UNIVERSITAT POLITÈCNICA DE VALÈNCIA

ESCUELA TÉCNICA SUPERIOR INGENIEROS INDUSTRIALES VALENCIA

TRABAJO FIN DE MASTER EN INGENIERÍA INDUSTRIAL

FORECASTING THE PRODUCTION VOLUME OF A MANUFACTURING COMPANY USING TIME-SERIES ANALYSIS METHODS

AUTHOR: JAVIER PÉREZ TRENCO

SUPERVISOR: dr hab. MAREK KARKULA Eng.

Academic year: 2017-18

AGH UNIVERSITY OF SCIENCE AND TECHNOLOGY

FACULTY OF MANAGEMENT

AGH

AKADEMIA GÓRNICZO-HUTNICZA
IM. STANISŁAWA STASZICA
W KRAKOWIE

FORCASTING THE PRODUCTION VOLUME OF A MANUFACTURING COMPANY USING TIME-SERIES ANALYSIS METHODS

MASTER'S DEGREE IN INDUSTRIAL ENGINEERING

(POLYTECHNIC UNVIERSITY OF VALENCIA)

MASTER'S THESIS

YEAR 2017/2018

Javier Pérez Trenco

June 2018

Supervisor: dr hab. Marek Karkula Eng.

ABSTRACT

Forecasting the production volume of a manufacturing company using time-series analysis methods

by

Javier Pérez Trenco

June, 2018

Cracow, Poland

Keywords: forecasting; time-series analysis; production volume; moving average; exponential smoothing; ARIMA; quantitative methods; qualitative methods; manufacturing company

The purpose of this thesis is to try different time-series analysis methods to forecast the future production volume of different products of a manufacturing company. Currently, the company uses the experience of managers to forecast the production volume. The aim of this study is to find if quantitative methods can improve such qualitative predictions.

In literature review, an introduction to forecasting is presented. Additionally, the most commonly time-series forecasting methods are described. At the end, different performance indicators are briefly explained.

In the empirical part of the study, first the information about the company and the data are presented. Secondly, six different methods are tried to forecast the production volume. Then, the performance of these methods is evaluated and compared with the qualitative predictions made by the company.

At the end, the results are analysed and a brief conclusion is presented. With the results and the conclusions, the company has a proof of how time-series analysis methods work with their products. They can examine the results and decide if it is beneficial for the company to implement these methods.

ACKNOWLEDGEMENT

I would first like to thank my thesis supervisor dr hab. Marek Karkula Eng. of the Faculty of Management at AGH University of Science and Technology. The door to Prof. Karkula office was always open whenever I had a question about my research, or when I needed some guidance to continue with my experiments.

I would also like to thank the AC Operations Network Manager of the company, Bronislav Bronislavovich Fialkovskij, who provided with the data and has been involved in the development of this thesis from the beginning. Without his cooperation and his passionate participation, the study could not have been successfully conducted.

Cracow, June 2018

Javier Pérez Trenco

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

NOTATION: SYMBOLS AND ABBREVIATIONS

$\boldsymbol{x_t}$: demand in period t

$\boldsymbol{a}$: level, or average demand per period

$\boldsymbol{b}$: trend

$\boldsymbol{\varepsilon_t}$: irregular random fluctuations with mean zero

$\boldsymbol{F_t}$: seasonal index in period t

$\boldsymbol{\hat{x}_{t,t+\tau}}$:forecast made at the end of period t, for a future period $t+\tau$, with $\tau > 0$

$\boldsymbol{e_{t,t+\tau}}$: error associated with the forecast made in period t for period $t+\tau$

$\boldsymbol{N}$: number of periods used in the moving average model

$\boldsymbol{\bar{x}_{t,N}}$: simple N-period moving average calculated at the end of period N

$\boldsymbol{\hat{a}_t}$: estimate of the level a, after observing the demand in period t

$\boldsymbol{\hat{b}_t}$: estimate of the trend b, after observing the demand in period t

$\boldsymbol{\hat{F}_t}$: estimate of the seasonal index F_t, after observing the demand in period t

$\boldsymbol{\alpha}$: smoothing constant between 0 and 1

$\boldsymbol{\beta}$: smoothing constant between 0 and 1

$\boldsymbol{\gamma}$: smoothing constant between 0 and 1

$\boldsymbol{p}$: order of the non-seasonal autoregressive component in ARIMA model

$\boldsymbol{d}$: number of non-seasonal differences applied to the time-series in ARIMA model

$\boldsymbol{q}$: order of the non-seasonal moving average component in ARIMA model

$\boldsymbol{P}$: order of the seasonal autoregressive component in ARIMA model

$\boldsymbol{D}$: number of seasonal differences applied to the time-series in ARIMA model

$\boldsymbol{Q}$: order of the seasonal moving average component

$\boldsymbol{\varphi_t}$: parameter for period t of the autoregressive model

$\boldsymbol{\theta_t}$: parameter for period t of the moving average model

$\boldsymbol{MSE}$: mean squared error

$\boldsymbol{RMSE}$: root mean squared error

$\boldsymbol{MAE}$: mean absolute error

$\boldsymbol{MAPE}$: mean absolute percentage error

1. INTRODUCTION

Every manufacturing company needs a production plan that ensures that they have sufficient raw materials, staff and other necessary items, to create finished products according to their specified schedule. The aim of production planning is to maximise the profitability while maintaining the customer satisfaction.

One of the first steps in production planning is to forecast the future sales, demand or production volume. Manufacturing companies need to make accurate predictions of how much they are going to sell or to produce in the future. It is the base of every production plan, and the success of the company depends strongly on the performance of the forecasts.

This thesis is a case study with real data from a manufacturing company that produces materials for construction, such as tile adhesives, grouts, plasters or silicones. Currently, the company uses qualitative forecasting methods based on the experience of the managers. They do not use any quantitative method for forecasting their future production volume. This is common practice in many manufactures, and sometimes it gives satisfactory results, especially for short-term forecasts such as weekly or monthly forecasts.

However, theory and research suggest that generally, quantitative methods give better forecasts than qualitative techniques. Nevertheless, this depends on the nature of the demand of each product. For some products, such as new products or products with a very irregular demand, most authors recommend using judgemental methods. On the other side, when enough data is available and the demand is relatively stable, experts suggest that quantitative methods produce forecasts that are more accurate.

Although in many companies managers are aware of this, they find these kind of methods intimidating and they do not consider that it is advantageous to use them, since their qualitative predictions are acceptable. With this thesis, the author expects to prove that there are some quantitative methods that, although being very simple and easy to understand and implement, they can give very satisfactory forecasts and improve the qualitative predictions made by the company.

The objective of the thesis is to analyse the data provided by the company and to try different forecasting methods. The company has provided with the production volume of all their products, and with an example of their qualitative predictions. The main idea is to generate forecasts with different types of quantitative forecasting techniques, and to compare their performances with the company's judgemental predictions. Some of the methods are very simple and can be easily implemented in a spreadsheet. Other methods are more complex and need a deeper background in statistics and mathematics to understand how they work.

With the results of the thesis, the company can evaluate if using quantitative forecasting methods is an advantage for them. The author does not expect to suggest which option is the best for the company. The aim is to show the results of each method compared with the company's predictions, and to give a brief explanation of how the method works. The company can judge the results and balance if the improvement achieved with the each method compensates the effort of implementing it.

2. LITERATURE REVIEW

2.1. What is forecasting

Forecasting is evaluating a variable, e.g., sales, demand, or production volume, to estimate its value in future periods. Usually, forecasting calculations are a source of information for production planning in manufacturing companies.

The goal of forecasting is to make accurate predictions for the future, minimizing the deviation between the actual demand and the forecast.

There are many other applications of forecasting apart from production planning. Forecasting is widely used in economy and in other fields, such as meteorology, earthquake predictions or eGain forecasting.

2.2. Why is forecasting important

Forecasting is one of the most essential steps in production planning. Production planning, as the name suggests, is a projection of the future production activity.

Every manufacturing company needs to make forecasts of the demand of the items they produce. It is crucial to make accurate predictions of the future demand to know how many raw materials to buy, or to optimise the inventory levels while meeting customer expectations. Production planners need forecasts in order to schedule production activities, order materials and establish inventory levels.

Forecasts based on time-series are the most widely used. These kind of forecasts are based only on historical data. These methods do not consider exogenous variables in the production process.

However, many manufacturing companies determine their production plans using subjective and intuitive judgment forecasts. In many cases, this is one of the factors that might lead to production inefficiency.

The accuracy of the forecasts is crucial for the company, since it significantly affects inventory levels and customer service levels. For that reason, it is highly important to choose the best forecasting method for each product. When the forecast is inaccurate, the production plan will be unreliable and may result in lack or excess of goods in stock.

Both of those situations are damaging for the company. When the actual demand is lower than the forecast, the company will have to deal with an excess of goods in stock, which leads to unnecessary inventory costs. On the other hand, when the actual demand is higher than the forecast, some clients will have to wait to receive their orders. This is even more detrimental, since it damages the image of the company and might result in the loss of potential clients.

2.3. Qualitative vs quantitative forecasting methods

Overall, there are three different kind of forecasting techniques: qualitative techniques, time-series analysis and projection methods, and causal models. The last two types are both quantitative techniques (Chambers, Mullick and Smith, 1971).

2.3.1. Qualitative methods

Qualitative forecasting techniques are widely used in manufacturing companies. These methods are based mostly in the experience of managers and the opinion of experts. They use qualitative data and may or may not consider the past data.

These methods are primarily used when there is no data, or not enough data, available. This is the only way to forecast the demand of a product that has just been introduced into a market. In order to predict the future demand, production planners must use human judgement to turn qualitative data into quantitative estimates.

Some of the most widely used qualitative methods in forecasting are the Delphi method, the market research, the panel consensus and the historical analogy.

2.3.2. Time-series analysis and projection methods

Time-series analysis methods rely entirely on historical data and focus only on patterns and pattern changes. The most used time-series analysis methods are moving averages, exponential smoothing and ARIMA or Box-Jenkins techniques. These methods will be explained in detail in this thesis.

2.3.3. Causal models

Causal models assume that the variable being forecast is related to other variables. The challenge is to find the relationships between the forecast variable and the others. These relationships are sometimes complicated, and complex mathematical models are needed.

2.4. Components of a time-series

"A time-series is a sequential set of data points, measured typically over successive time units. It is mathematically defined as a set of vectors $x(t)$, $t = 0,1,2,\ldots$, where t represents the time elapsed. The variable $x(t)$, is treated as a random variable. The measurements taken during an event in a time-series are arranged in a proper chronological order" (Adhikari and Agrawal, 2013, p. 12).

Depending on how many variables the time-series contains, it can be univariate (if it has only one single variable) or multivariate (if more than one variable is considered).

A time-series can be continuous or discrete. Continuous time-series are obtained when observations are recorded over some time interval. On the other hand, discrete time-series contain observations measured at discrete points of time.

There are five components that can be present in a time-series: the level, the trend, the seasonality, the cyclic variations and the irregular fluctuations. These five components might be present or not, and they can be separated from the original time-series.

2.4.1. Level

The level of a time-series is the average value that is assumed stable over time. If there were only the level, the time-series would be constant.

A constant model can represent some time-series. However, the level is not the only component in these models. Random deviations are also present. Nevertheless, the average value is relatively stable compared to the random deviations.

Usually, the letter a represents the level. A constant model can be represented as in equation (2.1).

$$x_t = a + \varepsilon_t \tag{2.1}$$

where:

x_t – demand in period t,

a – level, or average demand per period,

ε_t – irregular random fluctuations with mean zero.

In some cases, a constant model can represent the demand of a product. It is the case of products whose demand is not expected to follow a trend or a seasonal pattern. The most common examples are products that are in a mature stage of their life cycle, and are used regularly, e.g., toothpaste or toilet paper (Axsäter, 2006).

2.4.2. Trend

When the demand grows or decreases slowly over time, a trend model can represent it. The trend represents the long-term growth or decline of a time-series over time. A trend may, however, change its direction, which means that it can turn from increasing to decreasing, and vice versa.

In a trend model, the demand is represented by the equation (2.2).

$$x_t = a + bt + \varepsilon_t \tag{2.2}$$

where:

x_t – demand in period t,

a – average demand in period 0,

b – trend,

ε_t – irregular random fluctuations with mean zero.

2.4.3. Seasonality

"There is a seasonal pattern in a time-series when the series is influenced by seasonal factors (e.g., the quarter of the year, the month, or day of the week). Seasonality is always of a fixed and known period" (Hyndman and Athanasopoulos, 2013).

This type of variation is generally annual in period, whether measured weekly, monthly or quarterly, when similar patterns of behaviour are observed at particular times of the year. It is important to take into account the need to have enough historical data to determine if seasonality is present or not. (Chatfield, 2000). A seasonal model is only meaningful if the demand follows essentially the same pattern year after year (Axsäter, 2006).

"Seasonal variations can be of two kinds: (1) those resulting from natural forces and (2) those arising from human decisions or customs" (Silver, Pyke and Thomas, 2017, p. 73). An example of seasonal variations of the first kind are weather conditions. The demand of many products is strongly related with the weather conditions and the season of the year. For instance, the demand of ice creams is much higher in the summer than in the winter. On the other hand, various Christmas decorations are sold only during a very short period of the year.

An example of seasonal variations of the second kind are the increase in the demand of notebooks and pencils before the start of the school year. In this case, the seasonal variations depend on human decisions (Silver, Pyke and Thomas, 2017, p. 74).

The easiest way to identify seasonal patterns in time-series is the examination of the autocorrelation correlogram, which displays graphically and numerically the autocorrelation function.

Trend-seasonal model

In a trend-seasonal multiplicative model, the demand can be expressed using the equation (2.3) (Axsäter, 2006, pp. 10, 11).

$$x_t = (a + bt)F_t + \varepsilon_t \tag{2.3}$$

where:

- x_t – demand in period t,
- a – average demand in period 0,
- b – trend,
- F_t – seasonal index in period t. It indicates how much this period typically deviates from the annual average. At least one full season of data is required for the computation of the different F_t,
- ε_t – irregular random fluctuations with mean zero.

It is required that if there are T periods in one year, then $\sum_{k=1}^{T} F_{t+k} = T$ of any T consecutive periods.

Therefore, if for period t, $F_t = 1.2$, then the demand is expected to be 20% higher due to seasonal variations.

2.4.4. Cyclic variations

A cyclic pattern exists when the time-series exhibits rises and falls that are not of a fixed period. The duration of these variations is usually of two years or more (Hyndman and Athanasopoulos, 2013). Therefore, short and medium-term forecasts usually do not include cyclical effects.

An example of cyclic variations are business cycles, which duration is unknown beforehand, and it can last for several years.

In average, the length of the cycles is longer than the length of a seasonal pattern, and the magnitude of cycles tends to be more visible than the magnitude of seasonal patterns (Hyndman and Athanasopoulos, 2013).

2.4.5. Irregular fluctuations

Irregular fluctuations is the term often used to describe the residue that is left after the effects of the four components described above (level, trend, seasonality and cyclic variations) are removed from the time-series. Examples of irregular fluctuations are weather conditions and unexpected labour strikes (Silver, Pyke and Thomas, 2017, pp. 76, 77).

These variations are the result of unpredictable events. "They are completely random and they cannot be forecast. However, they may exhibit short-term correlation or include one-off discontinuities" (Chatfield, 2000, p. 23).

2.5. Multiplicative and additive models

Combining the five components of a time-series, two different kind of models can be formulated: multiplicative or additive models.

A multiplicative model of a time-series looks as follows (Silver, Pyke and Thomas, 2017, p. 77):

$$Demand = (Trend)(Seasonal)(Cyclic)(Irregular) = (b)(F)(C)(\varepsilon) \qquad (2.4)$$

On the other hand, an additive model can be written as (Silver, Pyke and Thomas, 2017, p. 77):

$$Demand = (Level) + (Trend) + (Seasonal) + (Cyclic) + (Irregular) \quad (2.5)$$
$$= a + bt + F_t + C_t + \varepsilon_t$$

There is a third option, which is partly additive and partly multiplicative, called mixed model.

The decision of which of these models should be used, depends on the nature of the time-series. For instance, the seasonality is said to be additive when it does not depend on the local mean level. On the other hand, the seasonality is multiplicative when the size of the seasonal variation is proportional to the local mean (Chatfield, 2000, p. 30).

For example, during the month of July, the sales of a particular ice cream may be one million dollars higher than the average in the whole year. Therefore, every year the forecaster would add to the forecasts for the month of July the amount of one million dollars over the average. In this case, the seasonality is additive.

Alternatively, during the month of July, the sales of a particular ice cream may be 1.4 times higher than the average. Thus, every year the forecast for the sales in July will be 40% higher than the average. Consequently, if the sales of that ice cream during that year are weak, the forecast for July will be also weak. In this case, the seasonality is multiplicative (Kalekar, 2004).

2.6. Forecasting methods

Time-series analysis techniques are used only when several years of data are available, and when relationships and trends are both clear and relatively stable (Chambers, Mullick and Smith, 1971). These methods depend only on historical data, and they use it to make projections for the future.

The most commonly used time-series analysis forecasting methods are described in this section.

2.6.1. Simple moving averages

In order to understand the concept of simple moving averages, it is necessary to know two of the simplest methods for predicting a model from historical data: the mean model and the random walk model.

The mean model

The mean model is often used to predict a variable whose values are independently and identically randomly distributed. With this model, we simply take the sample mean of all the previous values, as the forecast of the next value. The reason why the sample mean is used for forecasting the future values is because it is an unbiased predictor and it minimizes the mean squared forecasting error (Nau, 2018).

The random walk model

The random walk model is also one of the simplest models in time-series forecasting. In spite of its simplicity, it is often used in finance, physics, chemistry and biology, among many other fields. A time-series is said to follow a random walk if the first differences (difference from one observation to the next observation) are random. The time-series itself is not random, but its first differences are. A random walk model will predict that the following value of the time-series equals the last observed value (Imdadullah, 2013).

In summary, the mean model gives the same weight to all the previous values in the time-series to predict the next one, while the random walk gives all the weight to the most recent observation. However, there is a spectrum of possibilities between these two models described above. If we use a model that takes into account, for instance, the most recent five observations, we would have a mix of these two models. This is the concept of a moving average.

A moving average model improves both the mean model and the random walk model. It adapts better to cyclical patterns than the mean model, and it is less sensitive than the random walk model to random shocks from one period to the next one.

The simplest moving average model is computed giving the same weight to the most recent N observations. It is called simple moving average, and its idea is to calculate the forecast of the next value, as the average of the previous N values.

The notation that we will use to describe the moving average model is similar to the one used by Silver, Pyke and Thomas (2017):

- x_t is the demand observed in period t,

- $\hat{x}_{t,t+\tau}$ is the forecast made at the end of period t, for a future period $t+\tau$, with $\tau > 0$. It is usually called a "$lag - \tau$" forecast, since it is a forecast for a time period τ periods in the future. For example, $\hat{x}_{4,6}$ is the forecast made at the end of period 4 for the future demand in period 6, i.e., "lag-2" forecast made in period 4 (in this case, $t = 4$ and $\tau = 2$),

- $e_{t,t+\tau} = x_t - \hat{x}_{t,t+\tau}$ represents the forecast error associated with the forecast made in period t for period $t+\tau$. If the error is negative, it means that the forecast is higher than the actual demand observed.

