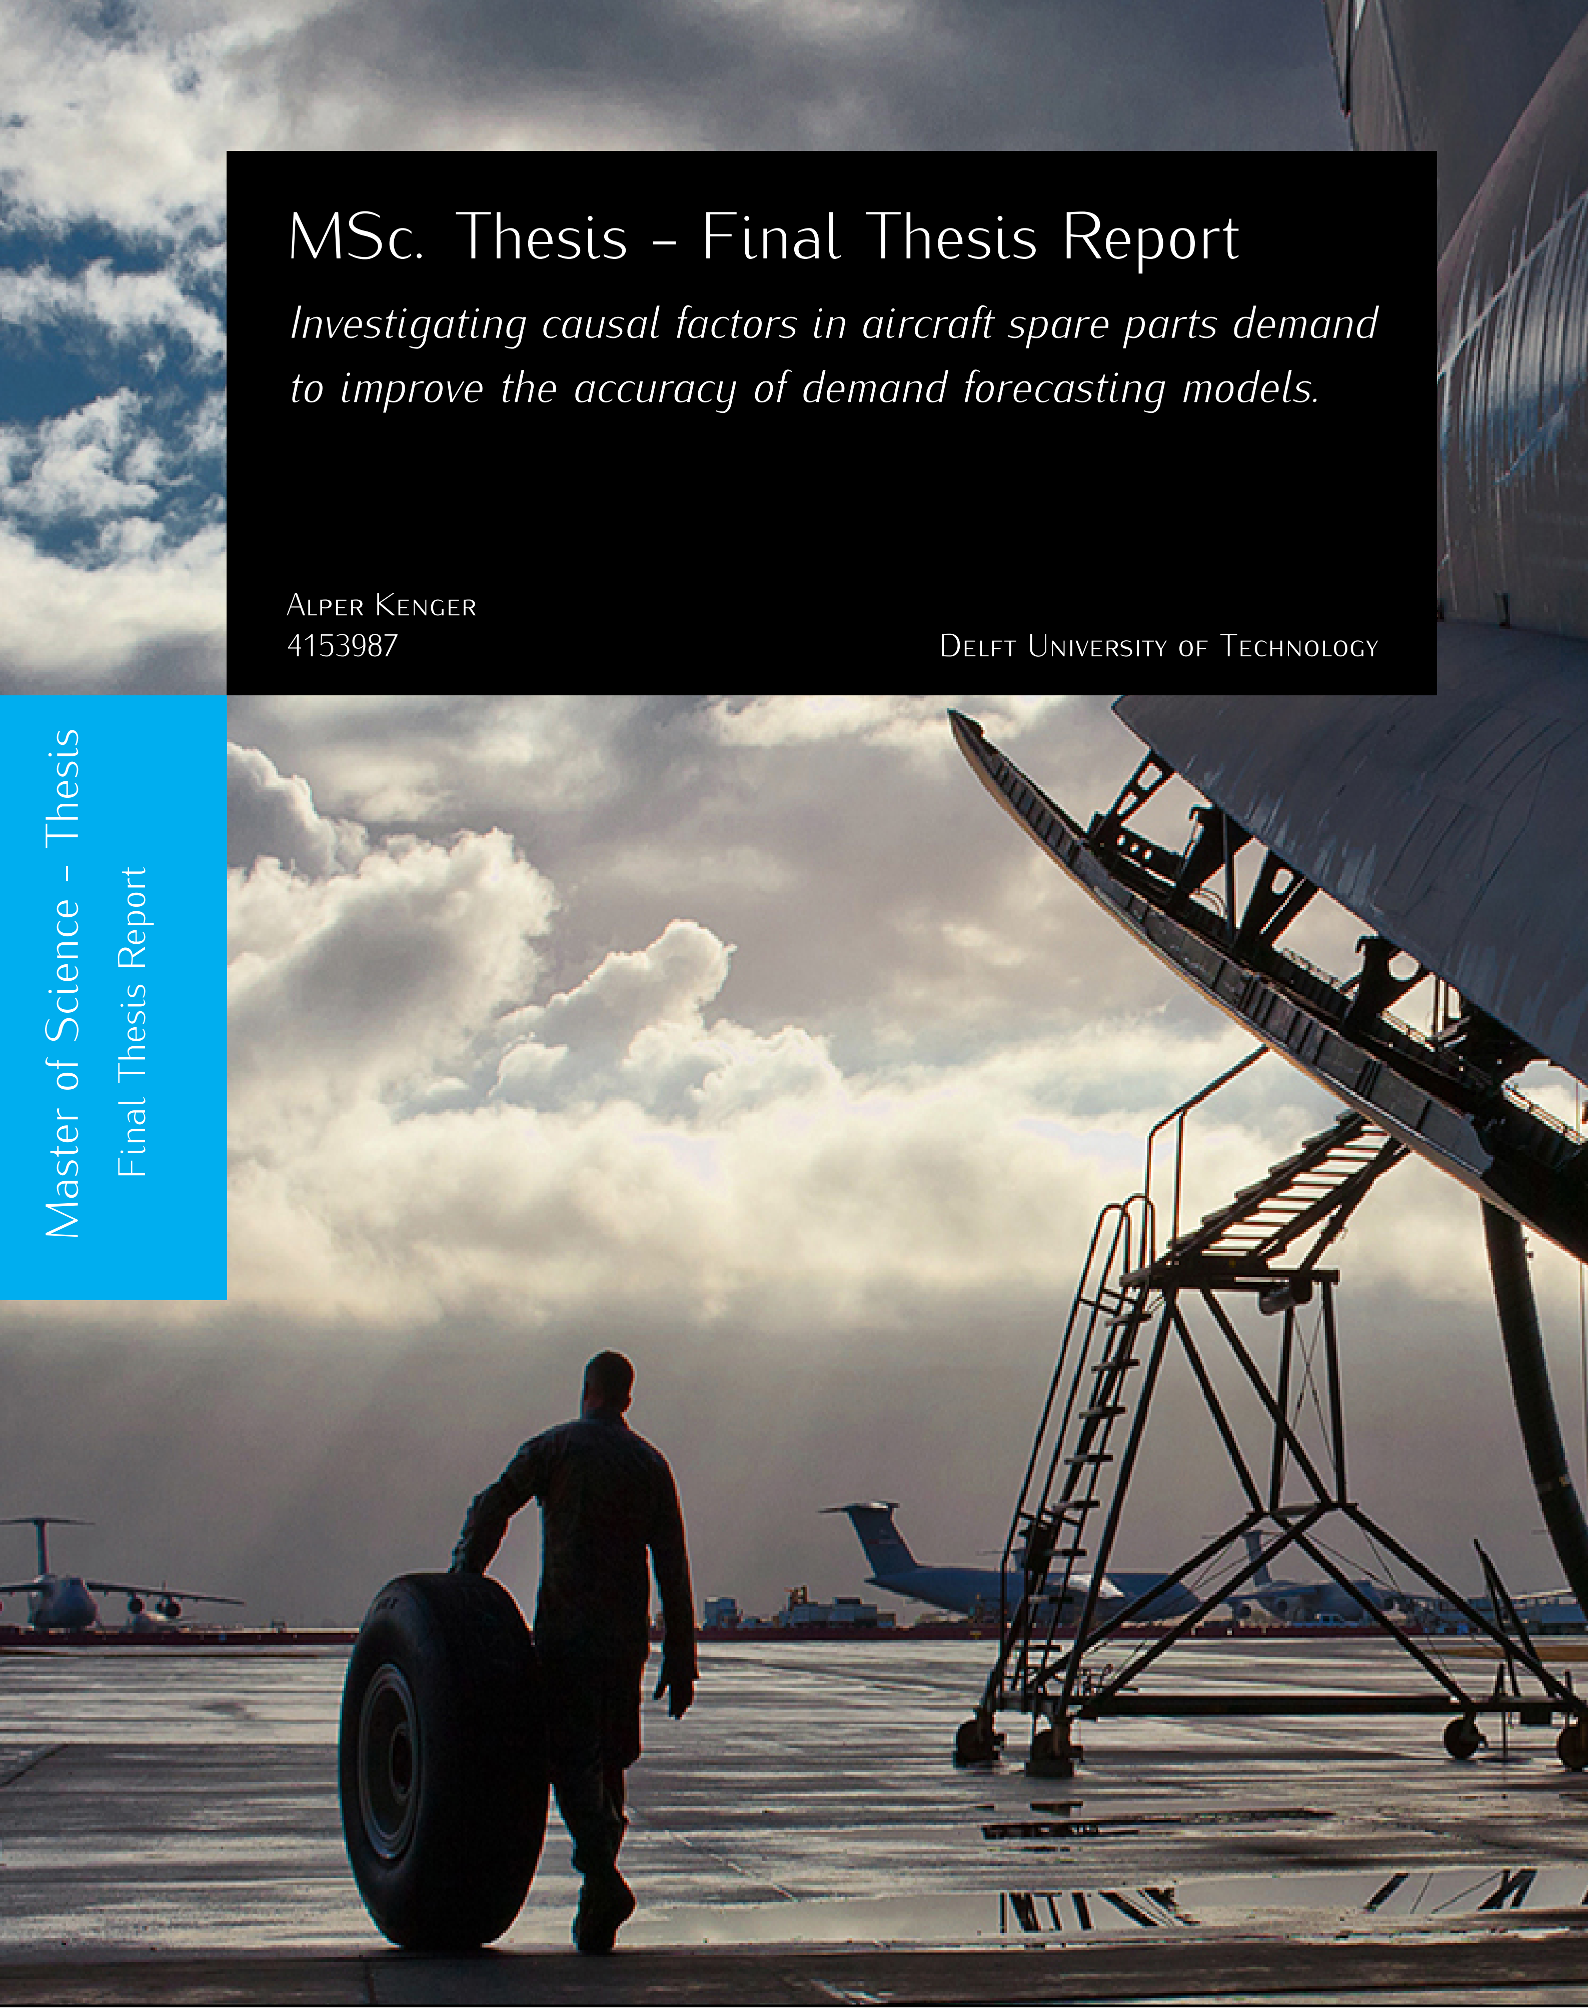

MSc. Thesis - Final Thesis Report

Investigating causal factors in aircraft spare parts demand to improve the accuracy of demand forecasting models.

Alper Kenger
4153987

Delft University of Technology

Master of Science - Thesis
Final Thesis Report

TU Delft Delft University of Technology

Supervisor: Dr. Ir. W.J.C. Verhagen

Faculty of Aerospace Engineering

May 1, 2019

Investigating causal factors in aircraft spare parts demand to improve the accuracy of demand forecasting models.

by

A. Kenger

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on May 15, 2019 at 1 PM.

Student number:	4153987	
Project duration:	July 16, 2018 – May 15, 2019	
Thesis committee:	Prof. dr. R. Curran,	TU Delft, section chair
	Dr. ir. W. J. C. Verhagen,	TU Delft, supervisor
	Dr. ir. E. van Kampen,	TU Delft

This thesis is confidential and cannot be made public until May 15, 2019.

TU Delft

Preface

In accordance with fulfilling the requirements of obtaining a Master of Science degree at the Delft University of Technology, the main findings of the Master thesis research are presented in this report. This report mainly deals with presenting the results following from an extensive analysis of an MRO database related to forecasting models within the airline maintenance operations domain. The findings are used to formulate an improved methodology for applying forecasting methods, by considering statistically correlated causal factors when forecasting aircraft spare parts with time-series methods. The results show that by implementing causal factors with time-series methods, the forecasting accuracy can be improved.

This report will be especially relevant for academia interested in optimising maintenance operations, forecasting spare parts demand and/or identifying underlying causal factors inherent to spare parts demand patterns. Finally I would like to personally thank Dr. ir. Wim Verhagen for his clear and thorough guidance and input throughout the execution of the thesis research.

