez, 2018). Both methods have also been combined with other traditional models in order to improve the accuracy of the forecasts. Methods such as SVM- GARCH (Zhang et al., 2015), ANNs integrated with FS (also called ANFIS) (Alameer at al., 2019), ANN-GFS (ANNs with Generic Fuzzy Systems, a category of FS) (Hadavandi et al, 2010; Guan at al., 2018; Barros and Medeiros, 2017) have been extensively used in the forecasting of electricity spot price as well as oil prices. These forecasting models might also be explored as apart of the future work.

Besides machine learning techniques, another method that has gathered a lot of attention from researchers because of its ability to deal with the complexities of dynamic systems is Chaos theory (CT). In fact, CT offers another alternative to the ever-challenging forecasting of real-world non-linear time series. In fact, CT ascertains that the chaotic behaviour of real-world time series produced from non-linear processes does actually follow a regular structure. This apparent randomness can be ‘treated’ via different methods to detect the hidden patterns that govern the

dynamics of the time series, allowing the observation of such properties to be the foundation of the estimation of future values (Picano et. al., 2019).

Ultimately, the forecasting of mineral commodity prices would remain an open-ended question and a constant process. Many factors, including but not limited to the availability of data, the forecasting accuracy we want to achieve, the time needed to conduct a thorough analysis and the forecasting time span (Chamber et. al., 1971), must be weighed in order to fully ascertain what the best forecasting technique is.

8. References

Agnolucci, P. (2009). Volatility in crude oil futures: A comparison of the predictive ability of GARCH and implied volatility models. *Energy Economics*, v. 31, issue 2, pp. 316-321, https://doi.org/10.1016/j.eneco.2008.11.001

Alameer, Z., Elaziz, M.A., Ewees, A.A., Ye, H., & Jianhua, Z. (2019). Forecasting Copper Prices Using Hybrid Adaptive Neuro-Fuzzy Inference System and Genetic Algorithms. *Natural Resource Research*, v, 28, pp. 1385–1401, https://doi.org/10.1007/s11053-019-09473-w

Arslanalp, S., Liao, W., Piao, S., & Seneviratne, D. (2016). China's Growing Influence on Asian Financial Markets. International Monetary Fund Working Paper, https://doi.org/10.5089/97814755251

Bakas, D., & Triantafyllou, A. (2019). Volatility forecasting in commodity markets using macro uncertainty. *Energy Economics*, v. 81, pp. 79-94, https://doi.org/10.1016/j.eneco.2019.03.016

Barros, M., Medeiros, L. (2017). Electricity spot price forecasting in Brazil using a hybrid neuro-fuzzy system and neural network approach. *International Journal of Energy and Statistics*, v. 05, issue 01, https://doi.org/10.1142/S2335680417500041

Bauwens, L., Hafner, C., & Laurent, S. (2012). Handbook of Volatility Models and Their Applications. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781118272039, https://onlinelibrary.wiley.com/doi/book/10.1002/9781118272039

Berlinger, E., Illés, F., Badics, M., Badics, M., & Banai, A. (2015). Mastering R for Quantitative Finance. Packt Publishing, O'Reilly Media, https://learning.oreilly.com/library/view/mastering-r-for/9781783552078/

Bernard, J., Khalaf, L., Kichian, M., & McMahon, S. (2006). *Forecasting Commodity Prices: GARCH, Jumps, and Mean Reversion*. Bank of Canada, Staff Working Paper 2006-14, https://www.bankofcanada.ca/2006/04/working-paper-2006-14/

Boateng, M. (2015). Volatility Assessment of Equities on the Ghana Stock Exchange. *International Journal of Statistics and Applications*, v. 5, issue 6, pp. 288-292, DOI: 10.5923/j.statistics.20150506.04, http://article.sapub.org/10.5923.j.statistics.20150506.04.html#Ref

Bönner, A. (2009). *Forecasting Models for the German Office Market*. Dissertation Universität St. Gallen, Gabler Edition Wissenschaft, https://doi.org/10.1007/978-3-8349-9402-8