Updating procedure

Using the described notation, a simple moving average model can be described with equation (2.6).

$$\bar{x}_{t,N} = \hat{a}_t = \frac{(x_t + x_{t-1} + x_{t-2} + \cdots + x_{t-N+1})}{N} \tag{2.6}$$

where $\bar{x}_{t,N} = \hat{a}_t$ is the simple N-period moving average calculated at the end of period N. It is just an average of the previous N observations, where x_t is the actual value of the time-series in period t (Silver, Pyke and Thomas, 2017).

Forecast

Using the simple average model, where the level is the only component of a time-series that is present, the forecast made at the end of period t for any future period $t + \tau$ is:

$$\hat{x}_{t,t+\tau} = \hat{a}_t \tag{2.7}$$

where $\hat{a}_t$ is the estimate of the level a, after observing the demand in period t.

In other words, in period t we calculate an estimate of the level, $\hat{a}_t$, using the previous N observations. Since we are considering a constant model, we can use $\hat{a}_t$ as the forecast for any future period, $\hat{x}_{t,t+\tau}$.

In a simple moving average model, all the N past observations get the same weight, which is of course $1/N$. This means that the larger N gets, the less is the weight given to each of the past observations. Therefore, for large values of N the model filters better the noise and gives more smoothed forecasts.

Nevertheless, there is a problem of using large values of N. The forecasts tend to lag behind in trying to follow trends or respond to turning points. The lag shown by the forecast depends strongly on N, since $lag = (N + 1)/2$. Therefore, the larger N is, the bigger will be the lag in our forecasts.

It is important to choose carefully the value of N. A big N value will filter more noise, but it will make the forecast slower to respond to trends and turning points. There is a need to make a trade-off between these two effects, in order to make forecasts that are more accurate.

For instance, if the level a is changing slowly and the stochastic deviations ε_t are large, we might prefer to use a large value of N. This way we will reduce the influence of the stochastic deviations. On the other hand, if a is varying quickly and the stochastic variations are small, we should rather choose a small value of N, which will allow the model to adapt better to the variations in the level a (Axsäter, 2006).

2.6.2. Simple exponential smoothing

Simple exponential smoothing is probably the most used statistical method for short-term forecasting. Its concept is in many ways similar to the moving average. Thus, it is very intuitive and simple to use. However, the way that the forecast is updated is not the same than in moving averages.

There is an undesired property in the updating procedure of moving averages: it gives the same weight to all the precedent observations, which means that they are treated equally. Nevertheless, it is obvious that the most recent data has a higher relation with the future. For this reason, it is appropriate to give a higher weight to those observations, and a lower weight to the oldest values (Nau, 2018).

In other words, the most recent observation should have the highest weight; the second most recent observation should get the second highest weight, and so on. This is what simple exponential smoothing does, and that is the reason why in most of the cases it makes better forecasts than moving averages.

Updating procedure

Assuming a constant model, we wish to estimate the parameter a, which represents the level of the time-series. We need to use a smoothing constant, a number between 0 and 1, which will determine the weight given to each past observation.

In most books and articles about forecasting, the Greek letter α is used to represent the smoothing constant.

Initialization

It is necessary to have an initial value of the level, $\hat{a}_{t-1}$ to start the forecasting procedure. If enough historical data exist, we can obtain this initial value of the level as the average demand in the last few periods. The number of periods that should be taken for this average depends on the nature of the time-series (Silver, Pyke and Thomas, 2017, pp. 86, 87).

Forecast

To update the forecast in period t, we use a linear combination of the previous forecast (the forecast made in period $t-1$) and the most recent observed demand, x_t:

$$\hat{x}_{t,t+\tau} = \hat{a}_t = (1-\alpha)\hat{a}_{t-1} + \alpha x_t \qquad (2.8)$$

where $t > \tau$ and α is the smoothing constant ($0 < \alpha < 1$) (Axsäter, 2006, p. 12).

It is important to note the meaning of using limit values of the smoothing constant α. Using $\alpha = 0$ results in $\hat{x}_{t,t+\tau} = \hat{a}_t = \hat{a}_{t-1}$, which means that the forecast made in period t is equal to the forecast made in period $t-1$. On the contrary, using $\alpha = 1$ results in $\hat{x}_{t,t+\tau} = \hat{a}_t = x_t$, which means that the forecast made in period t is equal to the actual demand observed in period t.

The geometric weighting of historical data in exponential smoothing can be expressed in a different way, if we substitute the corresponding relation for $\hat{a}_{t-1}$.

Since:

$$\hat{a}_{t-1} = (1-\alpha)\hat{a}_{t-2} + \alpha x_{t-1}, \tag{2.9}$$

Substituting (2.9) in (2.8), we have:

$$\hat{a}_t = (1-\alpha)\hat{a}_{t-1} + \alpha x_t = (1-\alpha)\big((1-\alpha)\hat{a}_{t-2} + \alpha x_{t-1}\big) + \alpha x_t \tag{2.10}$$

If we keep substituting (2.9), we obtain:

$$\begin{aligned}\hat{a}_t &= (1-\alpha)\hat{a}_{t-1} + \alpha x_t = (1-\alpha)\big((1-\alpha)\hat{a}_{t-2} + \alpha x_{t-1}\big) + \alpha x_t \\ &= \alpha x_t + \alpha(1-\alpha)x_{t-1} + (1-\alpha)^2\hat{a}_{t-2} = \cdots \\ &= \alpha x_t + \alpha(1-\alpha)x_{t-1} + \alpha(1-\alpha)^2 x_{t-2} + \cdots \\ &+ \alpha(1-\alpha)^n x_{t-n} + (1-\alpha)^{n+1}\hat{a}_{t-n-1}\end{aligned} \tag{2.11}$$

Comparing the previous equation with the equation of the moving average forecast, it is easy, after some simple calculations that we will skip here, to find a relation between the smoothing constant α and the number of periods N used in the moving average model:

$$\alpha = \frac{2}{N+1} \tag{2.12}$$

This equation might be used to estimate the value of α. For instance, the smoothing constant α corresponding to a moving average that is updated monthly with $N = 12$, is:

$$\alpha = \frac{2}{12+1} = \frac{2}{13} \approx 0.15$$

2.6.3. Exponential smoothing with a trend

The simple exponential smoothing procedure shown in the previous section is based on a constant model. Therefore, it is inappropriate to use it when there is a trend in the time-series. In this case, a more complicated procedure is needed. In case of a time-series that follows a trend model, we need to estimate two parameters. Apart from the level a, we also have to estimate the trend b. The model will also contain stochastic deviations ε_t, but we cannot predict them.

Updating procedure

a) Holt's linear exponential smoothing

Let us now describe the two most popular procedures to update the values of a and b. One of them is the procedure described by Holt. It is a natural extension of simple exponential smoothing. However, in this case we need two smoothing constants instead of one. The smoothing constants used in this method are α and β. Thus, the level and the trend are estimated using the equations (2.13) and (2.14).

$$\hat{a}_t = \alpha x_t + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1}) \qquad (2.13)$$

$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1} \qquad (2.14)$$

where α and β are smoothing constants between 0 and 1.

It is important to note that the difference $\hat{a}_t - \hat{a}_{t-1}$ is an estimate of the actual trend in period t. Therefore, we calculate $\hat{b}_t$ as a combination of this value and the estimate of the trend made in the previous period, $t-1$.

One of the most important decisions in every exponential smoothing method is the value of the smoothing constants. If there is more than one smoothing constant, the challenge is even greater. For that reason, it is crucial to understand how the value of these parameters affects the estimates.

As with simple exponential smoothing, high values of α and β make the forecasts react faster to changes, but it makes it more sensitive to stochastic deviations. Errors in the trend can give serious forecast errors for relatively long forecast horizons. Thus, it is recommended to choose low values of β (Axsäter, 2006, pp. 16, 17).

b) Brown's linear exponential smoothing

Brown described another procedure for estimating the level and the trend. It turns to be a special case of Holt's method. However, only one smoothing parameter is used. Brown's updating equations are (Silver, Pyke and Thomas, 2017, pp. 88, 89) (2.15) ad (2.16).

$$\hat{a}_t = [1-(1-\alpha)^2]x_t + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1}) \qquad (2.15)$$

$$\hat{b}_t = [\frac{\alpha^2}{1-(1-\alpha)^2}](\hat{a}_t - \hat{a}_{t-1}) + [1 - \frac{\alpha^2}{1-(1-\alpha)^2}]\hat{b}_{t-1} \quad (2.16)$$

where α is the single smoothing constant with value between 0 and 1.

Forecasting

Once we have estimated the level and the trend, we can make forecasts for future periods. However, there is something important to take into consideration. In the simple exponential smoothing, the forecast made in period t is the same for any future period, because the time-series is supposed to follow a constant model.

On the other hand, when the time-series follows a trend model, the forecast made in period t will be different for every future period $t+\tau$, with $\tau = 1,2,3,\dots$

The equation used to make forecasts with exponential smoothing with a trend is (2.17).

$$\hat{x}_{t,t+\tau} = \hat{a}_t + \hat{b}_t\tau \quad (2.17)$$

where $\hat{x}_{t,t+\tau}$ is the forecast made at the end of period t, of the demand in period $t+\tau$.

2.6.4. Holt-Winters exponential smoothing for a seasonal model

The methods that we have explained above are not the most suitable if the time-series shows a seasonal pattern. Models that follow the equation (2.3) can represent time-series with seasonality.

Since we cannot predict the value of the irregular fluctuations ε_t , there are three parameters that we must estimate: the level a, the trend b and the seasonal index F_t.

Winters (1960) suggested the following procedure that can be seen as a generalization and natural extension of the Holt procedure for a trend model. However, even though this method is intuitive, it is not optimal in the sense of minimizing mean-square forecast errors.

Updating procedure

The following equations are used for updating the parameters a, b and F_t:

$$\hat{a}_t = \alpha(x_t/\hat{F}_t) + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1}) \quad (2.18)$$

$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1} \quad (2.19)$$

$$\hat{F}_t = \gamma(x_t/\hat{a}_t) + (1-\gamma)\hat{F}_{t-T} \quad (2.20)$$

where:

α, β and γ are smoothing constants with value between 0 and 1,

$\hat{F}_{t-T}$ is the estimate of the seasonal index for the most recent (T periods earlier) equivalent period in the seasonal cycle β (Axsäter, 2006).

For example, if we want to calculate the seasonal index for the month of March, $\hat{F}_{t-T}$ would be the estimate of the seasonal index made for March of the past year.

It is important to understand that $a + bt$ represents the development of demand without taking into consideration the seasonal variations, and $x_t/\hat{F}_t$ is also an estimate of the deseasonalized demand.

Initialization

For time-series with both trend and seasonality, the initialization of the forecast procedure becomes more complicated than when only the trend is present. There are several initialization procedures, which consist on separating the effects of the trend and the seasonal effects. However, we are not going to show these methods in this paper, since it would make it too long. The objective of this thesis is not to dig in the theory behind the forecasting methods, but to analyse its performance with real data.

In Silver, Pyke and Thomas (2017), the reader might find a simple and easy to understand explanation of the initialization of this forecasting method.

Forecast

The equation that gives the forecasts for the Holt-Winters Exponential Smoothing is:

$$\hat{x}_{t,t+\tau} = (\hat{a}_t + \hat{b}_t \tau)\hat{F}_{t+\tau-T} \qquad (2.21)$$

where:

$\hat{x}_{t,t+\tau}$ is the forecast made at the end of period t, of the demand in period $t+\tau$,

$\hat{F}_{t+\tau-T}$ is the most recent estimate of the seasonal index for period $t+\tau$. In other words, it is the estimate of the seasonal index for period $t+\tau$ made T periods before.

2.6.5. Box-Jenkins: ARIMA models

ARIMA models provide a different approach to time-series forecasting. They aim to describe the autocorrelations in the data. ARIMA is an acronym with the first two letters, AR, standing for Autoregressive, the following I standing for Integration, and the last two letters, MA, standing for Moving Average.

The ARIMA models are appropriate for modelling time-series with trend characteristics, random walk processes, and seasonal and non-seasonal time-series. This family includes models that are combinations of autoregressive and moving average processes for stationary and non-stationary time-series.

A non-seasonal ARIMA model contains three terms: the autoregressive term (AR), the differencing, if needed, of the time-series (I), and the moving average term (MA).

A seasonal ARIMA model, SARIMA, has three additional terms: the seasonal autoregressive term, the seasonal differencing and the seasonal moving average term.

Autoregressive (AR) term

"The autoregressive component refers to the use of past values in the regression equation for the time-series" (Dalinina, 2017).

The description of the different terms of the ARIMA models is similar to the one used by Chatfield (2000).

A time-series $\{X_t\}$ is said to be an autoregressive process of order p, $AR(p)$, if it is a weighted linear sum of the past p values plus a random error:

$$X_t = \varphi_1 X_{t-1} + \varphi_2 X_{t-2} + \cdots + \varphi_p X_{t-p} + \varepsilon_t \quad (2.22)$$

where:

ε_t denotes a purely random process with zero mean and variance σ_ε^2,

φ_p are parameters of the model.

The auto-regressive parameter p specifies the number of lags used in the model. For example, $AR(2)$ or equivalently, $ARIMA(2,0,0)$, is represented as:

$$X_t = \varphi_1 X_{t-1} + \varphi_2 X_{t-2} + \varepsilon_t \quad (2.23)$$

Moving average (MA) processes

The moving average component represents the error of the model as a combination of the previous error terms, ε_t.

A time-series $\{X_t\}$ is said to be a moving average process of order q, $MA(q)$, if it is a weighted linear sum of the last q random errors:

$$X_t = \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \cdots + \theta_q \varepsilon_{t-q} \quad (2.24)$$

where

ε_t denotes a purely random process with zero mean and variance σ_ε^2,

θ_q are parameters of the model.

The order q determines the number of terms to include in the model. For example, $MA(2)$ or equivalently, $ARIMA(0,0,2)$, is represented as:

$$X_t = \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} \quad (2.25)$$

Integrated part (I)

The I in ARIMA is for "integrated". When the time-series is non-stationary, it is necessary to make some transformations to the data in order to make it stationary. These transformations are known as differencing.

The first differences of the time-series are calculated subtracting to each value, the previous value in the time-series. If this differencing process is performed twice, we have the second differences. Note that the second difference is not the difference from two periods ago. It is rather the first difference of the first difference.

If d=0: $X_t = x_t$,

If d=1: $X_t = x_t - x_{t-1}$,

If d=2: X_t=$(x_t - x_{t-1}) - (x_{t-1} - x_{t-2}) = x_t - 2x_{t-1} + x_{t-2}$.

Therefore, x_t are the values of the time-series and X_t are the values of the d^{th} difference of x_t.

A time-series can be differenced many times in order to obtain stationarity. However, differencing tends to introduce negative correlation and it should be done carefully.

ARIMA (p, d, q)

Given the d^{th} difference of a time-series of data x_t, X_t, where X_t are real numbers and t is an integer index, an $ARIMA(p, d, q)$ model is given by:

$$X_t - \varphi_1 X_{t-1} - \varphi_2 X_{t-2} - \cdots - \varphi_p X_{t-p} \qquad (2.26)$$
$$= \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \cdots + \theta_q \varepsilon_{t-q}$$

The moving average parameters, φ_p, are defined so that their signs are negative in the equation, following the convention introduced by Box and Jenkins (1976).

Order of the ARIMA model

When fitting non-seasonal ARIMA models, there are five parameters that we have to estimate: the coefficients φ and θ, the orders p and q, and the number of times that the time-series must be differenced, d. Nevertheless, the main difficulty is to choose the right order of the model rather than the value of the coefficients.

Fortunately, nowadays we have many computer programmes that will choose the optimal order of the ARIMA (p, d, q) model. However, the forecaster should know how this is calculated, in order to determine if the model is stable or not.

The key to determine the order of the different components is to analyse the Autocorrelation Function (ACF) plot, and the Partial Autocorrelation Function (PACF) plot. Since this is not a guide for forecasting, it is assumed that the reader is familiar with the ACF and the PACF concepts.

The procedure of analysing the ACF and PACF for determining the different orders of ARIMA models is long and it is not going to be explained here. The most used reference is Box and Jenkins (1976).

Particular cases of ARIMA model

Some particular cases of ARIMA models are equivalent to other models that we have explained before. Here are some examples:

ARIMA(1,0,0) – first-order autoregressive model,

ARIMA(0,1,0) – random walk,

ARIMA (1,1,0) – differenced first-order autoregressive model,

ARIMA (0,1,1) – simple exponential smoothing,

ARIMA(0,2,1) or ARIMA(0,2,2) without constant – linear exponential smoothing,

ARIMA (1,1,2) without constant – damped-trend linear exponential smoothing.

Seasonal ARIMA models (SARIMA)

Seasonal ARIMA models, also called SARIMA models, are used for modelling time-series that contain a seasonal pattern.

The idea is the same than for non-seasonal ARIMA models. However, there are six terms instead of three in these models:

p: order of the non-seasonal autoregressive component,

d: number of non-seasonal differences applied to the time-series,

q: order of the non-seasonal moving average component,

P: order of the seasonal autoregressive component,

D: number of seasonal differences applied to the time-series,

Q: order of the seasonal moving average component.

Therefore, a SARIMA model can be written as SARIMA (p, d, q)x(P, D, Q)s, where s is the number of seasons in the time-series.

2.6.6. Combination of forecasts

Some studies by Makridakis et. al. (1982), Makridakis et. al. (1993) and Makridakis and Hibon (2000) (cited in Silver, Pyke and Thomas, 2017, p. 125) prove that in most cases, a combination of some forecasting methods gives better results than any of the methods itself.

These studies involved 1.001 time-series of different natures, and different forecasting intervals (monthly, quarterly and annual) were used.

One of the conclusions was, as it has been said before, that it is recommended to combine the results of different forecasts. However, it is not necessary to try to find optimal weights when combining different forecasts. In most cases, a simple average of two or three methods gives good results.

2.7. Analysis of the accuracy of the methods used

We have reviewed different methods for forecasting the future values of a time-series. Some methods are simple and others are more sophisticated, and each of them will give very different results on each time-series.

It is important to analyse the accuracy of the forecasts, in order to know which method gives better results. The description of some of the most used performance indicators can be found below. The description of the performance indicators is similar to the one described by Armstrong (2001) and by Silver et. al. (2017)

2.7.1. Mean squared error (MSE)

The mean squared error (MSE) measures the average of the squared errors. It requires a target prediction (the real demand) along with a predictor or estimator (the forecast). The MSE is defined as the average of squares of errors, where the errors are the difference between the real demand and the forecast.

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(x_t - \hat{x}_{t-1,t})^2 = \frac{1}{n}\sum_{t=1}^{n}(e_{t-1,t})^2 \qquad (2.27)$$

where:

x_t is the actual observed demand in period t,

$\hat{x}_{t-1,t}$ is the one-period-ahead forecast,

$e_{t-1,t}$ is the error in the forecast made in period $t-1$ for period t,

n is the number of forecasts.