Keywords: *aircraft spare parts, spare parts forecasting, aircraft maintenance modelling, demand forecasting methods*

Contents

List of Symbols

Symbol	Unit	Description
α	$[-]$	Smoothing constant
A	$[-]$	Aspect ratio
c_{LD}	$[-]$	Correlation coefficient between component removals and aircraft landings
c_{PC}	$[-]$	Correlation coefficient between component removals and pilot complaints
e_t	$[-]$	Forecast error in month t
f_t	$[-]$	Forecast demand value for month t
F'	$[-]$	Baseline demand forecast
F^*	$[-]$	Adjusted demand forecast
k	$[-]$	Time-window for determining average value for PC_0 and LD_0
LD_0	$[-]$	Mean number of aircraft landings in certain timeframe
LD_1	$[-]$	Number of aircraft landings in the current month
m	$[-]$	Moving time-window used in computing the Moving Averages method
PC_0	$[-]$	Mean number of pilot complaints in certain timeframe
PC_1	$[-]$	Number of pilot complaints in the current month
s	$[-]$	Seasonal correction factor
U	$[-]$	Theil's U-statistic
δU	$[-]$	Change in Theil's U-statistic
Y_0	$[-]$	Mean seasonal trend in a certain timeframe
Y_1	$[-]$	Seasonal trend in current month
y_t	$[-]$	Actual demand value for month t

List of Abbreviations

AC	Aircraft
ACO	Aircraft Operator
ACT	Aircraft Type
ADI	Average inter-Demand Interval
ANOVA	Analysis Of Variation
AUR	Aircraft Utilisation Rate
CC	Correlation Cottoefficient
CM	Corrective Maintenance
COL	Component's Overhaul Life
CR	Component Removal
CT	Component Type
CV	Coefficient of Variation
DT	Delay Time
FC	Forecasting
IMS	Intelligent Maintenance System
LND	Aircraft Landings
MA	Moving Averages
MAPE	Mean Absolute Percentage Error
MRO	Maintenance Repair & Overhaul
MSE	Mean Square Error
MWBU	Main Wheel Brake Unit
MWT	Main Wheel Tire
NWT	Nose Wheel Tire
PIC	Pilot Complaints
PM	Preventive Maintenance
PMP	Primary Maintenance Process

RMSE	Root Mean Square Error
SBA	Syntetos-Boylan Approximation
SES	Single Exponential Smoothing
SPL	Seasonal Period Length
SSE	Sum of Squared Errors

List of Figures

List of Tables

CHAPTER 1

Introduction

Existing research has indicated that it is very challenging to plan and allocate resources accurately when dealing with spare parts, since the demand is uncertain for both the frequency and the volume of the demand. So it is not only uncertain when demand will occur, but the volume of the demand when it occurs is also uncertain. This leads to Maintenance, Repair and Overhaul (MRO) companies having to incorporate large spare parts buffers in their operations, in order to ensure having spare parts available at all times. This sub-optimal strategy can lead to very high holding costs, which, according to some estimates, can account for 40% of the total costs for MRO's [7]. Additionally, it is estimated that each year approximately $10 billion is invested in spare parts stocks [8]. Also on the other hand, having too few spare parts can also be very costly. According to Air Transport World [9], a delay of two hours can cost an airline close to $150,000. These figures emphasize the need for improved and more efficient operations and policies when dealing with forecasting spare parts demand and planning accordingly.

Therefore the main problem that this research aims to tackle can be defined as: *"The uncertain nature of spare parts demand makes it very challenging for MRO's to accurately forecast the need for spare parts, often leading to sub-optimal operations."* The objective of the proposed research is therefore to identify methods that will help reduce the demand uncertainty and with that, improve the accuracy of existing forecasting models and consequently improve the efficiency of maintenance policies and operations. The scope of this research project will be limited to spare parts demand forecasting and the thesis will focus on characterising the causal factors that may impact the demand for spare parts.

Currently, time-series forecasting methods are commonly used in practice, which rely heavily on consistent historic data and still perform rather poorly under lumpy or erratic demand patterns. Unique in this research is the fact that statistically correlated causal factors are taken into account with these time-series forecasting methods, so that the estimated demand sizes can be predicted more accurately. The identification and implementation of these causal factors are the main novel aspects of the research, and the corresponding improved methodology of these common time-series methods can be considered the main contribution to the academic state of the art.

This report is the Final Thesis Report which presents the methodology, results and main conclusions of the performed thesis research. Chapter 2 will summarise the relevant Literature study that was performed prior to the thesis research, and it will outline the research scope and relevant research questions of this thesis. Furthermore, Chapter 3 will focus on the description of the general methodology that is applied throughout the thesis project. After this, Chapters 4 through 7 will present the main findings and results obtained through each of the main phases of the thesis. Finally, Chapter 8 will conclude all findings of the thesis research and will e recommendations based on these findings.

CHAPTER 2

Academic background and research scope

Before the research of this thesis can commence, it is necessary to be aware of the academic background that the research deals with. This chapter will therefore describe the most relevant academic literature related to the research topic.

2.1 Relevant academic literature

This section will describe the most relevant academic sources related to the suggested research problem. The applied strategy in reviewing literature is first described, after which the three most important categories of research will be summarised with relevant sources of academic literature. Finally, the main shortcomings in the current state-of-the-art are discussed, before the research scope can be defined.

2.1.1 Applied methodology in reviewing literature

This subsection will detail the general philosophy or strategy that was applied throughout the execution of the search for credible and relevant literature sources. When looking for specific sources, the relevant topic and the problem were always kept in the background when initially selecting research papers based on their titles. The main research problem was defined as:

The uncertain nature of spare parts demand makes it very challenging for MRO's to accurately forecast the need for spare parts, often leading to sub-optimal operations.

To determine whether or not a source was deemed relevant for this literature study, a series of questions were asked as the contents of the documents were being identified. If the document in particular fails to positively respond to any of the questions, it would be deemed to be irrelevant, and as such it would be discarded and excluded from this literature review. Furthermore, the main objective of the paper would also be categorized into three different categories, all of which are relevant within the scope of this research. Any of the relevant papers would either concern itself with model building, model evaluation or the definition of driving factors.

Model building deals with the development of a completely new forecasting model, or an improvement of an existing model. Otherwise within Model evaluation, a paper could also focus on the evaluation of the accuracy of existing forecasting methods, by quantifying the forecasting errors and comparing between applicable models. Finally, within the category of Definition of driving factors, a relevant paper could also deal with investigating the main drivers or causal factors that define the characteristics of spare parts demand patterns. Figure 2.1 shows a visual representation of this entire selection procedure and the corresponding categories a research paper may be considered relevant for.