Botchkarev, A. (2018). Performance Metrics (Error Measures) in Machine Learning, Regression, Forecasting and Prognosis: Properties and Typology. *Interdisciplinary Journal of Information, Knowledge and Management,* v. 14, pp. 45-79, DOI: 10.28945/4184

Box, G., Jenkins, G., Reinsel, G., Ljung, G., & Ljung, G. (2016). Time Series Analysis: Forecasting and Control. John Wiley & Sons, Incorporated, ProQuest Ebook Central, https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=2064681

Brockwell, P., & Davis, R. (2016). Introduction to Time Series and Forecasting. Springer Texts in Statistics, Springer, New York, DOI 10.1007/978-3-319-29854-2, https://link.springer.com/book/10.1007%2F978-3-319-29854-2

Bulut, B. (2010). Forecasting the Prices of Non-Ferrous Metals with GARCH Models. Middle East Technical University, Ankara, Turkey, https://open.metu.edu.tr/handle/11511/19949

Čermák, M., Malec, K., & Maitah, M. (2017). Price Volatility Modelling – Wheat: GARCH Model Application. *Agris on-line Papers in Economics and Informatics*, v. 9, issue 4, DOI: 10.22004/ag.econ.27606, https://ageconsearch.umn.edu/record/276061/?ln=en

Chambers, J., Mullick, S., & Smith, D. (1971). How to Choose the Right Forecasting Technique. *Harvard Business Review*, July 1971 issue, https://hbr.org/1971/07/how-to-choose-the-right-forecasting-technique

Chatfield C. (2010). The Analysis of Time Series: Theory and Practice. John Wiley & Sons, Inc., New York, DOI 10.1007/978-1-4899-2925-9.

Chen, A., Leung, M., & Daouk, H. (2003). Application of neural networks to an emerging financial market: forecasting and trading the Taiwan Stock Index. *Computers & Operations Research*, v. 30, issue 6, pp. 901-923, https://doi.org/10.1016/S0305-0548(02)00037-0

Cuddington, J., & Jerrett, D. (2008). Super Cycles in Real Metal Prices. *International Monetary Funds Staff Papers*, v. 55, number 44, https://www.imf.org/external/pubs/ft/staffp/2008/04/cuddington.htm

Danielsson, J., & Danaielsson, J. (2011). Financial Risk Forecasting : The Theory and Practice of Forecasting Market Risk with Implementation in R and Matlab. John Wiley & Sons Incorporated, ProQuest Ebook Central, DOI:10.1002/9781119205869 https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=699340

De Gooijer, Jan (2017). Elements of Nonlinear Time Series Analysis and Forecasting. Springer Series in Statistics, Springer, New York, DOI: 10.1007/978-3-319-43252-6, https://www.springer.com/gp/book/9783319432519

Deloitte (2017). Commodity Price Risk Management. Deloitte India, https://www2.deloitte.com/in/en/pages/risk/articles/commodity-price-risk-management.html

Dooley, G., & Lenihan, H. (2005). An assessment of time series methods in metal price forecasting. *Resources Policy*, v. 30, issue 3, pp. 208-217, https://doi.org/10.1016/j.resourpol.2005.08.007.

Dritsaki, C. (2018). The Performance of Hybrid ARIMA-GARCH Modeling and Forecasting Oil Price. *International Journal of Energy Economics and Policy*, Econojournals, v. 8, issue 3, pp. 14-21, https://www.econjournals.com/index.php/ijeep/article/view/6437

Efimova, O., & Serletis, A. (2014). Energy markets volatility modelling using GARCH. *Energy Economics*, v. 43, pp. 264-273, https://doi.org/10.1016/j.eneco.2014.02.018

Engle, R. (2001). GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *Journal of Economic Perspectives*, vol. 15, issue 4, pp. 157–168, DOI: 10.1257/jep.15.4.157, https://www.aeaweb.org/articles?id=10.1257/jep.15.4.157