2.7.2. Root mean squared error (RMSE)

The root mean squared error (RMSE) is defined as the square root of the MSE.

$$RMSE = \sqrt{MSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(x_t - \hat{x}_{t-1,t})^2} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(e_{t-1,t})^2} \qquad (2.28)$$

where:

x_t is the actual observed demand in period t,

$\hat{x}_{t-1,t}$ is the one-period-ahead forecast,

$e_{t-1,t}$ is the error in the forecast made in period $t-1$ for period t,

n is the number of forecasts.

The RMSE is the square root of the average of squared errors. Since the errors are squared before they are averaged, the RMSE gives a higher weight to large errors.

2.7.3. Mean absolute error (MAE)

The mean absolute error (MAE) measures the average magnitude of the errors in a set of forecasts, without considering their direction. It takes the absolute value of the forecast errors and averages them over all forecast time periods.

For the n periods of data, the estimate of the MAE for one-period-ahead forecasts is presented in equation (2.29).

$$MAE = \frac{1}{n}\sum_{t=1}^{n}\left|x_t - \hat{x}_{t-1,t}\right|^2 = \frac{1}{n}\sum_{t=1}^{n}\left|e_{t-1,t}\right|^2 \tag{2.29}$$

where:

x_t is the actual observed demand in period t,

$\hat{x}_{t-1,t}$ is the one-period-ahead forecast,

$e_{t-1,t}$ is the error in the forecast made in period $t-1$ for period t,

n is the number of forecasts.

The MAE was useful in the past due to its computational simplicity. However, with the widespread availability of computers, the MAE has become of less practical importance. Nevertheless, it is still often used because it is intuitive (Silver, Pyke and Thomas, 2017 p. 105).

2.7.4. Mean absolute percentage error (MAPE)

The mean absolute percentage error (MAPE) is also widely used because it is an intuitive measure of variability. It is expressed as a percentage, thus it is generally not affected by the magnitude of the demand values (Silver, Pyke and Thomas, 2017 p. 106).

$$MAPE = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{x_t - \hat{x}_{t-1,t}}{x_t}\right| = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{e_t}{x_t}\right| \tag{2.30}$$

where:

x_t is the actual observed demand in period t,

$\hat{x}_{t-1,t}$ is the one-period-ahead forecast,

$e_{t-1,t}$ is the error in the forecast made in period $t-1$ for period t,

n is the number of forecasts.

Nevertheless, the MAPE has same disadvantages when it is used in forecasting. When forecasting the future demand or production volume, it is possible that, for some period, the value of the time-series is zero. In those cases, the MAPE cannot be used since there would be a division by zero.

In addition, there is no upper limit to the percentage error. The percentage error of forecasts that are too low (the actual demand is higher than the forecast), cannot exceed 100%. However, when the actual demand is lower than the forecast, the percentage error could be higher than 100%. For example, if our forecast for a particular period was 250 units, and the actual demand in that period is 800 units, the error made in that period would be an error of 220%.

The MAPE is also not recommended when the values of the time-series are very low. For example, if our forecast for a certain period is 1 unit, and the real demand for that period is 2 units, we are making an error of 1 unit. However, if we count that error as a percentage, it is a 100% error, which could be misleading (Silver, Pyke and Thomas, 2017 p. 106).

3. RESEARCH METHODOLOGY

3.1. The Company

The research has been conducted using the data provided by a manufacturing company. Since the company prefers to remain anonymous, in this paper we will use "the Company" instead of its real name. The name of the products has also been changed, since the data provided is confidential information.

The Company is located in an Eastern European country and it produces construction material products, only for the national market. Therefore, the demand of the products is strongly related to the national economy.

The Company produces products for tiling, as well as new construction products and construction renovation products. They produce a wide variety of products, such as tile adhesives, grouts, floor-levelling compounds, silicones, decorative wall plasters and waterproofing products.

The country where the Company is located suffers very cold and strong winters, with considerable amounts of snow and extremely low temperatures. For this reason, the construction activity is much lower in the winter than in the summer. Since the demand of the Company's products is highly related to the construction activity in the country, it strongly depends on the season of the year and the weather conditions. As a result, the demand and the production volume of the Company is expected to follow a clear seasonal pattern.

3.2. The data

The Company has provided the production volume of all their products from the year 2010 to the year 2017. In general, manufactures use the demand or the sales volume to make their forecasts. Companies need to have a prediction of how much they are going to sell in the next week, month, quarter or year, in order to know how much they should produce, how many raw materials they should buy, or how many resources they will need.

According to the operations manager of the company, the data of the production and the sales are almost identical, since they use a Make to Order manufacturing process. This means that they start producing a product once they receive the customer's order. Therefore, forecasting the production volume or forecasting the sales will give almost the same results. The data that the Company has provided is the production volume, thus we will forecast that variable and not the sales or the demand. The production volume is given in tons.

3.2.1. Data analysis

The first step in forecasting is always to analyse the data and to organise it in order to understand what might be improved. The data given by the company was very detailed and therefore it was difficult to find any pattern or discontinuity. The files with the data contain the production volume of each product for each day from the 1st January 2010 to 31st December 2017. Since the Company produces a high number of products, the data provided contains more than 55.000 values.

3.2.2. Aggregating the data

The second step is to aggregate the data. For obvious reasons, it is no very useful to forecast the production for every day during a whole year. For a manufacturer who prepares the production plan for the next year, it is not important to know how much they will produce on every single day. It might be more useful to have a prediction of how much they will produce on each month, or on each quarter. That is why the next step is to aggregate the data over time. We have analysed all the data provided by the Company and aggregated it monthly for every year.

It is also recommended to aggregate the products into different categories. The Company produces a variety of products that can be categorized into product groups or families. Individual products in the same product group generally have some common characteristics, such as the same demand pattern or relatively stable product mix. The general recommendation in those cases is to aggregate the demand or production volume of those products and forecast it together as a single product.

Since the Company produces more than 190 different products, it would be impossible in this paper to analyse all of them. After having a conversation with the manager about this issue, we decided to analyse only the 15 products with a higher demand.

Additionally, we decided that it would be better to aggregate these products. We have divided the products into four different categories, according to their properties. One of the main reasons why we decided to aggregate the demand is that some of the products are very similar, thus their demand is related to each other. In Figure 3.1. there is an example of this situation.

The graphs in Figure 3.1. represent the production volume of the products *Product 1* and *Product 2*. The production of *Product 1* started in period 43, which corresponds to July 2013. Around that time, there is a clear drop in the production of *Product 2*. The reason of this decrease is that the two products are very similar, and *Product 1* took part of the demand from *Product 2*. They have a very similar composition, similar price and the company considers them as one group, thus their demand is closely related.

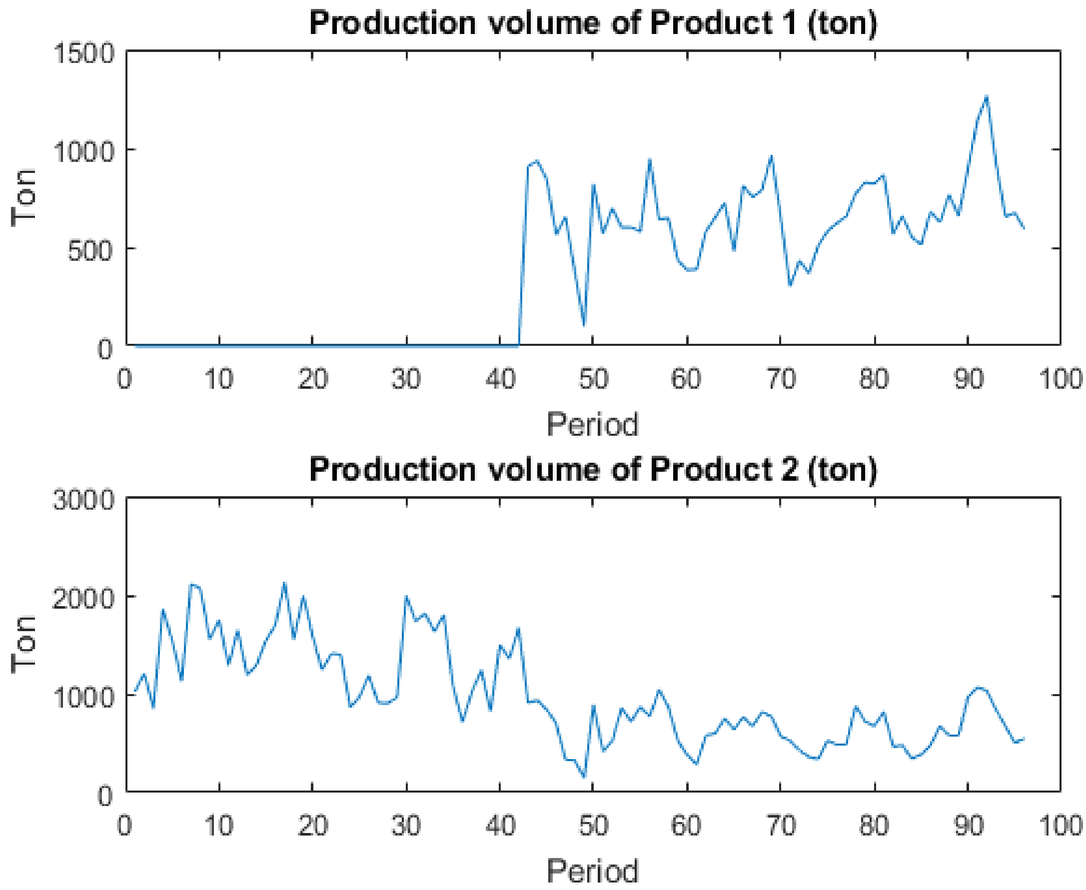

Figure 3.1. Production volume of Product 1 and Product 2 from 2010 to 2017

If we try to analyse these two products separately, it is difficult to explain why the production of *Product 2* suddenly drops between the years 2013 and 2014. However, if we aggregate these two products, the demand looks more stable and easier to analyse (Figure 3.2.).

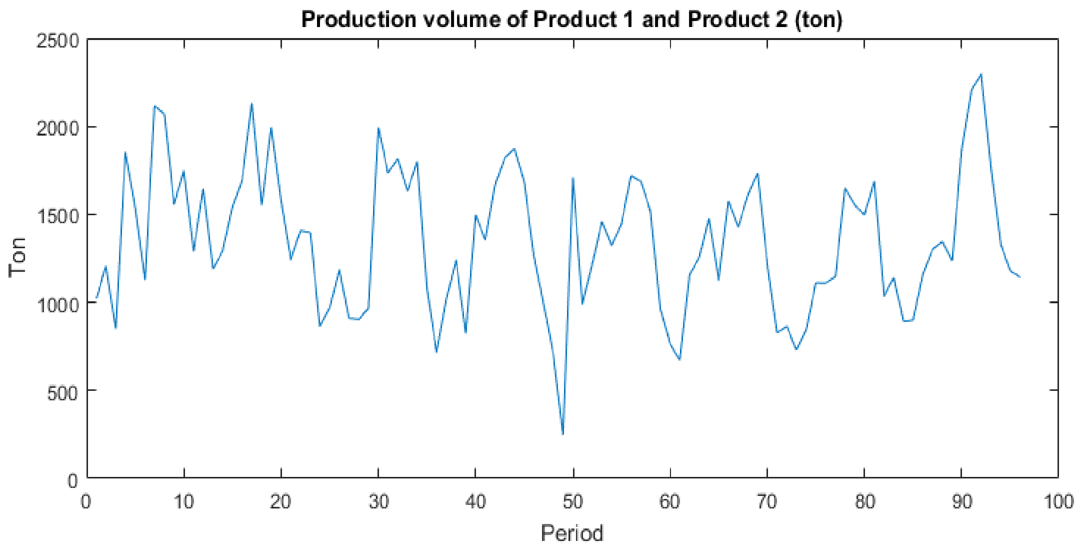

Figure 3.2. Aggregated production volume of Product 1 and Product 2 from 2010 to 2017

In general, the aggregated demand of these two products is relatively constant from the year 2010 to the year 2017, with a clear seasonal pattern with higher production in the middle of each year, which corresponds to the summer. The production in the year 2017 is higher than in the past data, and there is a drop in the production in January 2014. All these singularities in the data will be explained afterward.

In order to respect the privacy of the Company, neither the real name of the products nor the type of product are given in this thesis. Thus, the 15 products have been aggregated into 4 categories, which we have named *Product A*, *Product B*, *Product C* and *Product D*.

- **Product A**: Product 1, Product 2, Product 3, Product 4 and Product 5.
- **Product B**: Product 6, Product 7, Product 8 and Product 9.
- **Product C**: Product 10, Product 11, Product 12 and Product 13.
- **Product D**: Product 14 and Product 15.

For each one of the 15 products, and for each month, we have summed the production volume of each day. Here, in Table 3.1. is an example of the production volume of one of the aggregated products (*Product A*).

Table 3.1. Production volume of Product A aggregated monthly from 2010 to 2017

		PRODUCTION VOLUME OF PRODUCT A (ton)							
		Year							
		2010	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**	**2017**
Month	**JANUARY**	1172,65	1451,4	1231,8	1380,1	521,7	977,7	1282,5	1129,4
	FEBURARY	1384,8	1723,1	1527,8	1663,85	2637,4	1666,2	1180	1546,3
	MARCH	1110,7	2014	1323,75	1286,8	1622,4	1772,2	1504,3	1558,7
	APRIL	2045,2	2117,7	1254,75	2132,6	1945,1	2079,2	1547	1870
	MAY	1854,7	2494,5	1567,83	1957,5	2083,78	1658,1	1498,7	1784,8
	JUNE	1440,6	2011,6	2617,7	2473	2139,5	2303,2	2335	2502
	JULY	2561,7	2586,2	2409,22	2595	2318,5	2100,7	2306	3122,8
	AUGUST	2409	2031,85	2405,2	2878,1	2662,6	2387,5	2136,9	2987,6
	SEPTEMBER	1940	1773,1	2243,7	2485,5	2369,3	2427,9	2251,9	2508,9
	OCTOBER	2086,3	1815,3	2358,4	2044,8	2390,2	1858,1	1596,2	1865,7
	NOVEMBER	1738,5	1647,1	1597,2	1861,1	1434,6	1438,9	1470,4	1691,4
	DECEMBER	1866,2	1103,9	1127,2	1115,7	1209,9	1155,2	1266,4	1508,4

3.2.3. Forecasting methodology

The methods that we will use for forecasting the production volume are:

- moving averages,
- simple exponential smoothing,
- Brown's linear exponential smoothing,

- Holt's linear exponential smoothing,
- Holt-Winter's exponential smoothing,
- ARIMA – Box-Jenkins method.

Additionally, we will calculate an extra forecast as a combination of the two forecasts that give the best results. Research has proved that in most cases, this combination gives better forecasts than any method used independently.

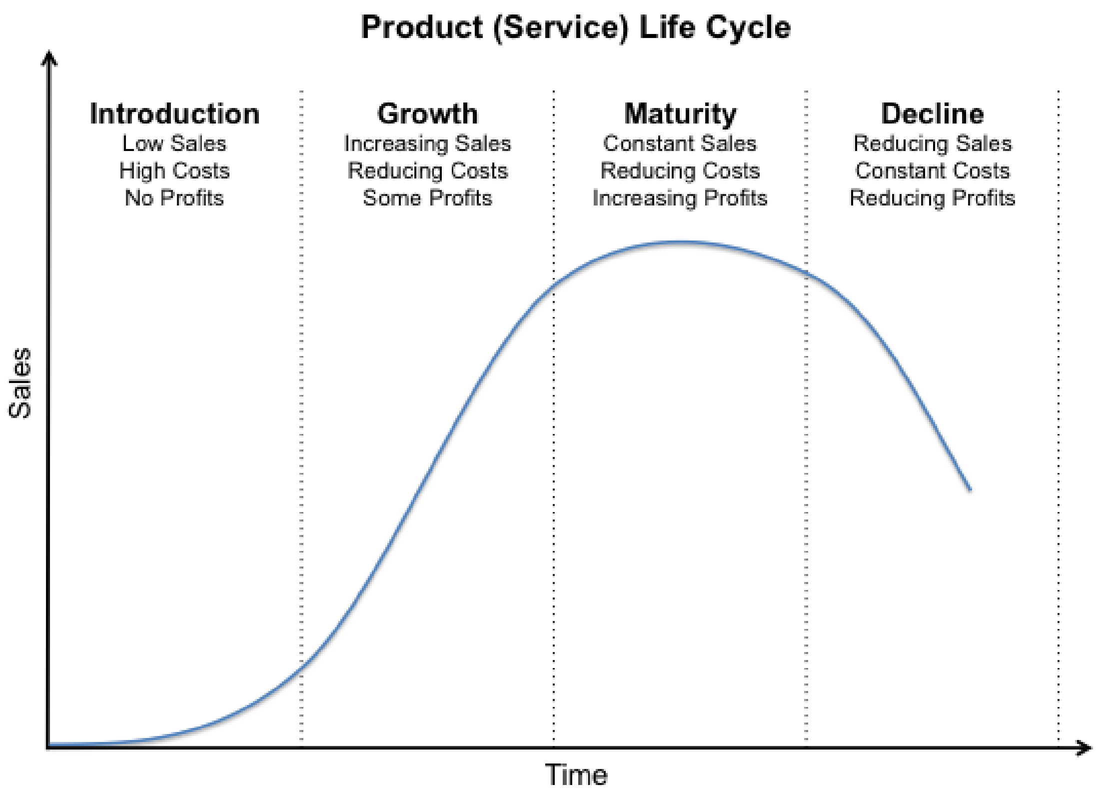

Figure 3.3. *Product life cycle* (Business Set Free Ltd, 2013)

All these methods are quantitative methods based on time-series analysis. The only variable used for making forecasts is the historical data. One of the main decisions before making any forecast is to decide how much data from the past we will use. The demand of every product changes along the different life cycle stages. For instance, it is not recommended to use de data from the growth stage (see Figure 3.3.) to make forecasts for the decline stage.

3.3. Forecasting horizon

Another crucial decision is to choose the forecasting horizon. Since we have aggregated the data month by month, there are three possible forecasting horizons: monthly, quarterly and yearly. Longer forecasting horizons are not recommended with time-series analysis methods. These methods give good results only for short-term or medium-term forecasts. Since they do not include the effects of cyclic variations, it would be dangerous to make long-term forecasts with them.

The Company also provided with their monthly, quarterly and yearly predictions for the year 2016. These predictions are based in the opinion of experts and managers of the Company, and no quantitative methods have been used. In Table 3.2., the reader can find the predictions made by the Company of *Product A* for every month in the year 2016.

The yearly predictions were made at the end of the year 2015, when the real production volume of December 2015 was already known.

The quarterly predictions were made before the beginning of every quarter of the year 2016. At the end of December 2015, the Company made the predictions for January, February and March 2016. At the end of March 2016, the Company made the predictions for April, May and June 2016. The same for the other two quarters.

The monthly predictions were made for every month at the end of the previous month. According to the operations manager of the Company, the monthly predictions are more a production plan than actual predictions. Therefore, it is not useful to make monthly forecasts, since this information would not be relevant for the Company.