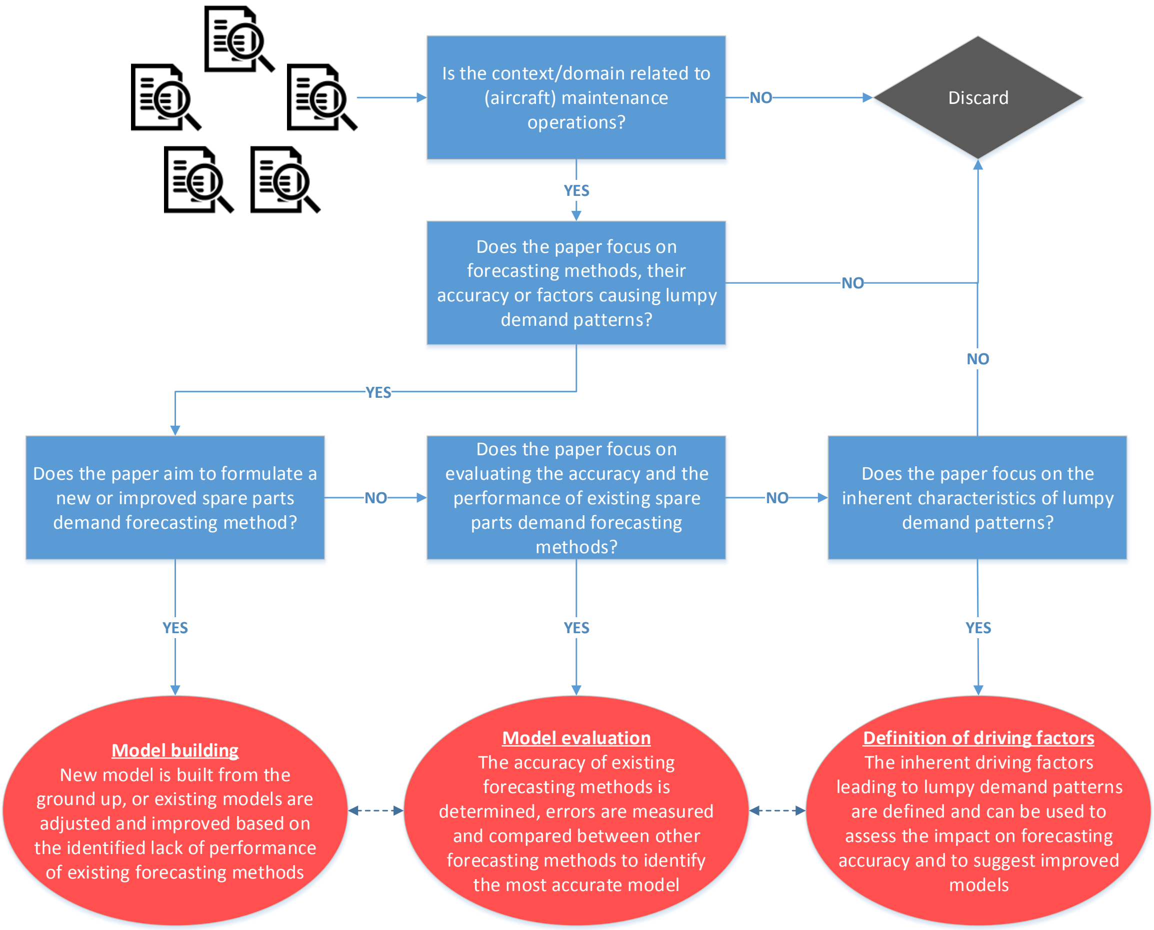

Figure 2.1: Flowchart depicting the selection process of relevant documents for the literature review

2.1.2 Spare parts demand forecasting: model building

As described by Wang and Syntetos (2011), "intermittent demand patterns are very difficult to deal with from a forecasting perspective because of the associated dual source of variation" [10]. According to Wang and Syntetos, Corrective Maintenance (CM) leads to demand being uncertain with regards to the time arrival, but usually deterministic in its size, while demand stemming from Preventive Maintenance (PM) is deterministic regarding arrival, but uncertain regarding the demand size [10].

Additionally, Wang and Syntetos state that currently, all of the forecasting methods developed in recent years mainly focus on coping reactively with demand patterns. The authors criticise the forecasting models for attempting to provide the most accurate modelling of lumpy demand patterns, without questioning the demand generation process itself. Thus they identify as a main gap in knowledge that as of yet, no efforts have been made "to characterise the very sources of such demand patterns for the purpose of developing more effective, pro-active mitigation mechanisms" (Wang and Syntetos, 2011).

They furthermore state that they believe that it would be possible to move away from the reactive nature of current maintenance procedures for spare parts to more pro-active methods, by studying the demand generating process itself. As such, the authors propose a new forecasting model that is based on both regularly planned PM and CM activities using the concept of "delay time". Delay time (DT) modelling is a method that has been discussed in previous literature as well [11] [12] [13] and is based on the principle that if a defective items arrives, it will lead to failure after some delay time.

The main results that the authors found is that the DT model yields more accurate results than the Syntetos-Boylan Approximation (SBA) method. These findings hold true for both the Block based inspection and the Age based inspection. For volumetric pumps, the average absolute error is 87.63 for SBA and 83.29 for DT, and for peristaltic pumps, the average absolute error was found to be 36.34 for SBA and 34.11 for DT. This indeed confirms that regarding forecasting errors, the proposed DT method does outperform the SBA method.

Wang and Syntetos also emphasize that conventional forecasting methods like a time-series based approach, rely heavily on the availability of past data. A maintenance-based approach is not dependent on past data, which is another advantage of using DT over SBA, especially when forecasting items that have little historic data available. However, for this method to work, the reliability characteristics of the items should be known beforehand, since these characteristics are linked to the input parameters that are used by the simulation.

Regarding the investigation to find out the underlying causes and factors of spare parts demand patterns, Wang and Syntetos unfortunately stayed rather superficial. Additionally, none of the results support their conclusion which states that their research "offers insights as to why demand for spare parts is intermittent". Especially in this area, a lot of research opportunities still exist, which is why that will be a predominant aspect within the project scope of the proposed thesis research. More details regarding the project scope can be found in Section 2.2.