Espí Rodríguez, J.A., & De La Torre, L. (2013). Factors influencing metal price selection in mining feasibility studies. *Mining Engineering*, v. 65 (n. 8), pp. 45-51. https://proxy.library.mcgill.ca/login?url=https://www-proquest-com.proxy3.library.mcgill.ca/trade-journals/factors-influencing-metal-price-selection-mining/docview/1554329830/se-2?accountid=12339

Fabozzi, F., Focardi, S., Rachev, S., & Arshanapalli, B. (2014). The Basics of Financial Econometrics: Tools, Concepts, and Asset Management Applications. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781118856406, https://onlinelibrary.wiley.com/doi/book/10.1002/9781118856406

Gordon, R., & Tilton, J. (2008). Mineral economics: Overview of a discipline. *Resources Policy*, v. 33, issue 1, pp. 4-11, https://doi.org/10.1016/j.resourpol.2008.01.003

Gocht, W., Zantop, H., & Eggert, R. (1988). International Mineral Economics. Springer, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-73321-5_8

Gregoriou, G., & Pascalau, R. (2018). Nonlinear Financial Econometrics: Forecasting Models, Computational and Bayesian Models. Palgrave Macmillan, New York, DOI 10.1057/9780230295223, https://www.palgrave.com/gp/book/9780230283657

Guan H., Dai Z., Zhao A., He J. (2018). A novel stock forecasting model based on High order-fuzzy-fluctuation Trends and Back Propagation Neural Network. *Plos One*, v. 13, issue 2, https://doi.org/10.1371/journal. pone.0192366

Guidolin, M., & Pedio, M. (2018). Essentials of Time Series for Financial Applications. Academic Press, O'Reilly Media, https://learning.oreilly.com/library/view/essentials-of-time/9780128134108/

Guzmán, J.I., & Silva, E. (2018). Copper price determination: fundamentals versus non-fundamentals. *Mineral Economics*, v. 31, pp. 283-300, https://doi.org/10.1007/s13563-017-0130-y

Hadavandi, E., Shavandi, H., Ghanbari, A. (2010). Integration of genetic fuzzy systems and artificial neural networks for stock price forecasting. Knowledge-Based Systems, v. 23, issue 8, pp. 800-808, https://doi.org/10.1016/j.knosys.2010.05.004

Heap, A. (2005). China – The Engine of a Commodity Super Cycle. New York, Citigroup, Smith Barney.

Hyndman, R.J., & Athanasopoulos, G. (2018). Forecasting: principles and practice, 2nd edition, OTexts: Melbourne, Australia, https://otexts.com/fpp2/nnetar.html

Hou, A., & Suardi, S. (2012). A nonparametric GARCH model of crude oil price return volatility. *Energy Economics*, v. 34, issue 2, pp. 618-626, https://doi.org/10.1016/j.eneco.2011.08.004

Hudson, R. S., & Gregoriou, A. (2015). Calculating and comparing security returns is harder than you think: A comparison between logarithmic and simple returns. *International Review of Financial Analysis*, v. 38, pp.151-162, https://doi.org/10.1016/j.irfa.2014.10.008

Irwin, S., & Sanders, D. (2011). Index Funds, Financialization, and Commodity Futures Markets. *Applied Economic Perspectives and Policy*, v. 33, number 1, pp. 1– 31, DOI:10.1093/aepp/ppq032, https://onlinelibrary.wiley.com/doi/pdfdirect/10.1093/aepp/ppq032

Jones, O., Aspandiar, M., Dugdale, A., Leggo, N., Glacken, I., & Smith, B. (2019). The Business of Mining: Mineral Deposits, Exploration and Ore-Reserve Estimation. CRC Press, London, https://doi.org/10.1201/9780429057540

Khashei, M., & Bijari, M. (2010). An artificial neural network (p,d,q) model for timeseries forecasting, *Expert Systems with Applications*, v. 37, issue 1, pp. 479-489, https://doi.org/10.1016/j.eswa.2009.05.044.