Table 3.2. *Predictions by the Company of the production volume of Product A*

	Product A			
	Production Volume	Predictions		
		Monthly	Quarterly	Yearly
January	1282,5	1184,5	1184,5	1.301,3
February	1180	1501	1496,5	1.633,5
March	1504,3	1588	1588	1.748,5
April	1547	1671,2	1671,2	1.891,0
May	1498,7	1621,5	1593,2	2.004,7
June	2335	1857	1737	2.122,4
July	2306	2124,5	2124,5	2.347,5
August	2136,9	2287	2167	2.515,5
September	2251,9	2016,5	2042,5	2.352,5
October	1596,2	1722,5	1722,5	2.030,4
November	1470,4	1379,5	1462	1.726,5
December	1266,4	1202,5	1336	1.594,3
	RMSE	207,73	223,10	316,50
	MAE	172,98	161,68	276,50
	MAPE	10,13%	9,38%	17,98%

Comparing the quarterly and yearly predictions for *Product A*, it is clear that the quarterly predictions are more accurate. This situation also happens for *Product D*, while yearly predictions are more accurate than quarterly predictions for *Product B* and *Product C*. However, according to the operations manager, it would be more beneficial for the Company to improve the yearly forecasts, since they are the base for all plans, such as the budget or the numbers presented to stakeholders.

For the reasons mentioned above, we decided that the forecasting horizon should be one year. At the beginning of the year 2016, we have calculated the forecasts of the production volume of each month for the whole year.

3.4. Procedure followed

For each one of the products, we first plot the data in order to find and correct irregularities. Real production data might contain some discontinuities because of reparations of machines, labour strikes, or changes in the production cycle for many different reasons.

Once the irregularities are corrected, we choose how much historical data we use for making the forecasts. If there is a clear growing trend and, after a while, the trend starts being descending, it is not recommended to include the data that is older than the change in the trend.

After this, for each one of the methods, we use the past data to build the model that fits the best our data. We use data from the years 2010 to 2015 to build the model.

Once we have built the best model, we use it for forecasting the monthly production for every month in the year 2016. We then analyse the accuracy of the forecasts with different performance indicators. Additionally, we calculate and evaluate a combination of the two forecasts that give the best results.

Finally, we compare the results with the predictions made by the Company, in order to discover if the forecasts that we have calculated over perform these predictions.

4. EXPERIMENTS

In this thesis, we only show the analysis and the results of every step of *Product A*. For the other three products, only the value of the performance indicators is shown. The production volume of the four products has a very similar pattern and the procedure followed to make the forecasts is very similar. It would be too long and unnecessary to show the analysis of every product, and to explain how the forecasts have been calculated.

The *Product A* is the aggregated production volume of five different products (*Product 1*, *Product 2*, *Product 3*, *Product 4* and *Product 5*) which have a similar composition. According to the manager of the Company, they can be aggregated together to make the forecasts.

4.1. Analysis and correction of the data

The first step is to plot the production volume of Product A from the year 2010 to the year 2017. It is important to notice that in Figure 4.1., the Y-axis starts from 500 tons, and not from zero.

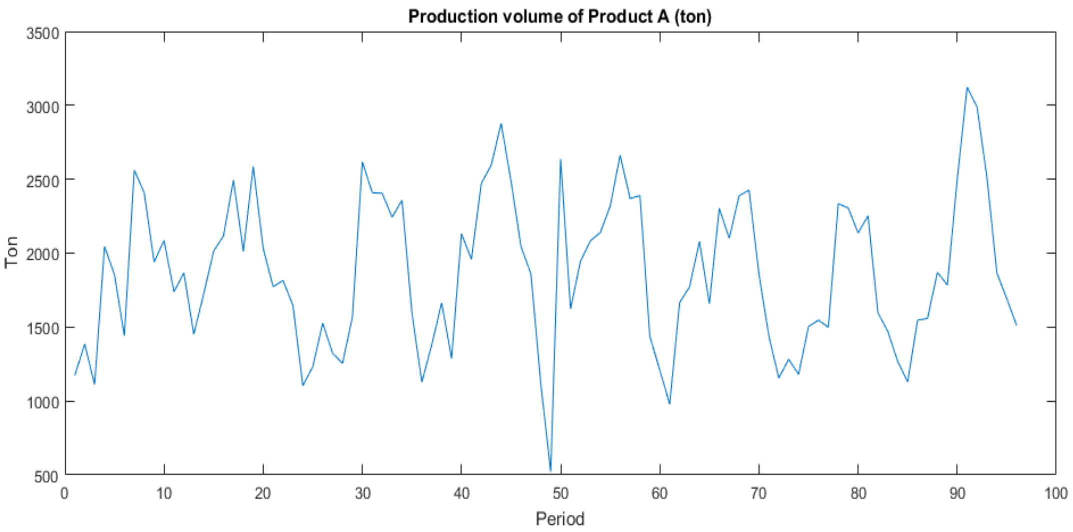

Figure 4.1. *Production volume of Product A from 2010 to 2017*

There is a sudden drop in the production in period 49, which corresponds to January 2014, and an unusually high level of production in period 50, which corresponds to February 2014. The reason of this irregularity in the production volume is that in January 2014 the Company changed some of the machines used for production. Therefore, the demand of January could not be produced in that month, and it was produced in February. The methods that we have used for forecasting do not understand these kind of irregularities in the time-series. These values should be corrected in order to have a more stable time-series and forecasts that are more accurate.

The production volume in January 2014 and in February 2014 needs to be readjusted. We have estimated the “correct” production volume for those months based on the production volume in January and February in the other years.

The graph of the production volume of *Product A* after this correction in represented in Figure 4.2.

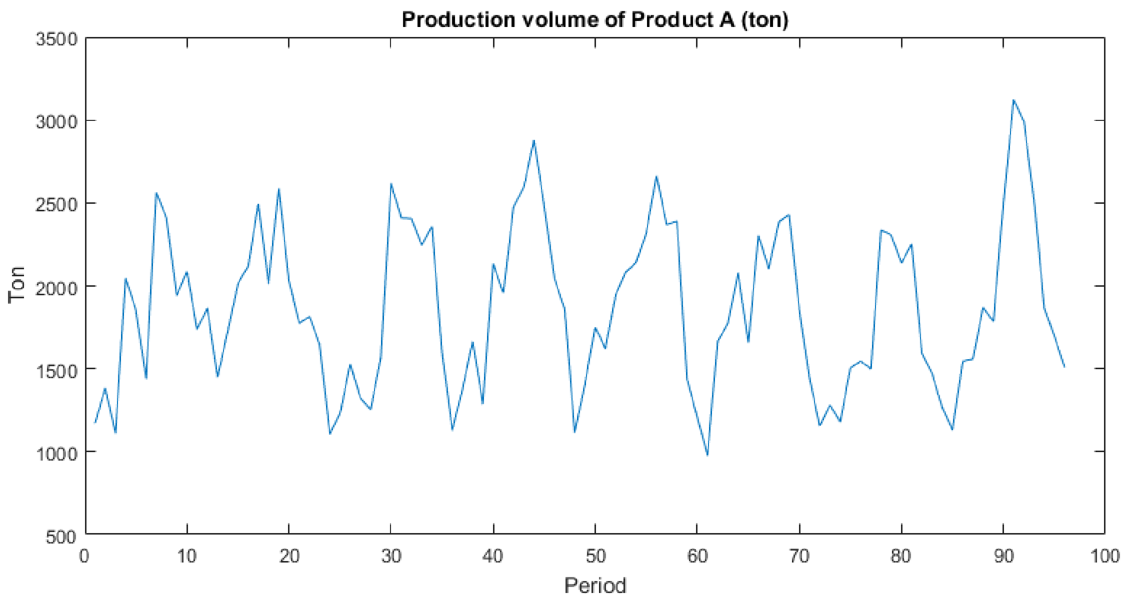

Figure 4.2. *Corrected production volume of Product A from 2010 to 2017*

Let us now analyse the different components in the time-series. In the first three years (2010–2012) there seems to be no trend. Without considering the seasonality, the production volume seems to be constant. However, in 2013, there seems to be an increase in the production volume and, after that, a decreasing trend until the year 2016. Again, there is a sudden increase of the production volume in the year 2017. According to the Company, this sudden increment in 2017 is strongly related with the recovery of the economy in the country where the Company is located – the market of the company is national, hence they only sell their products in the country where they produce it.

Additionally, a clear seasonality repeats every year. There is a higher production volume in the middle of every year, which corresponds to the summer or warm season. The products that the Company produces are used in construction; hence, their demand relies strongly on the weather conditions. In the country where the Company is located, the extremely low temperatures in the winter have a strong impact in the construction business. During the cold months, the activity of these companies decreases. Therefore, the demand and the production of the Company's products is also much lower during the winter.

The next step is to decide for which year we will make the forecasts. The methods that we are going to use are based on time-series analysis, which means that the only variable that they use for forecasting is the historical data. These methods make forecasts based on past data, thus the forecast is just a "projection" for the future of what has happened before. It is not possible to predict sudden changes in the demand, such as changes in the trend, sudden increases or sudden decreases in the production volume.

Using these methods for forecasting the production volume in the year 2017 will not give satisfactory results. Managers should take into account other exogenous variables that can affect the demand, such as the economy of the country, or the weather forecast. If a sudden change in the economy of the country is expected, the forecasts made with time-series analysis methods should be used very carefully, since the results might be misleading.

On the other hand, when the economy of the country and other external factors are stable, time-series analysis can give very adequate results. In the production volume of *Product A*, the production is stable from 2010 to 2016, but in the year 2017 there is an unexpected increase in the production. The objective of this thesis is to prove that time-series analysis methods are useful for stable time-series. Consequently, it would be interesting to make the forecasts for the year 2016 since the results will be more precise than for the year 2017.

Once we have decided that we will make the forecasts for the year 2016, we have to decide how much of the available historical data we will use. We have data from the year 2010, but in the year 2013, there is a change in the direction of the trend. From 2010 to 2012, it seems that there is no trend, but after the sudden increase in the production volume in 2013, there is a clear descending trend. For this reason, we will only use the production volume of the years 2013, 2014 and 2015 to make the forecasts for the year 2016.

This division of the data is a common practise in forecasting. In this case, the data from 2013 to 2015 is called *Training Set*, and it is used for building the model. The data of 2016 is called *Test Set*, and it is used to analyse the accuracy of the forecast method used. The results of the Test Set are used to choose which method is the best, and afterwards this method is used for making forecasts in the future.

Let us now analyse the production volume of the years 2013, 2014 and 2015. The first step is to plot the data to have a more clear idea of how the production volume of *Product A* it looks.

Looking at the Figure 4.3., we can have a general idea of how the different time-series components look. For instance, an obvious seasonal pattern repeats every year, with higher seasonal indexes in the middle of each year. In addition, it seems that there is a descending trend.

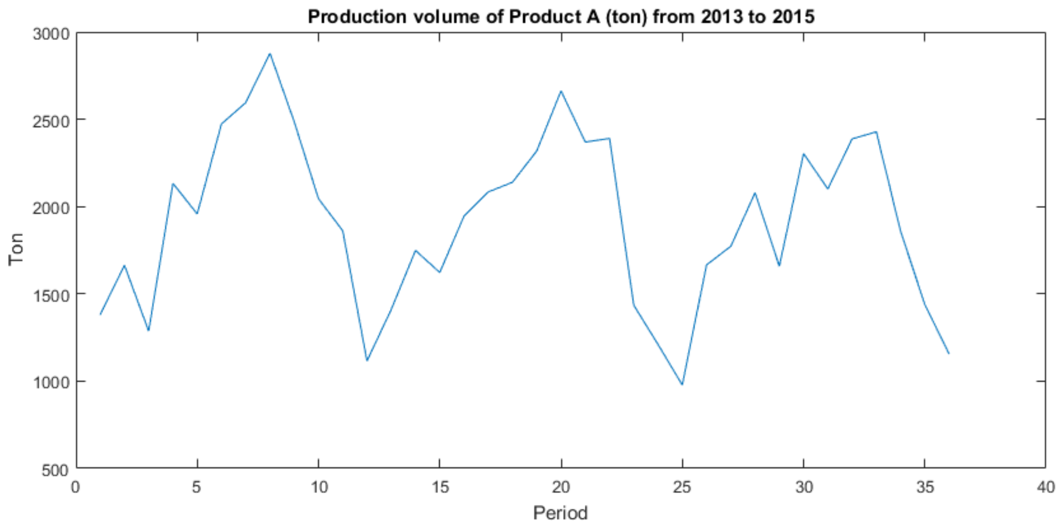

Figure 4.3. *Corrected production volume of Product A from 2013 to 2015*

4.2. Seasonal decomposition

We have used the Seasonal Decomposition option in Statgraphics to separate the different components of the time-series. This procedure applies a multiplicative seasonal decomposition to the data. The reason why we have used multiplicative seasonality instead of additive seasonality is that it gave better results after making different trials. The purpose of the decomposition is to separate the time-series into trend-cycle, seasonal, and random components. The Table 4.1. shows each step of the seasonal decomposition.

Table 4.1. *Seasonal decomposition*

Year	Month	Period	Data	Trend-Cycle	Seasonality	Irregular	Seasonally Adjusted
2013	**January**	1	1380,1				2267,6
	February	2	1663,85				1885,03
	March	3	1286,8				1455,63
	April	4	2132,6				2028,63
	May	5	1957,5				1998,46
	June	6	2473				2107,07
	July	7	2595	1990,77	1,30	104,874	2087,8
	August	8	2878,1	1995,57	1,44	102,433	2044,12
	September	9	2485,5	2013,08	1,23	100,556	2024,28
	October	10	2044,8	2019,25	1,01	90,4663	1826,75
	November	11	1861,1	2016,7	0,92	111,132	2241,2
	December	12	1115,7	2008,07	0,56	93,894	1885,46
2014	**January**	13	1410,58	1982,65	0,71	116,898	2317,68
	February	14	1748,52	1962,15	0,89	100,959	1980,96
	March	15	1622,4	1948,33	0,83	94,1964	1835,26
	April	16	1945,1	1957,88	0,99	94,5037	1850,27
	May	17	2083,78	1954,5	1,07	108,845	2127,38
	June	18	2139,5	1940,66	1,10	93,9332	1822,92
	July	19	2318,5	1926,54	1,20	96,8231	1865,34
	August	20	2662,6	1905,08	1,40	99,2642	1891,06
	September	21	2369,3	1907,89	1,24	101,14	1929,65
	October	22	2390,2	1919,72	1,25	111,23	2135,31
	November	23	1434,6	1907,57	0,75	90,5651	1727,59
	December	24	1209,9	1896,65	0,64	107,803	2044,65

Year	Month	Period	Data	Trend-Cycle	Seasonality	Irregular	Seasonally Adjusted
2015	**January**	25	977,7	1894,4	0,52	84,7988	1606,43
	February	26	1666,2	1873,86	0,89	100,738	1887,7
	March	27	1772,2	1864,84	0,95	107,5	2004,71
	April	28	2079,2	1845,11	1,13	107,193	1977,83
	May	29	1658,1	1823,12	0,91	92,8516	1692,8
	June	30	2303,2	1821,02	1,26	107,764	1962,4
	July	31	2100,7				1690,11
	August	32	2387,5				1695,68
	September	33	2427,9				1977,37
	October	34	1858,1				1659,95
	November	35	1438,9				1732,77
	December	36	1155,2				1952,21

The “Trend-Cycl”e column of Table 4.1. shows the results of a centered moving average of length 12 applied to the data. The reason why we need to apply a centered moving average is that the number of periods per season (12 months) is an even number. With this procedure, we smooth the data and remove the seasonality. The values in the “Trend-Cycle” column are an average that includes one value of each one of the 12 different periods. This way, we have the data smoothed and “free” of seasonality (Figure 4.4). We can understand the graph in Figure 4.4. as how the time-series would look if there were not a seasonal pattern and irregular fluctuations.

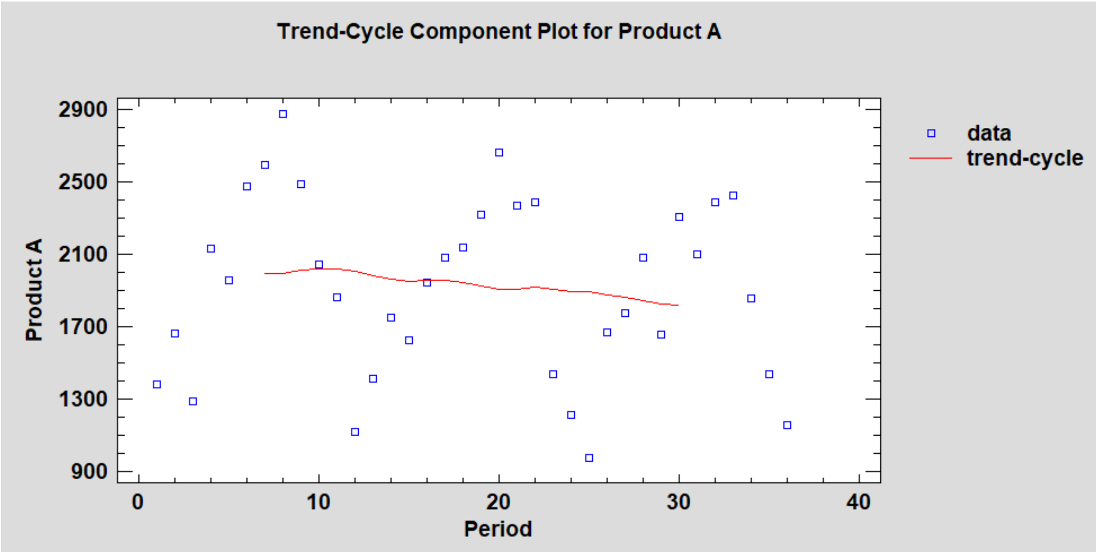

Figure 4.4. *Trend-Cycle component plot for Product A*

The "Seasonality" column is calculated dividing the "Data" column by the "Trend-Cycle" column. It represents, for each period, how much the data differs from the average. If this value is greater than one, it means that for this period the production volume is higher than the average. On the other hand, if this value is between zero and one, it means that for this period the production volume is lower than the average.

Seasonal indexes are computed for each season by averaging the ratios across all observations in that season, and scaling the indexes to make an average season equal 12.

In Table 4.2., the value in the column Index in January is the average of the values in column "Seasonality" of Table 4.1., which corresponds to January. The same applies to every month. However, the sum of all these seasonal indexes is higher than 12. The seasonal indexes need to be normalised in order to make the sum equal 12.

Table 4.2. *Normalised seasonal indexes*

Season	Index	Normalised Index
January	0,614	0,609
February	0,890	0,883
March	0,892	0,884
April	1,060	1,051
May	0,988	0,980
June	1,184	1,174
July	1,253	1,243
August	1,420	1,408
September	1,238	1,228
October	1,129	1,119
November	0,837	0,830
December	0,597	0,592
Sum	12,102	12,00

In Figure 4.5. there is a graphical representation of the normalised seasonal indexes.

As we expected, during the warmest months (from June to October) the seasonal indexes are higher than one (also in April, and in May very close to one). In Figure 4.5., the seasonal indexes are multiplied by 100. During the coldest months (from November to March), the seasonal indexes are significantly lower than one, especially in December and January.