Another relevant source that deals with model building, is a relatively recent paper released in 2013 that deals with developing a forecasting method that estimates the material consumption related to non-routine maintenance. In this paper, Zorgdrager et al. [3] focus on several regression and stochastic models to evaluate which model performs the most accurately for forecasting the demand for scheduled maintenance tasks. The main objective of their research is to propose a method that is able to predict material demand specifically for non-routine aircraft maintenance.

The authors first introduce how demand for aircraft maintenance is usually characterised. According to the authors, the classification of any demand pattern is related to its Coefficient of Variation (CV) and its Average Demand Interval (ADI). CV provides a measure of how divergent the demand volume is, i.e.: what is the variance of the demand relative to the average demand. The ADI tells something about how often demand occurs within a specific time frame, and provides a measure of what the average interval is between two demand occurrences. Using these CV and ADI values, specific demand patterns can be identified to be either of the following:

- ***Smooth demand*** (CV<0.49, ADI<1.32) : regular demand occurrence, low variance in demand volume, easy to forecast with low forecasting accuracy
- ***Erratic demand*** (CV>0.49, ADI<1.32) : regular demand occurrence, large variance in demand volume, difficult to predict demand volume
- ***Intermittent demand*** (CV<0.49, ADI>1.32) : irregular demand occurrence, low variance in demand volume, difficult to predict demand occurrence
- ***Lumpy demand*** (CV>0.49, ADI>1.32) : irregular demand occurrence, high variance in demand, difficult to predict both demand occurrence and demand volume

Figure 2.2 shows a graphical representation of these typical demand patterns that are described in the previous list.

According to Zorgdrager et al., the demand for non-routine material can typically be classified as being intermittent or lumpy, which is something they wish to confirm with the available dataset of KLM in their study. Furthermore, they mention that traditional forecasting methods give accurate results for

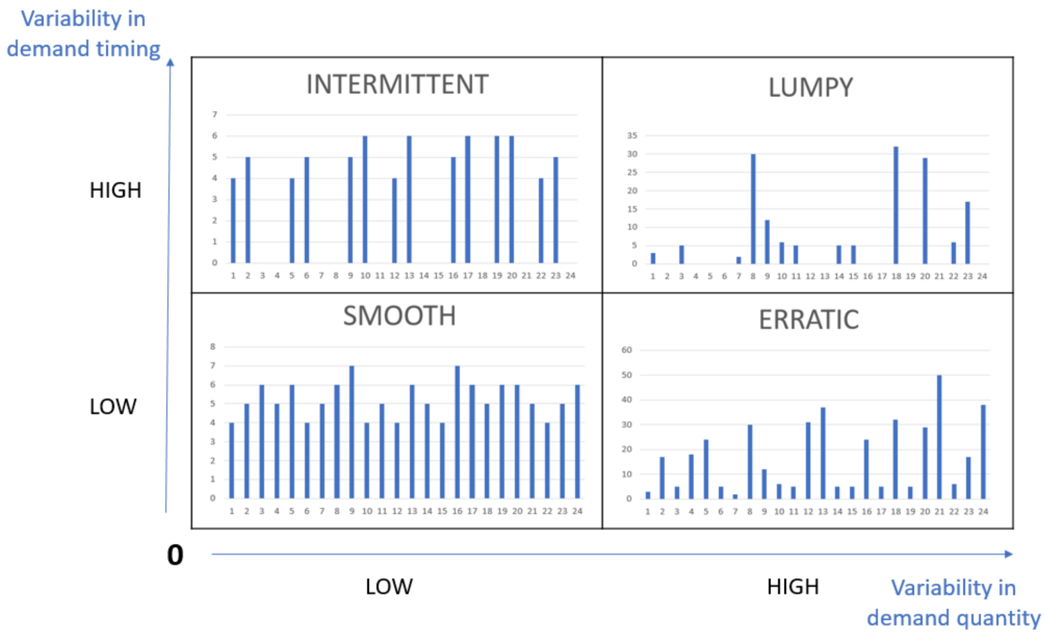

Figure 2.2: Typical demand patterns in the aircraft maintenance domain [2]

smooth demand, but yield inaccurate results for intermittent, lumpy or erratic data. For this reason, the authors consider a range of stochastic models that have shown to perform adequately for intermittent or lumpy demand data, as recommended by Ghobbar et al. [14]. Zorgdrager et al. will then analyse which of these models fits best with their data in the case study.

Subsequently, maintenance data from KLMs B737 fleet was used to analyse the demand predictability for non-routine maintenance checks. For the selected part numbers, the required material for non-routine maintenance was linked to scheduled maintenance tasks, thus effectively making the uncertain occurrence of non-maintenance tasks more predictable. Consequently, the authors computed for each forecasting model the Sum of Squared Errors (SSE) and Root Mean Squared Error (RMSE), to assess the accuracy of both the demand probability and quantity as predicted by the forecasting models. The results of this assessment are summarised in Figure 2.3.

140N2139-1: **Cabin Window Assembly**	Forecasting model	Probability		Quantity	
		SSE	**RMSE**	**SSE**	**RMSE**
	Weighted Mean	0.60	0.31	571.17	9.75
	Linear	0.07	0.13	158.5	6.29
	Weighted Linear	0.29	0.27	772	13.8
	Weighted Exponential	0.79	0.44	97.52	4.9
	Weighted 2nd Polynomial	0.27	0.30	206.70	8.30
	Weighted 5th Polynomial	0.00	0.00	0.00	0.00
	MA	0.05	0.09	52.80	2.96
	EMA	0.04	0.08	26.65	2.10
	Savitzky Golay Filter	0.04	0.09	15.58	1.61
	SES	0.00	0.00	0.00	0.00
	Croston	0.06	0.10	6.25	1.02
	SBA	0.50	0.29	250.65	6.46

Figure 2.3: Overview table showing the SSE and RMSE values for all forecasting models [3]

Based on the results presented in Figure 2.3, the authors conclude that when regarding overall forecasting accuracy, regression forecasting models are not suitable for predicting non-routine material demand, while stochastic models show significant better performance, partly due to the high reactiveness of these models and their ability to adapt to the irregular demand patterns captured by them. Overall, the authors choose the EMA method as the most suitable method for forecasting parts demand for non-routine maintenance tasks, due its low error values and its ability to capture general demand trends.