Knight, J., & Satchell, S. (2007). Forecasting Volatility in the Financial Markets. Elsevier Science & Technology, ProQuest Ebook Central, https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=287974

Kriechbaumer, T., Angus, A., Parsons, D., & Rivas Casado, M. (2014). An improved wavelet–ARIMA approach for forecasting metal prices. *Resources Policy*, v. 39, pp. 32-41, https://doi.org/10.1016/j.resourpol.2013.10.005

Kristjanpoller, W., & Minutolo, M. (2015). Gold price volatility: A forecasting approach using the Artificial Neural Network–GARCH model, *Expert Systems with Applications*, v. 42, issue 20, pp. 7245-7251, https://doi.org/10.1016/j.eswa.2015.04.058

Labys, W.C. (2006). Modeling and Forecasting Primary Commodity Prices. Routledge, London, https://doi.org/10.4324/9781315248783

Lama, A. (2015). Modelling and Forecasting of Price Volatility: An Application of GARCH and EGARCH Models. Agricultural Economics Research Review, v. 28, issue, pp 73-82, DOI: 10.5958/0974-0279.2015.00005.1

Lorente-Leyva, L., Pavón-Valencia, J., Montero-Santos, Y., Herrera-Granda, I., Herrera-Granda, E., & Peluffo-Ordóñez, D. (2019). Artificial Neural Networks for Urban Water Demand

Forecasting: A Case Study. *Journal of Physics: Conference Series*, DOI:10.1088/1742-6596/1284/1/012004, https://iopscience.iop.org/article/10.1088/1742-6596/1284/1/012004

Lütkepohl, H. (2005). New Introduction to Multiple Time Series Analysis. Springer, Springer Link, https://doi.org/10.1007/978-3-540-27752-1

Matyjaszek, M., Riesgo Fernández, P., Krzemień, A., Wodarski, K., & Fidalgo Valverde, G. (2019). Forecasting coking coal prices by means of ARIMA models and neural networks, considering the transgenic time series theory. *Resources Policy*, v. 61, pp. 283-292, https://doi.org/10.1016/j.resourpol.2019.02.017.

Mills, T. (2015). Time Series Econometrics: A Concise Introduction. Palgrave Texts in Econometrics, Springer Link, https://link.springer.com/book/10.1057%2F9781137525338

Mombeini, H., & Yazdani-Chamzini, A. (2015). Modeling Gold Price via Artificial Neural Network. *Journal of Economics, Business and Management*, v. 3, issue 7, DOI: 10.7763/joebm.v3.269, http://www.joebm.com/index.php?m=content&c=index&a=show&catid=45&id=572

Newbold, P., & Granger, C. (1974). Experience with Forecasting Univariate Time Series and the Combination of Forecasts. *Journal of the Royal Statistical Society*, Series A (General), v. 137, issue 2, pp. 131-165, DOI: 10.2307/2344546, https://www.jstor.org/stable/2344546?seq=1

Nau, R. (2019). Statistical Forecasting: Notes on Regression and Time Series Analysis. Fuqua School of Business, Duke University, Durham, http://people.duke.edu/~rnau/411home.htm

Omar, H., Hoang, V., & Liu, D. (2016). A Hybrid Neural Network Model for Sales Forecasting Based on ARIMA and Search Popularity of Article Titles. *Computational Intelligence and Neuroscience*, vol. 2016, Article ID 9656453, http://dx.doi.org/10.1155/2016/9656453

Panella, M., Barcellona, F., & D'Ecclesia, R. (2012). Forecasting Energy Commodity Prices Using Neural Networks. *Advances in Decision Sciences*, vol. 2012, Article ID 289810, https://doi.org/10.1155/2012/289810

Paolella, M. (2019). Linear Models and Time-Series Analysis: Regression, ANOVA, ARMA and GARCH. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781119432036, https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119432036

Pierdzioch, C., Rülke, J., & Stadtmann, G. (2013). Forecasting metal prices: Do forecasters herd? *Journal of Banking & Finance*, v. 37, issue 1, pp. 150-158, https://doi.org/10.1016/j.jbankfin.2012.08.016

Picano, B., Fantacci, R., & Han, Z. (2019). Passengers Demand Forecasting Based on Chaos Theory. *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, Shanghai, China, pp. 1-6, DOI: 10.1109/ICC.2019.8762041, https://ieeexplore.ieee.org/document/8762041