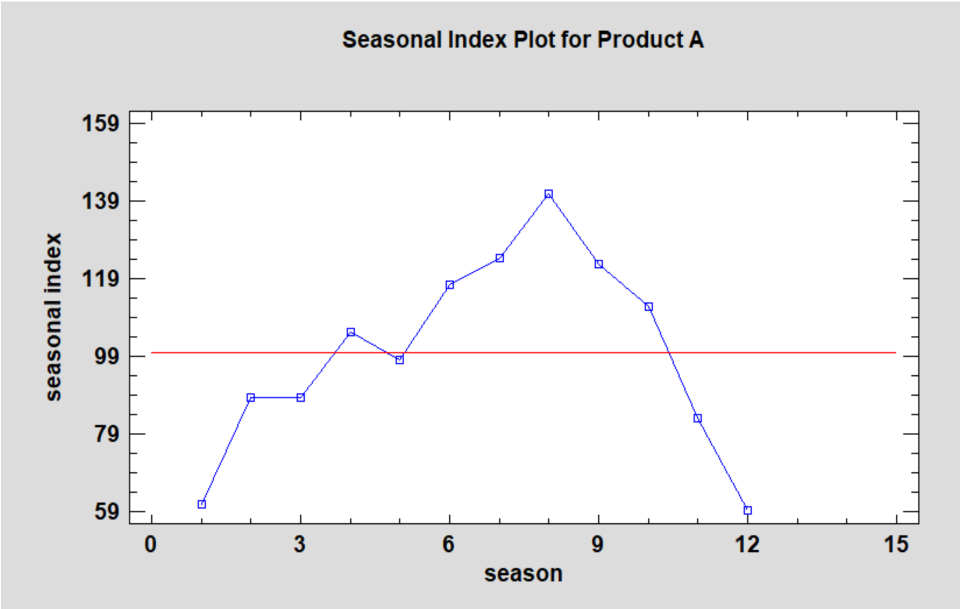

Figure 4.5. *Seasonal Index for Product A*

The “Data” is then divided by the “Trend-cycle” component and “Normalised Seasonal Indexes” to give the “Irregular” or residual component. This component is then multiplied by 100. The “Irregular” component is the residual variation remaining in the time-series after the trend-cycle and the seasonality have been extracted from the original time-series. Figure 4.6. shows a representation of the irregular component of the time-series.

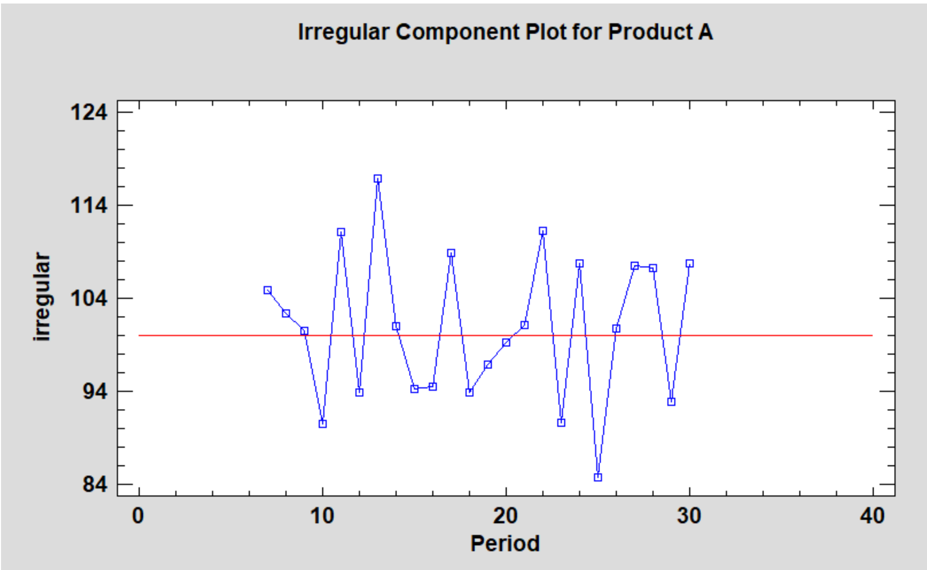

Figure 4.6. *Irregular component for Product A*

The last step in the seasonal decomposition is to obtain the seasonally adjusted data. We obtain these values by dividing the original time-series values, by the normalised seasonal indexes. These are the values shown in the last column of Table 4.1.

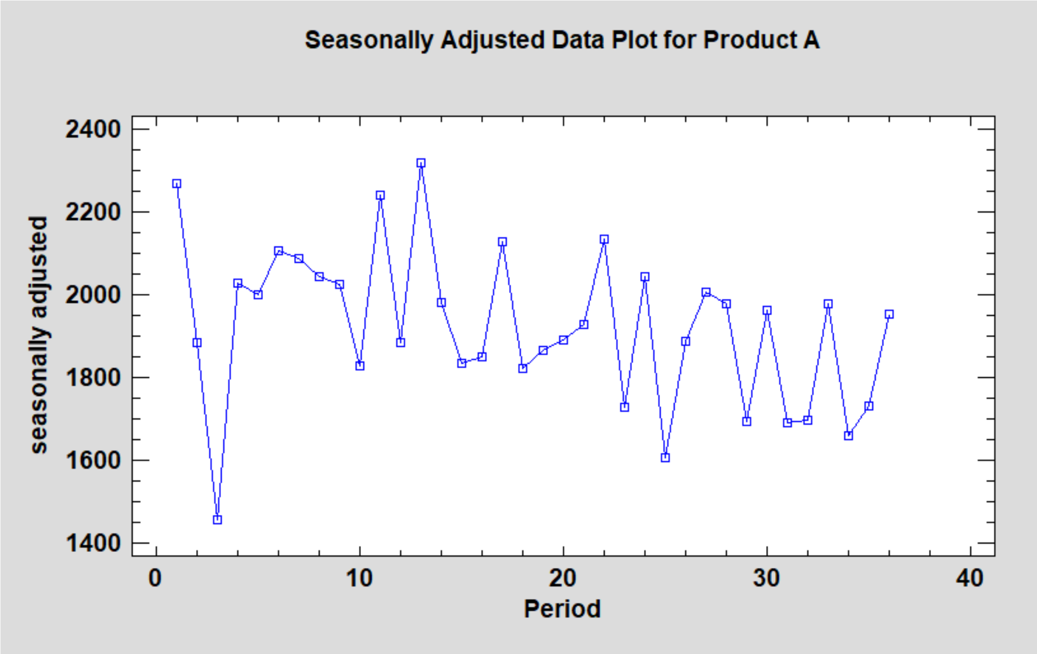

Figure 4.7. *Seasonally adjusted data plot for Product A*

Figure 4.7. represents the time-series without the seasonal component. Only the trend and the irregular component are present in the seasonally adjusted data.

4.3. Moving average

The first method that we have tried is the moving average. Although this method is very simple, it often gives satisfactory results. This method is intuitive and easy to apply.

The main decision that we have to make is the order or length of the moving average. The only way to find the optimal order of the moving average is to try different orders, and to compare the accuracy of the results.

We have applied in Statgraphics moving averages of different orders to the time-series from 2013 to 2015. We have analysed the accuracy with three different indicators: RMSE, MAE and MAPE. The indicator that we have used to decide which model is more accurate is an average of RMSE and MAE. We have also calculated MAPE because it is more intuitive and easy to understand, since it shows the average error in percentage.

Table 4.3. *Performance of moving averages of different orders*

Model	RMSE	MAE	MAPE
MA Order 2	254,87	176,51	10,00
MA Order 3	212,95	155,62	8,52
MA Order 4	210,43	145,10	7,93
MA Order 5	210,00	150,43	8,16
MA Order 6	197,05	141,65	7,78
MA Order 7	197,64	141,37	7,85
MA Order 8	205,54	144,76	8,14
MA Order 9	207,74	145,57	8,18
MA Order 10	215,42	147,30	8,28
MA Order 11	209,98	138,95	7,90

In Table 4.3., two of the indicators (RMSE and MAPE) show that the model that fits the data the best is a moving average of order 6. However, according to MAE the MA of order 11 is the best one. As we have explained before, we will use the average of RMSE and MAE (calculated in Table 4.4) for choosing the order. However, the values of RMSE and MAE must be normalised before making the average.

Table 4.4. *Average of RMSE and MAE with normalised values*

	Performance indicators		Normalised performance indicators		
Model	**RMSE**	**MAE**	**RMSE**	**MAE**	**Average**
MA Order 2	254,867	176,505	0,7732	0,7872	0,7802
MA Order 3	212,948	155,624	0,9254	0,8928	0,9091
MA Order 4	210,429	145,096	0,9364	0,9576	0,9470
MA Order 5	210	150,425	0,9383	0,9237	0,9310
MA Order 6	197,052	141,65	1,0000	0,9809	0,9905
MA Order 7	197,64	141,369	0,9970	0,9829	0,9899
MA Order 8	205,54	144,757	0,9587	0,9599	0,9593
MA Order 9	207,736	145,571	0,9486	0,9545	0,9515
MA Order 10	215,423	147,295	0,9147	0,9433	0,9290
MA Order 11	209,981	138,947	0,9384	1,0000	0,9692

The results in Table 4.4. show that the best method is a moving average of order 6. It is important to notice that a moving average of order 6 calculates the forecast for the next value of the time-series as the average of the six most recent values. However, since there is a seasonality of 12 periods in the time-series, Statgraphics uses the "Seasonally Adjusted Data", and not the original time-series, for the calculations.

Table 4.5. *Moving average of order 6 model applied to the training set*

MODEL				
Year	**Month**	**Production volume (ton)**	**Seasonally Adjusted**	**MA(6)**
2013	**January**	1380,1	2267,6	
	February	1663,85	1885,03	

Year	Month	Production volume (ton)	Seasonally Adjusted	MA(6)
2013	March	1286,8	1455,63	
	April	2132,6	2028,63	
	May	1957,5	1998,46	
	June	2473	2107,07	
	July	2595	2087,8	2432,51
	August	2878,1	2044,12	2713,35
	September	2485,5	2024,28	2398,73
	October	2044,8	1826,75	2292,91
	November	1861,1	2241,2	1673,05
	December	1115,7	1885,46	1216,15
2014	January	1410,58	2317,68	1228,35
	February	1748,52	1980,96	1815,27
	March	1622,4	1835,26	1808,75
	April	1945,1	1850,27	2117,8
	May	2083,775	2127,38	1977,1
	June	2139,5	1822,92	2346,75
	July	2318,5	1865,34	2472,3
	August	2662,6	1891,06	2694,46
	September	2369,3	1929,65	2331,31
	October	2390,2	2135,31	2142,96
	November	1434,6	1727,59	1629,21
	December	1209,9	2044,65	1121,53

Year	Month	Production volume (ton)	Seasonally Adjusted	MA(6)
2015	**January**	977,7	1606,43	1176,01
	February	1666,2	1887,7	1667,45
	March	1772,2	2004,71	1669,51
	April	2079,2	1977,83	1998,5
	May	1658,1	1692,8	1836,39
	June	2303,2	1962,4	2193,6
	July	2100,7	1690,11	2306,03
	August	2387,5	1695,68	2631,9
	September	2427,9	1977,37	2255,86
	October	1858,1	1659,95	2051,46
	November	1438,9	1732,77	1477,89
	December	1155,2	1952,21	1057,07

In Table 4.5. we can see the moving average of order six applied to the time-series. The values in column 5 "MA (6)" are calculated as an average of the six previous values in column 4 "Seasonally adjusted", multiplied by the normalised seasonal index of that period, obtained from Table 4.2.

For instance, in June 2014:

$$\hat{x}_{June\ 2014} = \left(\frac{2127{,}38 + 1850{,}27 + 1835{,}26 + 1980{,}96 + 2317{,}68 + 1885{,}46}{6}\right) * 1{,}174$$
$$= 2347{,}41\ tons$$

The calculations in Statgraphics were made with more decimals than the ones shown in the tables in this paper. That is the reason why the previous calculation is slightly different to the value shown in Table 4.5., 2346,75 ton.

Once we have chosen the moving average model that adjusts the best to the data, we have to make the forecast for the year 2016.

Table 4.6. *Forecast for the year 2016 using a moving average of order 6*

FORECAST			
Year	**Month**	**Production volume (ton)**	**MA(6)**
2016	**January**	1282,5	1086,19
	February	1180	1575,27
	March	1504,3	1577,69
	April	1547	1876,15
	May	1498,7	1748,1
	June	2335	2094,62
	July	2306	2218,25
	August	2136,9	2512,82
	September	2251,9	2191,31
	October	1596,2	1997,72
	November	1470,4	1482,01
	December	1266,4	1056,07

The forecast values are presented Table 4.6., in column 4 "MA (6)". In Figure 4.8., we can compare the real production volume with the forecast.

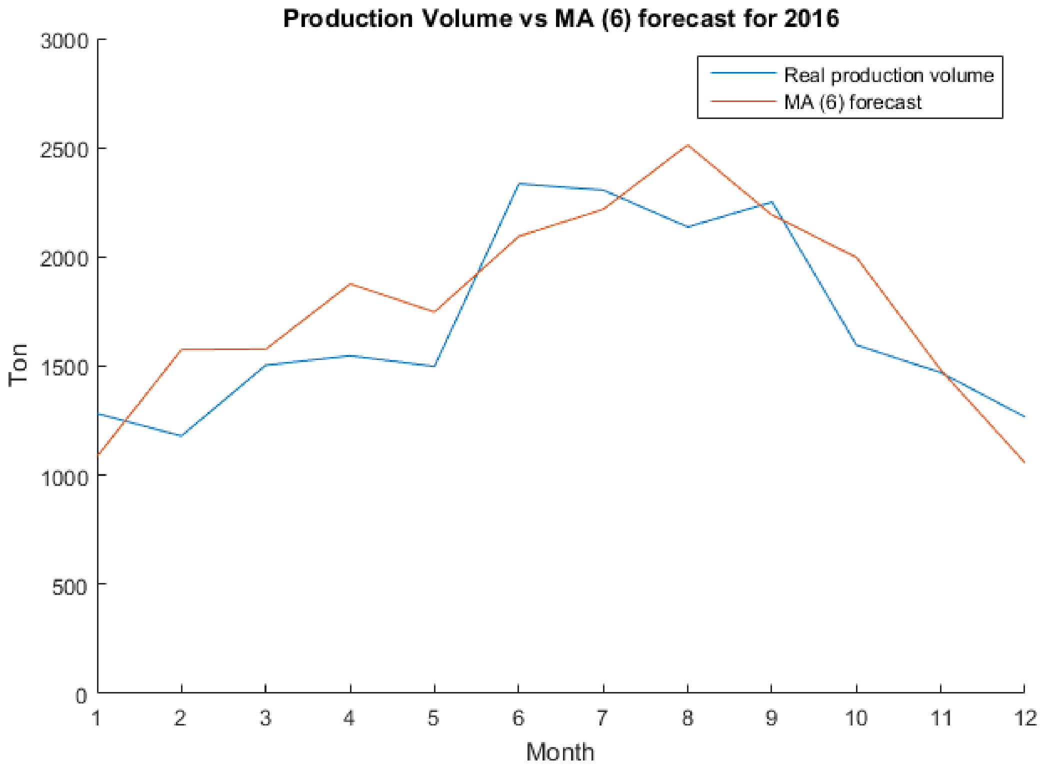

Figure 4.8. *Real production volume in 2016 compared with the forecasts made with MA (6)*

Different errors have been calculated in order to obtain the values of the performance indicators. Table 4.7. shows the comparison between the performance of the forecast and the predictions made by the company.

Table 4.7. *Accuracy of the MA (6) forecasts compared with the predictions made by the Company*

Forecast MA (6)		**Predictions by the Company**	
RMSE	256,05	**RMSE**	316,50
MAE	219,30	**MAE**	276,50
MAPE	14,04	**MAPE**	17,98

All the indicators in Table 4.7. confirm that the moving average forecast gives better results than the qualitative predictions made by the company.

4.4. Simple exponential smoothing

The simple exponential smoothing method is easy to apply, since there is only one parameter that should be calculated: the smoothing constant α. It is important to take into account that small values of α smooth out most of the random noise, but it should only be used when the time-series is considerably stable.

The first step in forecasting with simple exponential smoothing is to find the value of α that gives the optimal model. In our case, the optimal model is the one that gives the lowest RMSE. Statgraphics does not have an option to find the value of α that minimizes the average of RMSE and MAE. For this reason, we have chosen the option to minimize only RMSE.

We insert the production volume from the years 2013–2015 in Statgraphics and calculate the simple exponential smoothing model, choosing the option "Optimize" the smoothing constant α.

Statgraphics gives a value of $\alpha = 0{,}11$. Using the equation (2.10), this is "equivalent" to a moving average of period:

$$N = \frac{2}{\alpha} - 1 = \frac{2}{0{,}0918} - 1 = 17{,}18 \approx 17$$

It is a small value of α, which gives much smoothed results and removes most of the "noise" in the data.

Again, as it happened with the moving average model, the simple exponential smoothing is performed on the "Seasonally Adjusted Data", and then readjusted using the normalised seasonal indexes.

Table 4.8. *Simple exponential smoothing model applied to the training set*

MODEL				
Year	**Month**	**Production volume (ton)**	**Seasonally Adjusted**	**Simple ES**
2013	**January**	1380,1	2267,6	1200,94
	February	1663,85	1885,03	1770,28
	March	1286,8	1455,63	1761,27
	April	2132,6	2028,63	2032,39
	May	1957,5	1998,46	1903,95
	June	2473	2107,07	2288,42
	July	2595	2087,8	2444,99
	August	2878,1	2044,12	2788,36
	September	2485,5	2024,28	2440,2
	October	2044,8	1826,75	2229,17
	November	1861,1	2241,2	1638,66
	December	1115,7	1885,46	1185,14
2014	**January**	1410,58	2317,68	1211,08
	February	1748,52	1980,96	1788,23
	March	1622,4	1835,26	1786,6
	April	1945,1	1850,27	2103,1
	May	2083,775	2127,38	1943,37
	June	2139,5	1822,92	2347,1
	July	2318,5	1865,34	2461,44
	August	2662,6	1891,06	2770,5

Year	Month	Production volume (ton)	Seasonally Adjusted	Simple ES
2014	September	2369,3	1929,65	2405,67
	October	2390,2	2135,31	2189,49
	November	1434,6	1727,59	1640,66
	December	1209,9	2044,65	1152,97
2015	January	977,7	1606,43	1192,29
	February	1666,2	1887,7	1694,92
	March	1772,2	2004,71	1694,36
	April	2079,2	1977,83	2025,07
	May	1658,1	1692,8	1892,4
	June	2303,2	1962,4	2236,65
	July	2100,7	1690,11	2376,41
	August	2387,5	1695,68	2657,63
	September	2427,9	1977,37	2291,67
	October	1858,1	1659,95	2102,88
	November	1438,9	1732,77	1540,05
	December	1155,2	1952,21	1089,5

Once we have the model (results in Table 4.8.) that best fits the time-series, the next step is to calculate the forecasts for the year 2016 and evaluate the results. Again, we introduce the data in Statgraphics and we forecast the next 12 values using $\alpha = 0{,}11$. The forecasts for every month in the year 2016 are presented in Table 4.9.