With this method, the authors have successfully shown that it is actually possible to improve the predictability of the demand for parts due to non-routine maintenance by linking them to the scheduled maintenance tasks.

The results of this paper provide more insights on how to reduce uncertainty of a certain sub-set of demand for parts, specifically those required for non-routine maintenance. Even though the findings are restricted to non-routine maintenance demand, similar methods could be applied when dealing with more general demand patterns. Especially the insights regarding the grouping of parts in case of low availability of historic data, or linking the probability of demand occurrence to other events that are more predictable, can highly benefit the research methodology proposed for this thesis.

2.1.3 Spare parts demand forecasting: model evaluation

The sources detailed in this subsection are relevant for the category of evaluating the errors of spare parts demand forecasting models. Many of the sources look at existing forecasting models that are identified to perform well, and the authors try to determine the most accurate model by comparing forecasting errors between the models. These insights can then be used to further improve these forecasting models.

In 2003, Adel A. Ghobbar and Chris H. Friend conducted a research on developing a predictive model that can indicate which existing forecasting methods are most appropriate to be used by airline operators and MRO organisations. Starting with the definition of the state of the art, the authors describe demand forecasting to probably be the biggest challenge in the MRO industry, as airlines face a common problem of needing to know the short-term spare part demand with high accuracy [14]. The authors' work will focus on achieving the following two main objectives of their research [14]:

- To analyse the behaviour of different forecasting methods when dealing with lumpy and uncertain demand. According to the authors, the performance of a forecasting method should vary with the level and type of lumpiness (i.e., with the sources of lumpiness).
- Based on the forecast accuracy measurements and the results of their statistical analysis, a predictive model is developed successfully for each of the 13 forecasting methods analysed.

To reach these objectives, the authors have selected 13 forecasting methods to consider in their study. They use sample data from Fokker, BAe and ATR, taking into account only repairable parts with unpredictable and recurring demand behaviour. The weekly demand levels in these data sets were grouped together to give overviews of monthly and quarterly intervals of demand, with corresponding ADI- and CV-values. Almost all of the data were categorised to be either lumpy or intermittent. The Microsoft Excel tool *solver* was used to estimate the optimal smoothing parameters that will minimize forecasting errors, before initialising their forecasting methods and measuring the accuracy of the models by using the Mean Average Percentage Error (MAPE) metric.

An Analysis of Variance (ANOVA) was then used to determine the impact of ADI, CV, Seasonal Period Length (SPL) and Primary Maintenance Process (PMP) on the forecasting errors, specifically the MAPE metric, in order to gain an understanding of the significance of sources of lumpiness. The p-values resulting from this ANOVA were analysed to determine if a factor was significant or not, with a p-value being lower than a significance level of 0.01 or 0.05 (depending on the model and factor) indicating a significant relationship. Applying this methodology yielded the following main results:

- The highest forecasting error occurs when Winter's method (either AW or MW) has to forecast demand with high variation.
- Weighted moving averages is much superior compared to exponential smoothing
- WMA, EWMA and Croston's method show the best performance compared to the other models

- The impact of demand variability (ADI and CV^2) on forecast errors is significant, with an increasing demand variability leading to an increased MAPE.
- Generally, hard-time components show to have more effect on increasing the forecast error compared to condition-monitoring components.
- An increased SPL will reduce the average forecasting error for all methods.

The fact that the WMA approach is superior was also found by the research performed by Zorgdrager et al. Furthermore, the superiority of both WMA and Croston's method is once again confirmed. In addition, it is good to see that the authors add knowledge to the state-of-the-art by finding results that indeed confirm that the extent of lumpiness has an impact on the forecasting error, even though there was existing knowledge on the fact that lumpy and erratic patterns lead to inaccurate forecasting results in general. An interesting take in this research is that the authors did not only consider which forecasting model had the least errors, but they also evaluated the impact of some of the underlying factors of lumpy demand.

Another paper that deals with the evaluation of forecasting models, is a research conducted by A. A. Syntetos and J. E. Boylan in 2005 [15]. Like many other authors, Syntetos and Boylan start their paper by explaining the difficulties of forecasting spare parts, due to the demand patterns showing a dual source of variation.

According to Syntetos and Boylan, the current state of the art and the standard method in forecasting spare parts is Croston's method. Croston successfully proved the biased nature of SES models when applied in an intermittent context [16]. Even though Croston's method was claimed to be unbiased, Syntetos and Boylan did show that it was positively biased, therefore over-estimating mean demand.

Subsequently, they suggested an adjusted and improved version of Croston's method, which was deemed to be an approximately unbiased forecasting method called the Syntetos-Boylan Approximation (SBA) [17]. They also devised the SY method, which is another modification of Croston's method, which appeared to be exactly unbiased [18].

Therefore, it is necessary to find out which model actually shows the minimum variance and thus can be determined to be an unbiased estimator of mean demand. In their research paper, Syntetos and Boylan evaluate the variance explicitly for SES, Croston's method, the SY method and the SBA method. Unfortunately though, no actual results and values are computed, as the authors aimed to provide the relevant equations and relations, which could then potentially be used for further analytical work. As such, the authors can not conclude themselves which model shows the least variance, and thus is the most unbiased estimator.

The next model evaluation paper that will be outlined is that of Wallström and Segerstedt (2010) [4]. Unlike the other papers, this research deals with evaluating several forecasting error measurements, instead of focusing on the accuracy of the most commonly used forecasting methods. The authors point out that in existing research, evaluations of forecasting methods are often carried out using only one measure of error, most commonly with the Mean Absolute Deviation (MAD) or with the Mean Squared Error (MSE).