Piot-Lepetit , I., & M'Barek, R. (2011). Methods to Analyse Agricultural Commodity Price Volatility. Springer, New York, DOI 10.1007/978-1-4419-7634-5 https://www.springer.com/gp/book/9781441976338

Pirrong, C. (2011). Commodity Price Dynamics: A Structural Approach. Cambridge: Cambridge University Press, DOI:10.1017/CBO9781139018142, https://www.cambridge.org/core/books/commodity-price-dynamics/9C67EFBF989C2B5A8E9B8D8E12843476

Priolon, J. (2018). Financial Markets for Commodities. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781119579274, https://onlinelibrary.wiley.com/doi/book/10.1002/9781119579274

Radetzki, M., & Wårell, L. (2016). A Handbook of Primary Commodities in the Global Economy, Cambridge University Press, pp. 28-53. https://doi.org/10.1017/9781108886529.003

Radetzki, M. (1983). Long-run Price Prospects for Aluminium and Copper. *Natural Resources Forum*, v. 7, issue 1, pp. 23-36, https://doi.org/10.1111/j.1477-8947.1983.tb00016.x

Rossen, A. (2015). What are metal prices like? Co-movement, price cycles and long-run trends. *Resources Policy*, v. 45, pp. 255-276, https://doi.org/10.1016/j.resourpol.2015.06.002

Ru, Y., & Ren, H. J. (2012). Application of ARMA Model in Forecasting Aluminium Price. *Applied Mechanics and Materials*, v. 155-156, pp. 66-71, https://doi.org/10.4028/www.scientific.net/amm.155-156.66

Rudenno, V. (2012). The Mining Valuation Handbook: Mining and Energy Valuation for Investors and Management. John Wiley & Sons, Incorporated, ProQuest Ebook Central, https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=918592

Ruppert, D. (2011). GARCH Models. Statistics and Data Analysis for Financial Engineering. Springer Texts in Statistics, Springer, New York, https://doi.org/10.1007/978-1-4419-7787-8_18

Schaeffer , P. (2008). Commodity Modeling and Pricing: Methods for Analyzing Resource Market Behavior. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781118267905, https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118267905

Shafiee, S., & Topal, E. (2010). An overview of global gold market and gold price forecasting. *Resources Policy*, v 35, issue 3, pp. 178-189, DOI: 10.1016/j.resourpol.2010.05.004, https://doi.org/10.1016/j.resourpol.2010.05.004

Takatoshi, I., & Rose, A. (2011). Commodity Prices and Markets. *National Bureau of Economic Research East Asia Seminar on Economics*, v. 20 https://chicago.universitypressscholarship.com/view/10.7208/chicago/9780226386904.001.0001/upso-9780226386898

Tan, Z., Zhang, J., Wang, J., & Xu, J. (2010). Day-ahead electricity price forecasting using wavelet transform combined with ARIMA and GARCH models, *Applied Energy*, v. 87, issue 11, pp. 3606-3610, https://doi.org/10.1016/j.apenergy.2010.05.012

Tapia Cortez, C.A., Saydam, S., Coulton, J., & Sammut, C. (2018). Alternative techniques for forecasting mineral commodity prices, *International Journal of Mining Science and Technology*, v. 28, issue 2, pp. 309-322, https://doi.org/10.1016/j.ijmst.2017.09.001

Tang, Z., & Fishwick, A. (1998). Feedforward Neural Nets as Models for Time Series Forecasting. INFORMS, Journal on Computing, v.5, issue 4, pp. 374-385, http://dx.doi.org/10.1287/ijoc.5.4.374

Tauber, M., Feth, M., & Möncks, T. (2015). Managing Commodity Price Uncertainties. The Boston Consulting Group (BCG), Canada, www.bcg.com/en-ca/publications/2015/metals-mining-managing-commodity-price-uncertainties. Accessed 8th Jan 2021.