Table 4.9. *Forecast for the year 2016 using simple exponential smoothing*

FORECAST			
Year	**Month**	**Production volume**	**Simple ES**
2016	**January**	1282,5	1128,01
	February	1180	1635,92
	March	1504,3	1638,43
	April	1547	1948,38
	May	1498,7	1815,4
	June	2335	2175,26
	July	2306	2303,65
	August	2136,9	2609,56
	September	2251,9	2275,67
	October	1596,2	2074,63
	November	1470,4	1539,06
	December	1266,4	1096,73

Since the time-series is seasonal, the forecasts are not calculated using the equation (2.8). It is necessary to modify this equation to take into account the seasonality. After making the necessary modifications, the equation (2.8) looks as follows:

$$\hat{x}_{t,t+\tau} = \hat{a}_t = [(1-\alpha)\frac{\hat{a}_{t-1}}{F_{t-1}} + \alpha x_t]F_{t+1} \tag{4.1}$$

This means that the previous forecast, $\hat{a}_{t-1}$, should be divided by its normalised seasonal index, F_{t-1}, to obtain the "deseasonalized forecast". Then the whole result should be multiplied by the seasonal index of the period for which we are making the forecast, F_{t+1}.

For example, the forecast for October 2013 made at the end of September 2013, is calculated as follows:

$$\hat{x}_{October\ 2013} = \left[(1 - 0{,}11) * \frac{2788{,}36}{1{,}408} + 0{,}11 * 2024{,}28\right] * 1{,}119 = 2221{,}44 \text{ ton}$$

The difference between this value and the value shown in Table 4.8. (2229,17 ton) is due to the number of decimals used in the calculations.

The real production volume and the forecast for the year 2016 are presented in Figure 4.9.

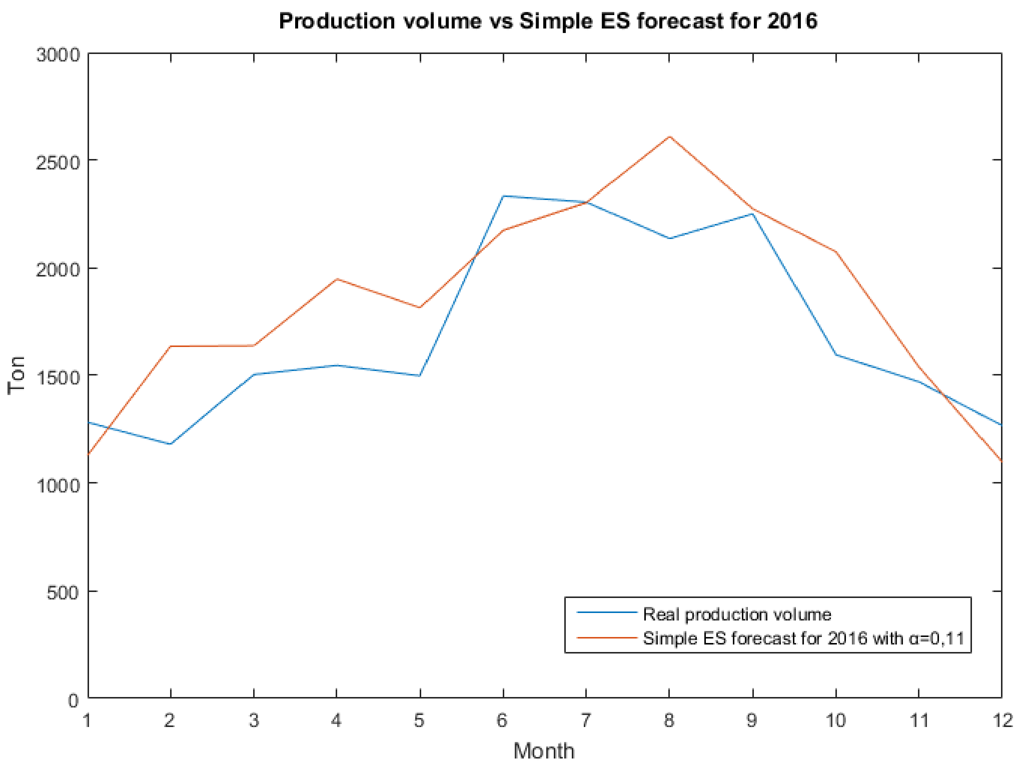

Figure 4.9. *Real production volume in 2016 compared with the forecasts made with simple ES*

In Table 4.10. there is a comparison between the performance of the forecast and the predictions made by the Company, for the year 2016.

Table 4.10. *Accuracy of the simple ES forecasts compared with the predictions made by the Company*

Forecast Simple ES		**Predictions by the Company**	
RMSE	291,98	**RMSE**	316,50
MAE	236,49	**MAE**	276,50
MAPE	15,40	**MAPE**	17,98

All the indicators confirm that the exponential smoothing gives better results than the qualitative predictions made by the company. However, the indicators of the forecasts made with moving average gave better results. It is important to take into account that a Simple ES is not as appropriate as other methods for this time-series, since there is a trend and seasonality. That is the reason why we have tried other exponential smoothing methods that are expected to give better forecasts.

4.5. Brown's Linear exponential smoothing

With this method, not only the level of the time-series is updated with each forecast, but also the trend. We can expect better results than with the simple exponential smoothing, since there is a trend in our data. However, it is important to take into account that the seasonal indexes are not updated; the method uses always the same indexes.

The advantage of Brown's forecast updating method is that it uses only one smoothing constant, α, which makes it almost as simple as Simple ES.

We proceed the same way as we did with simple exponential smoothing. We introduce the data in Statgraphics and we obtain the model that adjusts the best to our time-series. It is the model with the lowest RMSE, with a value of $\alpha = 0{,}07$.

One more time, the mode and forecasts are calculated taking the seasonally adjusted data. The results of the Brown's Linear ES applied to our data are shown in Table 4.11.

Table 4.11. *Brown's linear exponential smoothing model applied to the training set*

MODEL				
Year	**Month**	**Production volume (ton)**	**Seasonally Adjusted**	**Brown's Linear ES**
2013	**January**	1380,1	2267,6	1210,31
	February	1663,85	1885,03	1791,42
	March	1286,8	1455,63	1779,16
	April	2132,6	2028,63	2036,43
	May	1957,5	1998,46	1909,81
	June	2473	2107,07	2296,69
	July	2595	2087,8	2459
	August	2878,1	2044,12	2808,86
	September	2485,5	2024,28	2460,1
	October	2044,8	1826,75	2248,26
	November	1861,1	2241,2	1648,5
	December	1115,7	1885,46	1196,64
2014	**January**	1410,58	2317,68	1220,63
	February	1748,52	1980,96	1810,41

Year	Month	Production volume (ton)	Seasonally adjusted	Brown's linear ES
2014	**March**	1622,4	1835,26	1807,46
	April	1945,1	1850,27	2121,73
	May	2083,775	2127,38	1955,81
	June	2139,5	1822,92	2366,31
	July	2318,5	1865,34	2474,56
	August	2662,6	1891,06	2779,61
	September	2369,3	1929,65	2409,95
	October	2390,2	2135,31	2191,65
	November	1434,6	1727,59	1646,21
	December	1209,9	2044,65	1152,28
2015	**January**	977,7	1606,43	1193
	February	1666,2	1887,7	1686,24
	March	1772,2	2004,71	1684,27
	April	2079,2	1977,83	2015,33
	May	1658,1	1692,8	1884,54
	June	2303,2	1962,4	2218,59
	July	2100,7	1690,11	2359,04
	August	2387,5	1695,68	2628,4
	September	2427,9	1977,37	2258,88
	October	1858,1	1659,95	2076,48
	November	1438,9	1732,77	1515,04
	December	1155,2	1952,21	1069,51

Once we have our model and the optimal value of α, we calculate the forecast for the year 2016. The values for each month are presented in Table 4.12.

Table 4.12. *Forecast for the year 2016 using Brown's linear exponential smoothing*

FORECAST			
Year	**Month**	**Production volume (ton)**	**Brown's linear ES**
2016	**January**	1282,5	1109,51
	February	1180	1605,59
	March	1504,3	1604,55
	April	1547	1903,93
	May	1498,7	1770,1
	June	2335	2116,32
	July	2306	2236,3
	August	2136,9	2527,69
	September	2251,9	2199,41
	October	1596,2	2000,66
	November	1470,4	1480,9
	December	1266,4	1052,93

The graph in Figure 4.10. compares the real production volume in the year 2016 with the forecast made with Brown's Linear ES.

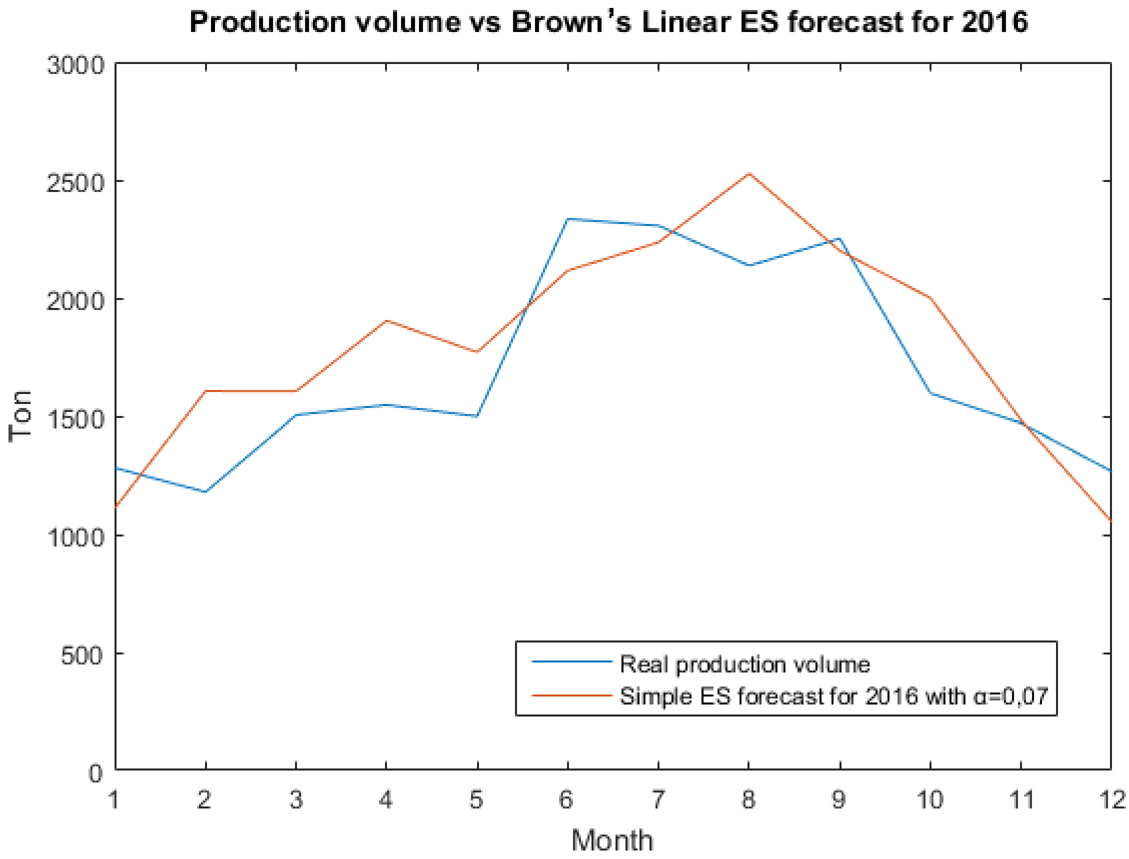

Figure 4.10. *Real production volume in 2016 compared with the forecasts made with Brown's linear ES*

The last step is to calculate the value of the performance indicators used to evaluate the forecasts.

The results shown in Table 4.13. are very similar to the ones obtained with simple exponential smoothing, but there is an improvement in the forecasts. However, the forecasts made with the moving average method are still slightly more accurate.

Table 4.13. *Accuracy of the Brown's linear ES forecasts compared with the predictions made by the Company*

Forecast Brown's Linear ES		**Predictions by the Company**	
RMSE	264,51	**RMSE**	316,50
MAE	223,94	**MAE**	276,50
MAPE	14,44	**MAPE**	17,98

4.6. Holt's linear exponential smoothing

With this method, again both the level and the trend are updated with each new forecast, while the seasonal indexes remain constant. The difference between this method and Brown's linear ES is that the procedure described by Holt for updating the level and the trend uses two smoothing constants, instead of one.

We will use α and β for the smoothing constants, as described in the literature review and in most of the books, articles and softwares about forecasting.

The optimal values of α and β, calculated with Statgraphics, are:

$$\alpha = 0{,}1451$$

$$\beta = 0{,}1968$$

Which gives the results presented in Table 4.14. for the years 2013–2015.

Table 4.14. *Holt's linear exponential smoothing model applied to the training set*

MODEL				
Year	**Month**	**Production volume (ton)**	**Seasonally adjusted**	**Holt's linear ES**
2013	**January**	1380,1	2267,6	1262,03
	February	1663,85	1885,03	1868,33
	March	1286,8	1455,63	1848,85
	April	2132,6	2028,63	2091,3
	May	1957,5	1998,46	1945,63
	June	2473	2107,07	2323,56
	July	2595	2087,8	2477,8
	August	2878,1	2044,12	2823,25
	September	2485,5	2024,28	2467,84
	October	2044,8	1826,75	2251,59
	November	1861,1	2241,2	1643,29
	December	1115,7	1885,46	1194,53
2014	**January**	1410,58	2317,68	1215,56
	February	1748,52	1980,96	1810,17
	March	1622,4	1835,26	1808,46
	April	1945,1	1850,27	2117,48
	May	2083,775	2127,38	1944,14
	June	2139,5	1822,92	2351,97
	July	2318,5	1865,34	2449,77
	August	2662,6	1891,06	2739,8

Year	Month	Production volume (ton)	Seasonally adjusted	Holt's linear ES
2014	September	2369,3	1929,65	2365,59
	October	2390,2	2135,31	2144,53
	November	1434,6	1727,59	1613,25
	December	1209,9	2044,65	1124,55
2015	January	977,7	1606,43	1165,11
	February	1666,2	1887,7	1636,38
	March	1772,2	2004,71	1630,13
	April	2079,2	1977,83	1952,28
	May	1658,1	1692,8	1829,56
	June	2303,2	1962,4	2148,6
	July	2100,7	1690,11	2289,22
	August	2387,5	1695,68	2544,85
	September	2427,9	1977,37	2180,26
	October	1858,1	1659,95	2009,47
	November	1438,9	1732,77	1463,11
	December	1155,2	1952,21	1031,54

One more time, the forecasts are calculated based on the seasonally adjusted data.

The forecast for every month in the year 2016 calculated with the above-mentioned values for the smoothing constants is presented in Table 4.15.

Table 4.15. *Forecast for the year 2016 using Holt's linear exponential smoothing*

FORECAST			
Year	**Month**	**Production volume (ton)**	**Holt's linear ES**
2016	**January**	1282,5	1074,24
	February	1180	1550,45
	March	1504,3	1545,31
	April	1547	1828,71
	May	1498,7	1695,58
	June	2335	2021,71
	July	2306	2130,47
	August	2136,9	2401,42
	September	2251,9	2083,73
	October	1596,2	1890,13
	November	1470,4	1395,14
	December	1266,4	989,135

In Figure 4.11., we compare the results of the real production volume and the forecast:

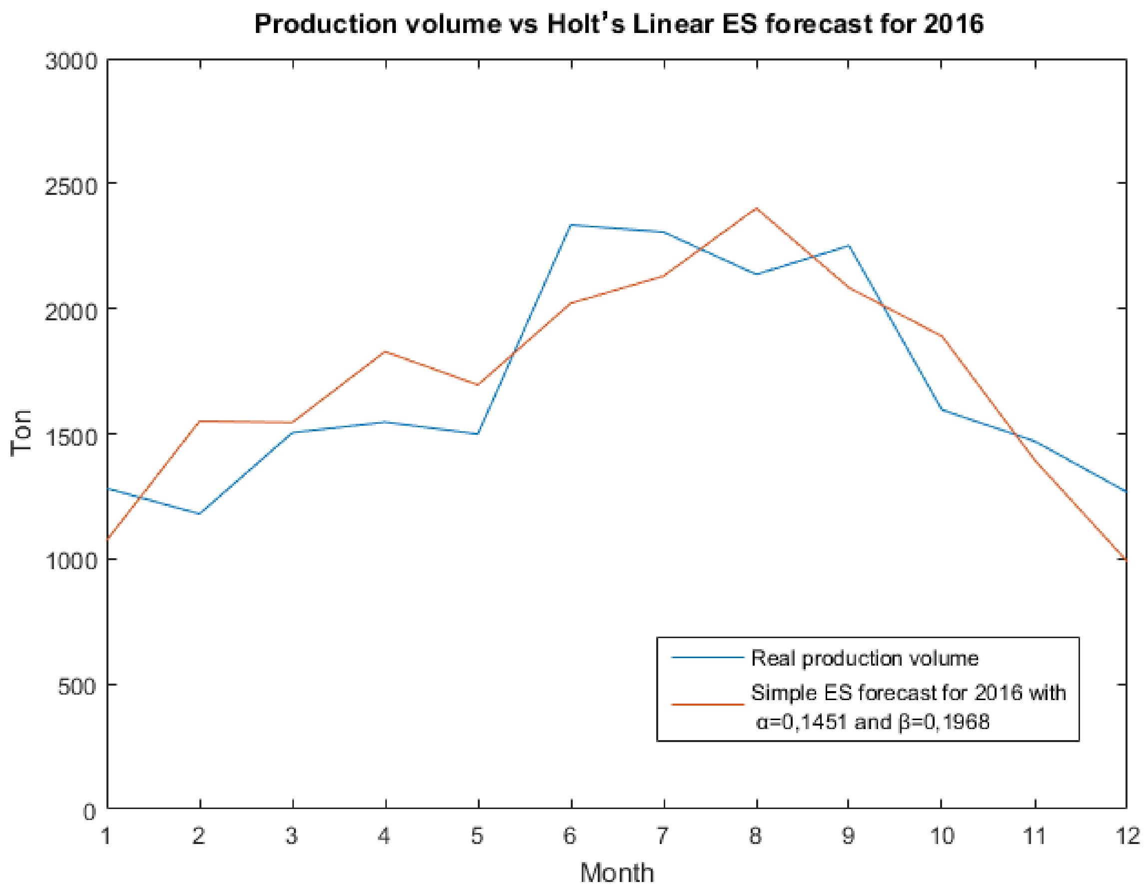

Figure 4.11. *Real production volume in 2016 compared with the forecasts made with Holt's linear ES*

The performance indicators for the forecasts have the values presented in Table 4.16.

Table 4.16. *Accuracy of the Holt's linear ES forecasts compared with the predictions made by the Company*

Forecast Holt's linear ES		Predictions by the Company	
RMSE	240,93	**RMSE**	316,50
MAE	222,19	**MAE**	276,50
MAPE	14,00	**MAPE**	17,98

The results of this method are, in general, better than any of the other methods tried until now. However, the MAE of the moving average forecasts is lower than the MAE of Holt's ES, which means that it is better.

4.7. Holt-Winters exponential smoothing

The Holt-Winters Exponential Smoothing method should only be used for time-series with a seasonal pattern. The Holt's Linear ES only updates the level and the trend on each new forecast. However, as we have seen before, we can also include seasonal indexes in the method. Nevertheless, these seasonal indexes remain constant and they are not updated with each new forecast.