The authors start their paper by briefly describing the governing equations which are used in the four forecasting methods that will be applied in their research. SES is explained to be very efficient in providing short-term forecasts for smooth demand patterns, depending on the proper selection of a smoothing constant. As stated by other research papers, the authors also state that SES is very inaccurate under intermittent circumstances, and describe Croston's method to be effective for these demand patterns. Since the Croston method was still shown to be biased [18], a model adjusted by Syntetos and Boylan is also described, which the authors call $CrSyBo$. Finally, another modified Croston method is described, which is a model that forecasts the demand rate directly, and like SES also requires one smoothing constant.

For the four forecasting methods, the number of items showing the lowest error for each type of error measure was then counted. An example of an overview of MSE results is presented in Figure 2.4. Looking at this figure, it can be seen that regarding forecasting performance, SES showed the lowest MSE values. In a similar fashion, results are presented and discussed for remaining error measures as well.

MSE	Mean				Naive			
	ModCr	Croston	*CrSyBo*	*SES*	*ModCr*	Croston	*CrSyBo*	*SES*
0.025–0.30	8	4	16	44	7	6	24	35
0.025	7	7	15	43	6	9	21	36
0.05	4	3	31	34	6	10	27	29
0.075	5	1	48	18	5	2	44	21
0.10	4	1	58	9	4	2	55	11
0.15	2	3	64	3	4	2	60	6
0.20	0	3	68	1	2	2	66	2
0.25	0	0	72	0	2	2	67	1
0.30	0	0	72	0	1	0	70	1

Figure 2.4: One example of a results table from Wallström and Segerstedt's research [4]

The authors conclude their paper by stating that while evaluating the several forecasting methods *ModCr* showed the most bias errors, followed by Croston. Based on the results, ModCr would overestimate the demand consistently, thus making it the least suitable method. The authors also mention that none of the forecasting methods are completely free of bias in all cases, and at some point will show bias. Therefore they suggest that it should always be important to have methods that can detect the bias (and not only the error), so proper corrections can be implemented in the forecasting methods.

The final research paper that was found to be relevant regarding the evaluation of forecasting models is the one written by Regattieri et al. in 2005 [5]. Using data from Alitalia, Regattieri et al. analyse the behaviour of forecasting models under lumpy conditions, and they identify the effectivity and accuracy of models that are used to forecast aircraft spare parts. Referring to Ghobbar and Friend's research [19], the authors mention that only 10% of companies actually use forecasting models, while the majority of airlines usually base their predictions on their operational experience, annual budgets and recommendations provided by manufacturers.

The proposed methodology by Regattieri et al. first starts with measuring the degree of lumpiness in their data set, continued by the selection of forecasting models to be evaluated, and concluded with an evaluation of error values. Figure 2.5 shows the degree of lumpiness by plotting the ADI and CV values for each of the five components. Since all of the components have ADI and CV values over 1.32 and 0.49 respectively, it becomes very apparent that all of these components show lumpy demand patterns. It can also be seen that item w is the most lumpy (largest CV and ADI values), while item z is the least lumpy (smallest CV and ADI values).

Furthermore, Figure 2.6 shows the performance of the forecasting methods in general, with the position scores either summed or averaged to indicate the accuracy of each method. In this case, the lowest total score represents the best performance. Based on these results, the conclusion can be drawn that the WMA method performs best across all boards (at least regarding forecasting accuracy), followed by Croston's method. Additionally, the error values are also graphed for each item and forecasting method combination, as shown in Figure 2.7.

From this graph it can already be seen on first glance that regardless of item type, WMA and Croston show the lowest error values. However, a more interesting note is that the item lumpiness is actually the determinant factor for the magnitude of errors, and not necessarily type of forecasting method. These results are also in accordance with Figure 2.5, as item w had the highest lumpiness and as such shows the highest error values, while item z had the lowest lumpiness and shows the lowest error

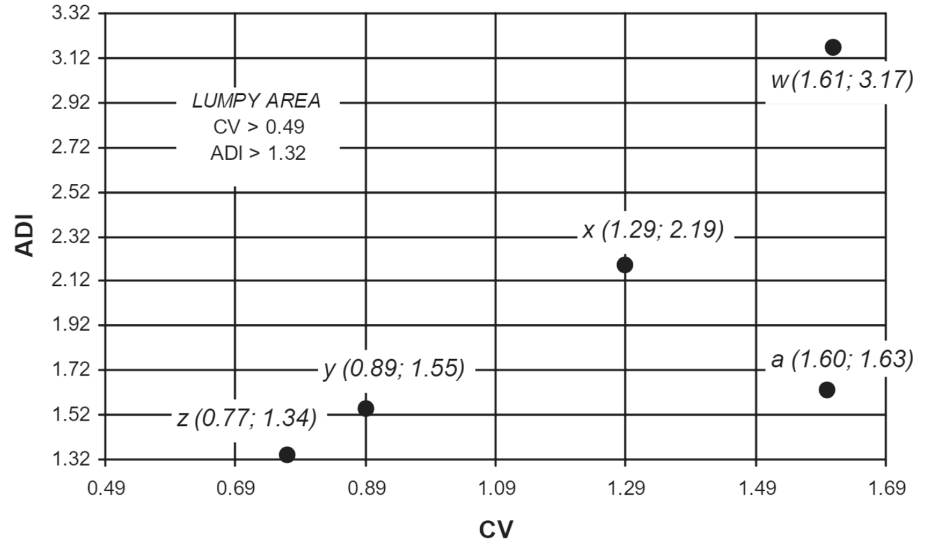

Figure 2.5: Degree of lumpiness for each component type (based on monthly CV and ADI) [5]

Method	Total score	Average score
WMA	8	1.6
CROSTON	18	3.6
EWMA	21	4.2
TAES	21	4.2
SRM	24	4.8
MW[a]	24	6.0
AW	31	6.2
SES	34	6.8
MA(12)	54	10.8
MA(11)	56	11.2
MA(9)	58	11.6
MA(10)	58	11.6
MA(8)	60	12.0
MA(7)	61	12.2
MA(4)	70	14.0
MA(5)	73	14.6
MA(6)	76	15.2
MA(3)	89	17.8
DES	95	19.0
MA(2)	98	19.6

Figure 2.6: Overall performance of forecasting methods, with total and average scores [5]