Tealab, A. (2018). Time series forecasting using artificial neural networks methodologies: A systematic review. *Future Computing and Informatics Journal*, v. 3, issue 2, pp. 334-340, https://doi.org/10.1016/j.fcij.2018.10.003

The London Metal Exchange (2019). A Guide to the London Metal Exchange. https://www.lme.com/en-GB/Education;and-events/Online-resources/Brochures#tabIndex=0

Triantafyllou, A., Dotsis, G., & Sarris, A. (2015). Volatility Forecasting and Time-varying Variance Risk Premiums in Grains Commodity Markets. *Journal of Agricultural Economics*, v. 66, pp. 329-357, https://doi.org/10.1111/1477-9552.12101

Tsay, R. (2012). An Introduction to Analysis of Financial Data with R. John Wiley & Sons, Incorporated, ProQuest Ebook Central, https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=1771583

Tully, E., & Lucey, B. (2007). A power GARCH examination of the gold market. *Research in International Business and Finance*, v. 21, issue 2, pp. 316-325, https://doi.org/10.1016/j.ribaf.2006.07.001

Viens, F., Mariani, M., & Florescu, I. (2012). Handbook of Modeling High-Frequency Data in Finance. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781118204580 https://onlinelibrary.wiley.com/doi/book/10.1002/9781118204580

Virginia, E., Ginting, J., & Elfaki, F. (2018). Application of GARCH Model to Forecast Data and Volatility of Share Price of Energy (Study on Adaro Energy Tbk, LQ45). *International*

Journal of Energy Economics and Policy, Econjournals, v. 8, issue 3, pp. 131-140, https://www.econjournals.com/index.php/ijeep/article/view/6352

Wang, Y., & Wu, C. (2012). Forecasting energy market volatility using GARCH models: Can multivariate models beat univariate models? *Energy Economics*, v.34, issue 6, pp. 2167-2181, DOI: 10.1016/j.eneco.2012.03.010

Weron, R. (2013). Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach. John Wiley & Sons, Incorporated, Wiley Online Library, DOI:10.1002/9781118673362, https://onlinelibrary.wiley.com/doi/book/10.1002/9781118673362

Wongsathan, R., & Seedadan, I. (2016). A Hybrid ARIMA and Neural Networks Model for PM-10 Pollution Estimation: The Case of Chiang Mai City Moat Area, *Procedia Computer Science*, v. 86, pp. 273-276, https://doi.org/10.1016/j.procs.2016.05.057

Yaziz, S., Azizan, N., Zakaria, R., & Ahmad, M. (2013). The Performance of Hybrid ARIMA-GARCH Modelling in Forecasting Gold Price. In 20th International Congress on Modelling & Simulation, pp. 1201-1207, http://www.mssanz.org.au/modsim2013/F2/yaziz.pdf

Xiao, L., & Aydemir, A. (2007). Volatility modelling and forecasting in finance. In Quantitative Finance, Forecasting Volatility in the Financial Markets, pp. 1-45, https://doi.org/10.1016/B978-075066942-9.50003-0

Zhang, J., Zhang, Y., Zhang, L. (2015). A novel hybrid method for crude oil price forecasting. Energy Economics, v. 49, pp. 649–659, https://doi.org/10.1016/j.eneco.2015.02.018

Zhang, G., Patuwo, B., & Hu, M. (1998). Forecasting with artificial neural networks: The state of the art. *International Journal of Forecasting*, v. 14, issue 1, pp. 35-62, https://doi.org/10.1016/S0169-2070(97)00044-7

Zhang, K., Olawoyin, R., Nieto, A., & Kleit, A. (2018). Risk of commodity price, production cost and time to build in resource economics. *Environment, Development and Sustainability*, v. 20, pp. 2521–2544, DOI: 10.1007/s10668-017-0003-0, https://doi.org/10.1007/s10668-017-0003-0

Zhang, P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, v. 50, pp.159-175, https://doi.org/10.1016/S0925-2312(01)00702-0

Zsidisin, G. (2011). Managing Commodity Price Risk : A Supply Chain Perspective. Business Expert Press, ProQuest Ebook Central, https://ebookcentral.proquest.com/lib/mcgill/detail.action?docID=876643