On the other hand, in the Holt-Winters ES, there is one more smoothing constant, γ, that is used for updating the value of the seasonal indexes. The equations for updating the level and the trend, (2.18) and (2.19), are the same than in the Holt's linear ES.

The value of the three smoothing constants calculated in Statgraphics, for the time-series from 2013 to 2015 is:

$$\alpha = 0{,}0174$$

$$\beta = 1$$

$$\gamma = 0{,}6665$$

It is important to notice that this time there was no need to apply a multiplicative seasonal decomposition before using the method. The previous methods that we have used are not meant to be used for time-series with a seasonal pattern. Therefore, we need to eliminate the seasonal component before using the methods, and afterwards we can readjust the forecast applying the seasonality.

However, the Holt-Winters method is designed for seasonal time-series. Therefore, the equations already take into account the seasonality. The results of the model for the years 2013-2015 are presented in Table 4.17.

Table 4.17. *Holt-Winters linear exponential smoothing model applied to the training set*

MODEL			
Year	**Month**	**Production volume (ton)**	**Holt-Winters ES**
2013	**January**	1380,1	
	February	1663,85	
	March	1286,8	
	April	2132,6	
	May	1957,5	
	June	2473	
	July	2595	
	August	2878,1	
	September	2485,5	
	October	2044,8	
	November	1861,1	
	December	1115,7	
2014	**January**	1410,58	1380,10
	February	1748,52	1665,13
	March	1622,4	1290,53
	April	1945,1	2160,59

Year	Month	Production volume (ton)	Holt-Winters ES
2014	**May**	2083,775	1987,56
	June	2139,5	2525,05
	July	2318,5	2648,11
	August	2662,6	2930,40
	September	2369,3	2522,42
	October	2390,2	2067,31
	November	1434,6	1886,66
	December	1209,9	1121,56
2015	**January**	977,7	1405,31
	February	1666,2	1701,63
	March	1772,2	1474,54
	April	2079,2	1952,60
	May	1658,1	1972,68
	June	2303,2	2152,50
	July	2100,7	2294,95
	August	2387,5	2582,97
	September	2427,9	2255,87
	October	1858,1	2117,14
	November	1438,9	1455,87
	December	1155,2	1078,46

The forecast for every month in the year 2016, calculated with the Holt-Winters Exponential Smoothing method is shown in Table 4.18.

Table 4.18 *Forecast for the year 2016 using Holt-Winters linear exponential smoothing*

FORECAST			
Year	**Month**	**Production volume (ton)**	**Holt-Winters ES**
2016	**January**	1282,5	1021,50
	February	1180	1532,91
	March	1504,3	1524,35
	April	1547	1845,76
	May	1498,7	1591,80
	June	2335	2030,07
	July	2306	1945,57
	August	2136,9	2202,13
	September	2251,9	2125,84
	October	1596,2	1741,74
	November	1470,4	1295,15
	December	1266,4	1012,23

As we can see in Table 4.19., the results of this method are the best of all the methods that we have tried until this point. This is not surprising, since this is the only method that updates the seasonal indexes.

Table 4.19. *Accuracy of the Holt-Winters linear ES forecasts compared with the predictions made by the Company*

Forecast Holt-Winters		Predictions by the Company	
RMSE	232,83	**RMSE**	316,50
MAE	204,79	**MAE**	276,50
MAPE	12,96	**MAPE**	17,98

As we expected, this method over performs all the other exponential smoothing methods. The reason is that Holt-Winters ES is designed for time-series with seasonality, while the other exponential smoothing methods had to be adapted to include the seasonal effects.

The graph in Figure 4.12. is a comparison of the real production volume with the forecast made with the Holt-Winters ES method.

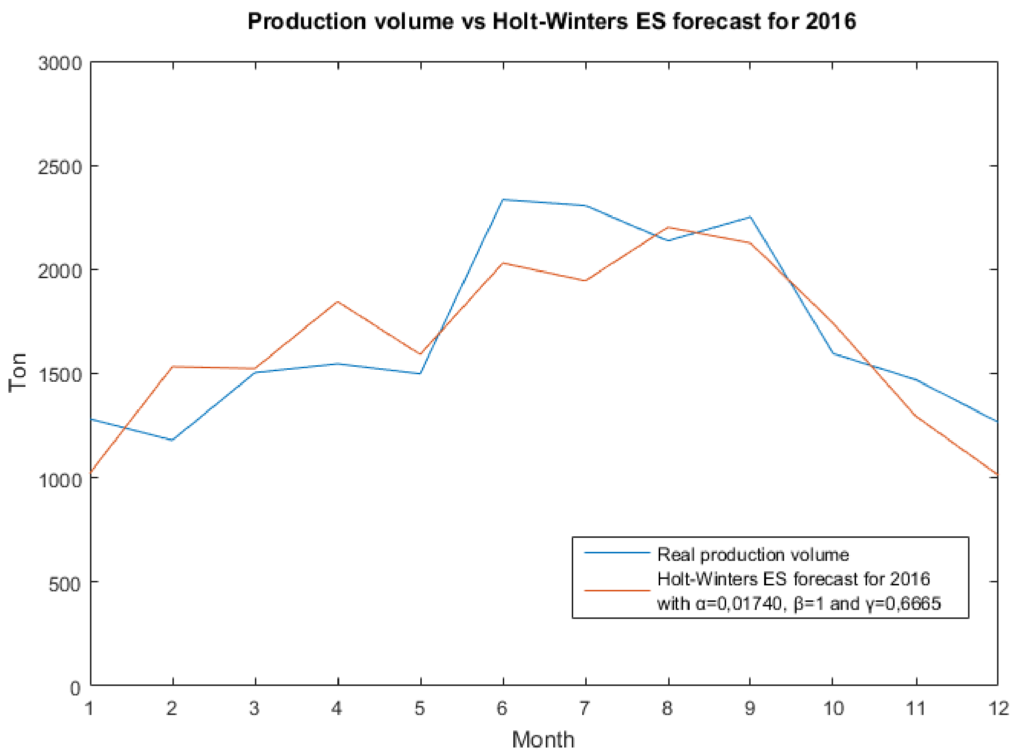

Figure 4.12. *Real production volume in 2016 compared with the forecasts made with Holt-Winters linear ES*

4.8. ARIMA

Fitting an ARIMA model to a time-series is not as easy as the previous models that we have tried. It requires a more advanced knowledge of the method and the procedure to follow.

Since this paper is not a guide for forecasting, it is assumed that the reader has a certain knowledge of time-series analysis. The moving average and exponential smoothing methods, however, are simple and easy to understand for a person with a medium/high level in mathematics and statistics.

Nevertheless, to fit correctly an ARIMA model it is required a higher understanding of the topic. The most common procedure for fitting ARIMA models is the one described by Box and Jenkins (1976).

The biggest challenge in fitting an ARIMA model to a time-series it to determine the number of times that the data should be differenced, and the order of the moving average term and the autoregression term. If the time-series contains seasonality, it is even more challenging since there are seven parameters, instead of three, that must be determined.

In other words, if the time-series is seasonal, then the ARIMA model is defined by $(p, d, q)x(P, D, Q)s$, where:

p is the non-seasonal autoregression term order,

d is the number of non-seasonal differences,

q is the order of the non-seasonal moving average term,

P is the order of the seasonal autoregression term,

D is the number of seasonal differences,

Q is the order of the seasonal moving average term,

s is the number of periods on each season.

The procedure described by Box and Jenkins (1976), which is not explained here, uses the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots to determine both the number of differences needed and the orders of the moving average and autoregression terms.

We have followed the procedure described by Nau (2018) for fitting ARIMA models, which is based on the method described by Box and Jenkins (1976).

The first step is to determine if there is a seasonal pattern present in the data. We already know that this is true, from the plot of the time-series in Figure 4.2. However, it is also possible to detect the seasonality by plotting the ACF and the PACF of the time-series (Figure 4.13.). In our case, for *Product A* the number of periods per season is 12, thus $s = 12$.

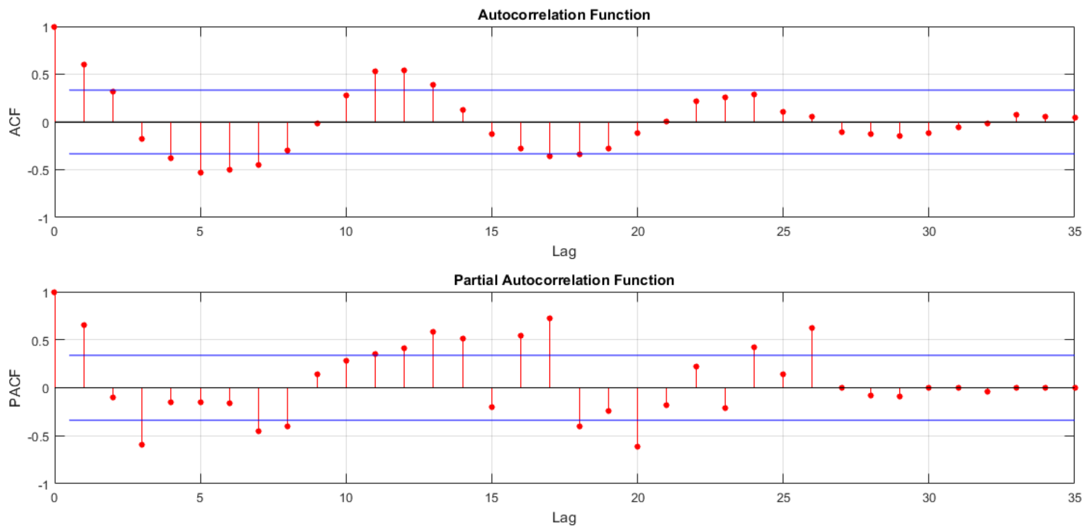

Figure 4.13. *ACF plot and PACF plot of the production volume of Product A*

From the plots in Figure 4.13., there is an obvious seasonality and the time-series is non-stationary (there are nine values of the ACF out of the limits of stationarity). Thus, it must be differenced. The question is how many times it should be differenced. For seasonal time-series, the recommendation is to use only one or two differences. In most of the cases, one seasonal difference is enough. Sometimes, however, it is necessary to add one non-seasonal difference.

A non-seasonal difference consists on subtracting the previous value to each value of the time-series ($y_i' = y_i - y_{i-1}$). A seasonal difference consists on subtracting to each value of the time-series, the value corresponding to the same period but one season before ($y_i' = y_i - y_{i-s}$).

We will first apply one seasonal difference to the time-series and plot the ACF and the PACF of the result of this difference (Figure 4.14.).

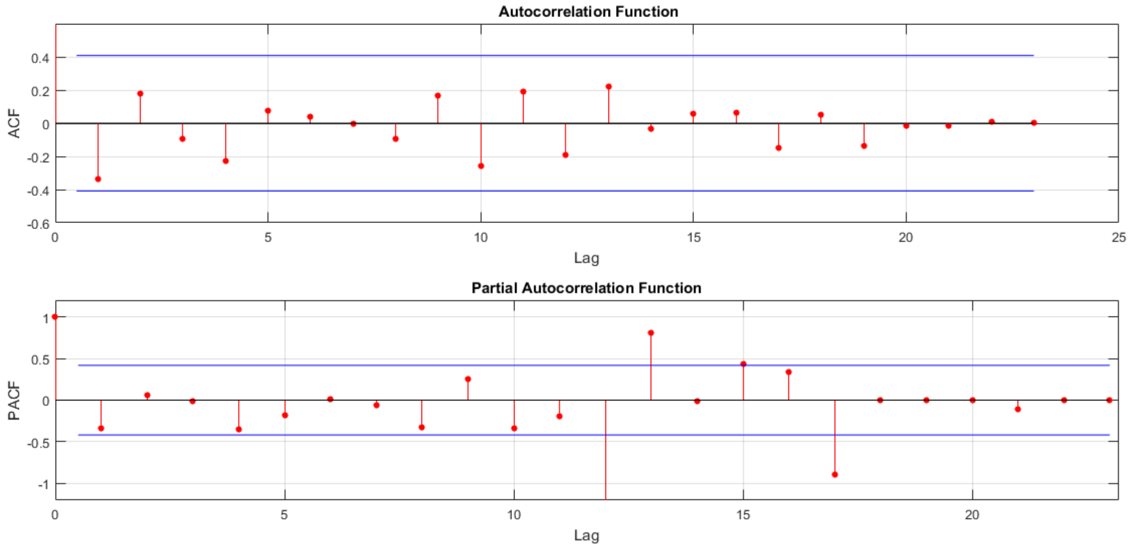

Figure 4.14. *ACF plot and PACF plot of the first seasonal difference of the production volume of Product A*

From the observation of the ACF plot in Figure 4.14., it seems that only one seasonal difference is enough to make the time-series stationary (all the values are inside the limits). However, it is always possible to try to add later one non-seasonal difference and check if the model is more accurate.

Therefore, we have decided that initially $d = 0$ and $D = 1$. The next step is to determine the order of the seasonal autoregression and moving average terms, SAR and SMA.

We have to look at the ACF and the PACF for this (Figure 4.14.). Basically, a pure SAR(1) process has spikes in the ACF at lags s, 2s, 3s, etc., while the PACF cuts off after lag s. On the other hand, a pure SMA(1) process has spikes in the PACF at lags s, 2s, 3s, etc., while the ACF cuts off after lag s.

Since it is possible that our model will contain both SAR and SMA terms, we will use the above-described approach to determine an "initial solution" that afterwards we will try to improve.

There is a clear spike in lag 12 in the PACF plot, while the ACF suddenly decays after lag 12. This indicates that we should include at least one SMA term in the model.

The next step is to fit this model to our time-series in Statgraphics and validate the results obtained. Afterwards, we must try to add AR, MA and SAR terms, and try these combinations with SMA(2).

After trying all these combinations, we obtained that the model that gives the best RMSE is the ARIMA (0,0,0)x(0,1,2)12. There were other models that gave similar results, like ARIMA (1,0,0)x(0,1,2)12 or ARIMA (2,0,0)x(0,1,2)12. However, the best is to keep the model as simple as possible. For this reason, we will make the forecast with the model ARIMA (0,0,0)x(0,1,2)12.

The results of applying the model to the Test Set of the time-series (from 2013 to 2015) are presented in Table 4.20.

Table 4.20. *ARIMA model applied to the training set*

MODEL				
Year	**Month**	**Production volume (ton)**	**Seasonally adjusted**	**ARIMA**
2013	**January**	1380,1	2267,6	
	February	1663,85	1885,03	
	March	1286,8	1455,63	
	April	2132,6	2028,63	

Year	Month	Production volume (ton)	Seasonally adjusted	ARIMA
2013	May	1957,5	1998,46	
	June	2473	2107,07	
	July	2595	2087,8	
	August	2878,1	2044,12	
	September	2485,5	2024,28	
	October	2044,8	1826,75	
	November	1861,1	2241,2	
	December	1115,7	1885,46	
2014	January	1410,58	2317,68	1347,82
	February	1748,52	1980,96	1643,33
	March	1622,4	1835,26	1393,9
	April	1945,1	1850,27	1936,54
	May	2083,775	2127,38	1953,62
	June	2139,5	1822,92	2235,85
	July	2318,5	1865,34	2433,66
	August	2662,6	1891,06	2749,84
	September	2369,3	1929,65	2419,23
	October	2390,2	2135,31	2163,06
	November	1434,6	1727,59	1536,94
	December	1209,9	2044,65	1150,98
2015	January	977,7	1606,43	1133,63
	February	1666,2	1887,7	1717,96

Year	Month	Production volume (ton)	Seasonally adjusted	ARIMA
2015	**March**	1772,2	2004,71	1685,9
	April	2079,2	1977,83	1938,58
	May	1658,1	1692,8	1788,45
	June	2303,2	1962,4	2210,26
	July	2100,7	1690,11	2128,62
	August	2387,5	1695,68	2483,65
	September	2427,9	1977,37	2384,1
	October	1858,1	1659,95	2079,19
	November	1438,9	1732,77	1449,64
	December	1155,2	1952,21	1169,76

The results of the forecasts with ARIMA (0,0,0)x(0,1,2)12 are shown in Table 4.21.

Table 4.21. *Forecast for the year 2016 using ARIMA*

FORECAST			
Year	**Month**	**Production volume (ton)**	**ARIMA**
2016	**January**	1282,5	845,506
	February	1180	1517,05
	March	1504,3	1539,86
	April	1547	2028,95
	May	1498,7	1488,09

Year	Month	Production volume (ton)	ARIMA
2016	June	2335	2304,4
	July	2306	2074,97
	August	2136,9	2339,44
	September	2251,9	2351,39
	October	1596,2	1609,73
	November	1470,4	1469,58
	December	1266,4	1002,76

In Figure 4.22., there are the performance indicators of the forecasts made with ARIMA, compared with the predictions made by the Company.

Table 4.22. *Accuracy of the Holt-Winters linear ES forecasts compared with the predictions made by the Company*

Forecast ARIMA		Predictions by the Company	
RMSE	243,78	RMSE	316,50
MAE	178,65	MAE	276,50
MAPE	11,98	MAPE	17,98

The graph in Figure 4.12. is a comparison of the real production volume with the forecast made with the Holt-Winters ES method.

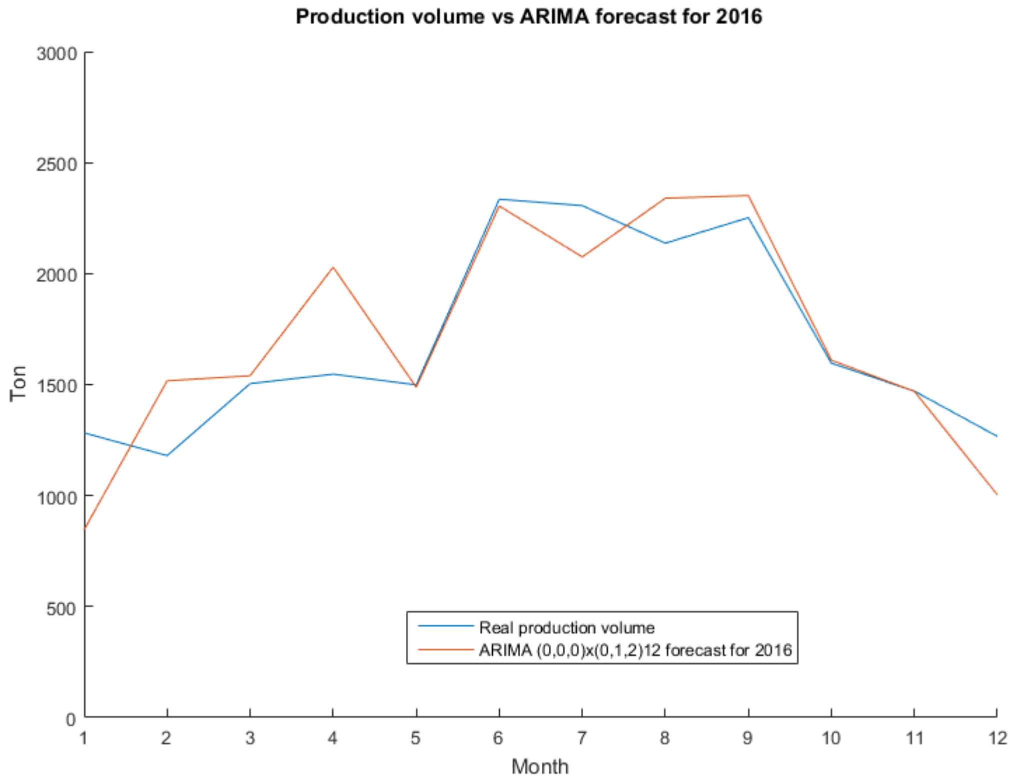

Figure 4.15. *Real production volume in 2016 compared with the forecasts made with Holt-Winters linear ES*

In this case, it is difficult to decide if the results are better than the results obtained with Holt-Winters ES. The RMSE is lower with Holt-Winters ES, but the MAE and the MAPE are lower with ARIMA. The decision of which method is better, depends on what is more important for the forecaster.

The RMSE gives more weight to the highest errors, since the errors are squared. If the forecaster wants to avoid high errors, the Holt-Winters ES would be the best option in this situation. On the other hand, if the forecast prefers to minimize the average error, the forecasts should be made with ARIMA.

5. RESULTS

In the previous section, we have analysed the performance of six different time-series analysis forecasting methods. These methods have been applied to the data of one of the products produced by the Company, *Product A*. The objective is to determine which one of those methods gives better results, and how many of them outperform the predictions made by the Company.

However, it would be dangerous to attempt to reach successful conclusions analysing only one product. We have applied the six methods to three more products – *Product B*, *Product C* and *Product D* – in order to broader results.

It is important to take into account that not all the methods are equally accurate for every product. Each product has a different demand, and some methods work well for some products, while they give bad results for other products.

The procedure followed with *Product B*, *Product C* and *Product D* is the same than with *Product* A. For this reason, it is not going to be described step by step.

In Table 5.1., we present the production volume of 2016, the predictions made by the Company for that year, and the forecasts made for that year with each one of the methods. There are also presented the values of the performance indicators – RMSE, MAE and MAPE – for the predictions and for every method.

Table 5.1. *Production volume, predictions and all forecasts for Product A in 2016*

PRODUCT A (ton) – YEAR 2016							
Production volume	**Prediction**	**MA(6)**	**Simple ES**	**Brown's linear ES**	**Holt's linear ES**	**Holt-Winters ES**	**ARIMA**
1282,5	1301,3	1086,19	1128,01	1109,51	1074,24	1021,50	845,506
1180	1633,5	1575,27	1635,92	1605,59	1550,45	1532,91	1517,05
1504,3	1748,5	1577,69	1638,43	1604,55	1545,31	1524,35	1539,86
1547	1891	1876,15	1948,38	1903,93	1828,71	1845,76	2028,95
1498,7	2004,7	1748,1	1815,4	1770,1	1695,58	1591,80	1488,09
2335	2122,4	2094,62	2175,26	2116,32	2021,71	2030,07	2304,4
2306	2347,5	2218,25	2303,65	2236,3	2130,47	1945,57	2074,97
2136,9	2515,5	2512,82	2609,56	2527,69	2401,42	2202,13	2339,44
2251,9	2352,5	2191,31	2275,67	2199,41	2083,73	2125,84	2351,39
1596,2	2030,4	1997,72	2074,63	2000,66	1890,13	1741,74	1609,73
1470,4	1726,5	1482,01	1539,06	1480,9	1395,14	1295,15	1469,58
1266,4	1594,3	1056,07	1096,73	1052,93	989,135	1012,23	1002,76
RMSE	316,50	256,05	291,98	264,51	240,93	232,83	243,78
MAE	276,50	219,30	236,49	223,94	222,19	204,79	178,65
MAPE	17,98%	11,63%	13,62%	12,56%	13,10%	13,08%	12,34%

As we have mentioned before in this thesis, the MAPE will not be used for making any decision. It is presented only because its values are intuitive and give an easy to understand estimate of the performance of the forecasts.

In Table 5.2., only the values of the performance indicators are shown. We have used a colour scale in order to identify more easily the best methods. The scale goes from red (highest values, which means worst results), to green (lowest values, best results).

Table 5.2. *Performance indicators values of Product A*

	Product A						
	Prediction	MA	Simple ES	Brown's linear ES	Holt's linear ES	Holt-Winters ES	ARIMA
RMSE	316,50	256,05	291,98	264,51	240,93	232,83	243,78
MAE	276,50	219,30	236,49	223,94	222,19	204,79	178,65

Both RMSE and MAE show that every method gives better results than the predictions made by the Company. According to the RMSE, the best method for Product A is Holt-Winters ES. However, according to the MAE, the method that gives the best results is ARIMA.

In order to have values that are easier to understand, we have normalised the values in Table 5.2. These values are represented in Table 5.3.

Table 5.3. *Normalised performance indicators values of Product A*

	Normalised Product A						
	Prediction	MA (6)	Simple ES	Brown's linear ES	Holt's linear ES	Holt-Winters ES	ARIMA
RMSE	0,74	0,91	0,80	0,88	0,97	1,00	0,96
MAE	0,65	0,81	0,76	0,80	0,80	0,87	1,00
Total	0,69	0,86	0,78	0,84	0,89	0,94	0,98

We have in Table 5.3. the normalised values of the performance indicators in a scale from 0 to 1, which gives value 1 to the best value, which is the lowest value. Additionally, we have averaged the results of both methods, obtaining normalised values that represent the accuracy of each method.

The best method is ARIMA, followed by Holt-Winters ES. Nevertheless, we have tried one more method. As we stated in section 2.6.6., some studies suggest that the combination of the two best methods, gives better results than any of the methods itself.

For each month of the year 2016, we have averaged the forecasts made with ARIMA and with Holt-Winters ES methods. This numbers are presented in Table 5.4.

Table 5.4. *Combination of the two best forecasting methods for Product A*

	PRODUCT A (ton)		
2016	**Holt-Winters ES**	**ARIMA**	**Combination**
January	1021,50	845,506	933,50
February	1532,91	1517,05	1524,98
March	1524,35	1539,86	1532,11
April	1845,76	2028,95	1937,36
May	1591,80	1488,09	1539,95
June	2030,07	2304,4	2167,23
July	1945,57	2074,97	2010,27
August	2202,13	2339,44	2270,78
September	2125,84	2351,39	2238,61
October	1741,74	1609,73	1675,74
November	1295,15	1469,58	1382,37
December	1012,23	1002,76	1007,49
RMSE	232,83	243,78	225,55
MAE	204,79	178,65	182,54
MAPE	13,08%	12,34%	12,70%

Once again, in Table 5.5., we have calculated the normalised values of the performance indicators, this time including the combination of the two best methods.

Table 5.5. *Performance indicators values, before and after normalization, for Product A*

	Product A							
	Prediction	MA (6)	Simple ES	Brown's linear ES	Holt's linear ES	Holt-Winters ES	ARIMA	Combination of the 2 best methods
RMSE	316,50	256,05	291,98	264,51	240,93	232,83	243,78	225,55
MAE	276,50	219,30	236,49	223,94	222,19	204,79	178,65	182,54
	Normalised Product A							
	Prediction	**MA (6)**	**Simple ES**	**Brown's linear ES**	**Holt's linear ES**	**Holt-Winters ES**	**ARIMA**	**Combination of the 2 best methods**
RMSE	0,71	0,88	0,77	0,85	0,94	0,97	0,93	1,00
MAE	0,65	0,81	0,76	0,80	0,80	0,87	1,00	0,98
Total	0,68	0,85	0,76	0,83	0,87	0,92	0,96	0,99

As we had assumed, the combination of ARIMA and Holt-Winters ES is the method that gives the best forecasts, followed by ARIMA and by Holt-Winters ES. All the methods over perform the qualitative predictions made by the Company.

For *Product B*, *Product C* and *Product D*, only the normalised values of the performance indicators are presented in Table 5.6., Table 5.7. and Table 5.8.

The results for *Product B* (Table 5.6.) are very similar to the ones obtained with *Product A*. All the methods give better results than the predictions made by the Company, and the best methods are ARIMA and Holt-Winters ES. Additionally, the combination of the two best methods gives better results than any of the methods.

Table 5.6. *Normalised performance indicators values for Product B in 2016*

	Normalised Product B							
	Prediction	MA (5)	Simple ES	Brown's linear ES	Holt's linear ES	Holt-Winters ES	ARIMA	Combination of the 2 best methods
RMSE	0,725	0,803	0,795	0,796	0,725	0,880	1,000	0,992
MAE	0,674	0,783	0,786	0,792	0,776	0,799	0,931	1,000
Total	0,699	0,793	0,790	0,794	0,751	0,839	0,965	0,996

The results of *Product C* (Table 5.7.) are not as optimistic as with *Product A* and *Product B*. Surprisingly, the method which gives the best results is the moving average of order 10. In spite of its simplicity, this method over performs not only the predictions made by the Company, but also more sophisticated methods such as ARIMA or Holt-Winters ES.

Apart from the moving average, the only "method" that gives better results than the predictions made by the Company, is the combination of the two best methods: the moving average and the Simple ES. This is a proof that sometimes even simple quantitative methods can give satisfactory results.

It is important to notice that in this case, the combination of the two best methods is in the second place. It did not over perform the moving average of order 10.

Table 5.7. *Normalised performance indicators values for Product C in 2016*

	Normalised Product C							
	Prediction	MA (10)	Simple ES	Brown's linear ES	Holt's linear ES	Holt-Winters ES	ARIMA	Combination of the 2 best methods
RMSE	1,00	0,98	0,86	0,68	0,70	0,67	0,90	0,92
MAE	0,84	1,00	0,95	0,69	0,74	0,70	0,84	0,97
Total	0,92	0,99	0,90	0,69	0,72	0,68	0,87	0,94

Finally, for *Product D*, only the ARIMA and the Simple ES forecasts improve the predictions made by the Company (Table 5.8.). The combination of these two methods also over performs the predictions, but once again, it does not give the best results. For *Product D*, ARIMA gives the best results.

Table 5.8. *Normalised performance indicators values for Product D in 2016*

	Normalised Product D							
	Prediction	**MA (7)**	**Simple ES**	**Brown's linear ES**	**Holt's linear ES**	**Holt-Winters ES**	**ARIMA**	**Combination of the 2 best methods**
RMSE	0,748	0,731	0,757	0,715	0,664	0,638	1,000	0,865
MAE	0,744	0,730	0,757	0,714	0,658	0,635	1,000	0,862
Total	0,746	0,730	0,757	0,715	0,661	0,637	1,000	0,864

6. DISCUSSION OF THE RESULTS

In section 5, we have presented the results of the forecasts made with different time-series analysis methods. In this section, we will briefly analyse the results.

Firstly, it is important to know which are the methods that give the best results for each product. Additionally, it is valuable to know how many of the methods give forecasts that are better than the qualitative predictions made by the Company. This information is presented in Table 6.1.

Table 6.1. *Best methods for each product*

	Best forecast	Second best forecast	Number of methods better than the prediction
Product A	Combination	ARIMA	7
Product B	Combination	ARIMA	7
Product C	MA (10)	Combination	2
Product D	ARIMA	Combination	3

The combination of the two best methods turns to be the best option for forecasting the production volume of *Product A* and *Product B*, while ARIMA is the second best option. For both of these products, all the methods tried gave better results than the Company's predictions.

For *Product C*, the best option is a moving average of order 10, after applying a multiplicative seasonal decomposition to the time-series. The second best option is a combination of a moving average of order 10 and a Simple ES. Two of the methods gave better results than the Company's predictions.

For *Product D*, ARIMA (0,1,0)x(0,1,2)12 model is the best method, while the combination of ARIMA (0,1,0)x(0,1,2)12 and a moving average of order 7 is the second best option. Three of the methods gave better results than the Company's predictions.

Finally, we have evaluated each one of the methods for every product. In Table 6.2., the crossed cells indicate that the method outperformed the Company's predictions for the forecasts of that product.

Table 6.2. *Number of times that a method outperforms the Company's predictions*

	Product A	**Product B**	**Product C**	**Product D**	**Times that forecast is better than prediction**
MA	x	x	x		3
Simple ES	x	x		x	3
Brown's linear ES	x	x			2
Holt's linear ES	x	x			2
Holt-Winters ES	x	x			2
ARIMA	x	x		x	3
Combination of the 2 best methods	x	x	x	x	4

All the methods exceeded the accuracy of the qualitative predictions made by the Company for at least two of the products.

Additionally, in Table 6.3., we show the improvement in the forecasts of each product. For each product, we have calculated the RMSE and MAE of the Company's predictions, and of the method that gave the best results.

We have calculated the average of RMSE and MAE, and then we have used these values to determine the improvement (in percentage) in the forecasts of each product.

Table 6.3. *Improvement in the forecast of each product*

	PRODUCT A		PRODUCT B	
	Prediction	**Best forecast**	**Prediction**	**Best forecast**
RMSE	316,50	225,55	147,37	107,74
MAE	276,50	182,54	106,73	71,91
Average	296,50	204,05	127,05	89,82
Improvement	45,31%		41,45%	

	PRODUCT C		PRODUCT D	
	Prediction	**Best forecast**	**Prediction**	**Best forecast**
RMSE	234,33	239,74	39,99	29,91
MAE	198,31	167,30	35,51	26,41
Average	216,32	203,52	37,75	28,16
Improvement	6,29%		34,05%	

Product A shows the highest improvement, with 45,31%. On the other hand, *Product C* has the lowest enhancement, which is still 6,29% higher than the Company's predictions. On average, the increase in the performance of the forecasts is of 25,42%.

7. CONCLUSIONS

The results of this thesis confirm our hypothesis that quantitative methods make better forecasts than qualitative methods, when enough data is available. All the products that we have analysed are in a mature stage of their life cycle. For that reason, there is sufficient data to apply time-series analysis methods, and the results obtained are better than the predictions made by the Company.

For the four products that we have used for the experiments, at least two of the six methods tried (seven, if we count the combination of the two best forecasts), outperformed the predictions made by the Company. For two of those products, all the methods, even the simplest ones, improved the Company's judgemental forecasts.

One of the simplest time-series analysis methods, the moving average, was able to improve the qualitative predictions by almost 25% in one of the products. The simplicity of this method (it can be easily implemented in a spreadsheet), make it an attractive option for managers with a shallow knowledge of forecasting, or with the lack of a background in mathematics.

The results obtained prove that, in general, ARIMA models give the best results for the Company's products. For some of the products, they can improve by 40% the yearly predictions made by the Company. Nevertheless, ARIMA models are more sophisticated and difficult to implement. They require a higher forecasting expertise and a deeper background in statistics and mathematics. In addition, expensive software is necessary for forecasting with ARIMA models.

However, all these conclusions should be considered carefully. All the products that we have analysed show a relatively stable demand. Moreover, there are six years of historical data available, which is enough for using time-series analysis methods. On the other hand, when the forecaster has only one or two years (or less, for new products) of data, it is not advised to depend purely on these methods. In those cases, the best option might be the usage of judgemental forecasts and qualitative techniques, perhaps combined with some quantitative method. The same conclusion should be taken into account when sudden changes in the demand are expected due to cyclic variations. It is important to notice that time-series analysis forecasting methods do not consider cyclic variations. Therefore, forecasts with a forecasting horizon longer than one or two years might lead to erroneous results.

Furthermore, it is important to take into account that only a yearly forecasting horizon has been used. Time-series analysis forecasting methods can be also used for quarterly or monthly forecasts. Nonetheless, the qualitative predictions made by the Company are much more accurate the shorter the forecasting horizon is. For this reason, some forecasters might consider to use quantitative techniques only for longer forecasting horizons (yearly or biannual).

In general, the results of the thesis indicate that time-series analysis methods should definitely be considered for forecasting the demand or the production volume of the Company's products. The improvement that these methods can bring to the Company are more accurate forecasts, which results in a better production planning. Consequently, the Company would avoid having a lack or an excess of inventory levels. A lack of goods in inventory might have disastrous effects on the Company, since it could lead to unsatisfied customers and the loss of potential clients. An excess of goods in inventory results in unnecessary costs for the Company. Avoiding these two situations are key for the success of every manufacturing company, and optimizing the forecast is the first step that they should take.

LIST OF REFERENCES

Adhikari, R. and Agrawal, R. (2013). *An Introductory Study on Time-series Modeling and Forecasting*. LAP Lambert Academic Publishing, pp. 12–17.

Armstrong, J. (2001). *Principles of Forecasting: A Handbook for Researchers and Practitioners*. 1st ed. Boston, Mass.: Kluwer Academic Publishers, pp. 443–465.

Axsäter, S. (2006). *Inventory control*. 2nd ed. New York: Springer, pp. 7–37.

Box, G. and Jenkins, G. (1976). *Time-series analysis: forecasting and control*. San Francisco, Calif.: Holden-Day.

Business Set Free Ltd (2013). *4 Stages of the Small Business Product Life Cycle*. [image] Available at: http://www.businesssetfree.com/small-business-product-life-cycle/ [Accessed 20 May 2018].

Chambers, J., Mullick, S. and Smith, D. (1971). How to Choose the Right Forecasting Technique. *Harvard Business Review*, [online] pp.6–17. Available at: https://hbr.org/1971/07/how-to-choose-the-right-forecasting-technique [Accessed 31 Jan. 2018].

Chatfield, C. (2000). *Time-series forecasting*. Boca Raton [Fla.]: Chapman & Hall/CRC, pp. 10–115, 157–185.

Dalinina, R. (2017). *Introduction to Forecasting with ARIMA in R*. [online] datascience.com. Available at: https://www.datascience.com/blog/introduction-to-forecasting-with-arima-in-r-learn-data-science-tutorials [Accessed 27 Mar. 2018].

Hyndman, R. and Athanasopoulos, G. (2013). *Forecasting: principles and practise*. 2nd ed., Otexts.

Imdadullah, M. (2013). *Time-series Analysis and Forecasting*. [online] http://itfeature.com. Available at: http://itfeature.com/time-series-analysis-and-forecasting/time-series-analysis-forecasting [Accessed 30 Mar. 2018].

Kalekar, P. (2004). *Time-series Forecasting using Holt-Winters Exponential Smoothing*. [online] Available at: https://labs.omniti.com/people/jesus/papers/holtwinters.pdf [Accessed 17 Mar. 2018].

Makridakis, S. and Hibon, M. (2000). The M3-Competition: results, conclusions and implications. *International Journal of Forecasting*, 16(4), pp. 451-476.

Makridakis, S., Andersen, A., Carbone, R., Filder, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E. and Winkler, R. (1982). The accuracy of extrapolation (time-series) methods: Results of a forecasting competition. *Journal of Forecasting*, 2(1), pp. 111–153.

Makridakis, S., Chatfield, C., Hibon, M., Lawrence, M., Mills, T., Ord, K. and Simmons, L. (1993). The M2-competition: A real-time judgmentally based forecasting study. *International Journal of Forecasting*, 9(1), pp. 5–22.

Nau, R. (2018). *Statistical forecasting: notes on regression and time-series analysis*. [online] people.duke.edu. Available at: https://people.duke.edu/~rnau/411home.htm [Accessed 25 Mar. 2018].

Silver, E., Pyke, D. and Thomas, D. (2017*). Inventory and production management in supply chains*. 4th ed. Boca Raton: CRC Press-Taylor & Francis Group, pp. 73–137.

Winters, P. (1960). Forecasting Sales by Exponentially Weighted Moving Averages. *Management Science*, 6(3), pp. 324–342.