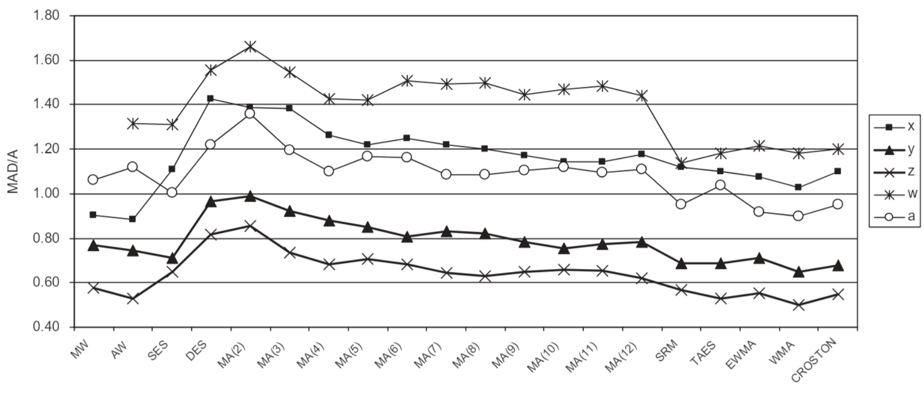

Figure 2.7: Plotted MAD/A values to indicate method accuracy for each item [5]

values, accordingly.

The research of Regattieri et al. is very relevant for this thesis research, as the authors not only present results that show which method performs best, but they also underline that in the general picture, item lumpiness is the main factor that impacts forecasting inaccuracies, while the specific forecasting method is of secondary importance. Similar to the findings of Regattieri et al., a study conducted by Kostenko and Hyndman (2006) [20] also confirms that the magnitude of CV^2 impacts the accuracy of the selected forecasting method. Additionally, in a research performed by Petropoulus et al. (2014) [21], the best-performing forecasting methods are selected based on not only CV^2, but also on other demand characteristics such as the length of the series, the seasonal period length and the forecasting horizon.

With regard to selecting and applying suitable forecasting methods, some authors also apply bootstrapping methods which have shown advantages in certain conditions [22] [23], but they are computationally demanding since the calculations are rather complex. This is also why they are not often implemented in practice. Furthermore, a recent study by Syntetos [24] has shown that the advantages of these bootstrapping methods over conventional methods are questionable. This is why improving time-series forecasting methods will be a focal point in this thesis research

2.1.4 Definition of driving factors

This subsection will outline the most important literature that focuses the inherent characteristics generating spare parts demand. Even though research on this specific topic is rather scarce, some relevant sources could still be identified. Some of the few authors that are particularly concerned with the actual underlying sources of demand patterns, are A. A. Ghobbar and C. H. Friend. Their papers on evaluating model errors are already discussed in the previous subsection, but they also have an interesting piece of work on the investigation of sources of demand lumpiness [25].

Ghobbar and Friend believe that environmental factors can have an impact on the extent of lumpiness of demand. To verify this hypothesis, they select a number of factors to investigate whether or not they have an effect on lumpiness in spare parts demand. The factors that are included in their experiment are the following:

- Primary maintenance process (PMP)
- Aircraft utilization rate (AUR)
- Component's overhaul life (COL)
- Square coefficient of variation (CV^2)
- Average inter-demand interval (ADI)

By using an ANOVA method, their aim is to find and compare p-values, which quantify the level of impact a factor can have on lumpiness. Their findings show that all factors and their interactions were highly significant, thus implying that these factors most likely have an impact on demand lumpiness. It also appears that the coefficient for AUR is positive, which implies a positive correlation between aircraft utilisation rate and demand size.

The authors conclude their paper by stating that AUR, COL and PMP are major sources in increasing the demand size, which they believe can aid material managers in providing a clearer picture and could therefore lead to substantial benefits. Additionally, the authors mention that understanding the sources of lumpiness is important in choosing a proper forecasting method.

The findings presented by the authors are very much in line with the proposed thesis research, and can thus form a fundamental basis for the methodology to be executed at a later stage. Especially the fact that there is a way that implementing these insights could contribute to improvements in practice, is an important result that further solidifies the need for the proposed research, and it is an indication that the research could lead to promising results.

The paper called *"Reliability and operations: keys to lumpy aircraft spare parts"*, written by A. F. Lowas III and F. W. Ciarallo in 2015 [6] is one of the first research papers that mainly focuses on the reasons for aircraft spare parts to show lumpy demand patterns. These insights are then used to provide suggestions on how to improve the regularity of spare parts demand, thus allowing opportunities to improve forecasting accuracy. The authors start their paper by reviewing existing studies that deal with the difficulties involved in forecasting for lumpy demand patterns. Like other authors, Lowas III and Ciarallo investigate the existing forecasting methods and how to deal best with intermittent or lumpy demand patterns, as they summarise the main findings of existing research.

The main objective of Lowas III and Ciarallo's research is stated to be to empirically demonstrate the underlying factors for lumpy spare parts demand, by uncovering probable reasons that affect the lumpiness of spare parts demand. Furthermore, the authors use Weibull-based models to simulate the failure of (and therefore, demand for) replaceable aircraft components. Since 93% of non-structural components are cited to exhibit a constant failure rate [26], the failure probability density function can be modeled by an Exponential function with constant failure rate. With that, the scope of the research is limited to non-repairables components fitting the Weibull distribution of failure models.

In this research, Buy Period (BP) is considered to be an inherent characterizing factor that may impact demand lumpiness, and it is assumed that each aircraft has a life of 20 years, and the aircraft in the fleet are acquired evenly over a BP of 1, 2, 4, 8 or 16 years. Another characterising factor that is considered is the Fleet size, which is assumed to consist of 8, 32, 128, 512 or 2048 aircraft. Each simulation will also be replicated 50 times to ensure statistical significance, and with 3000 unique combinations of the previously mentioned variables and Weibull parameters, the total number of simulations will amount to 150,000.

The results showed that 76% of the cases had output that could be characterised to be lumpy. Based on the results, it can also be stated that there is a strong correlation between ADI and CV, meaning that a higher ADI will usually also come with a higher CV. The appropriateness of selecting a Weibull distribution to simulate the results is also proven by fitting the Weibull graphs onto actual engineering data for a C-135 ruddervator, F-15 speed brake and a F-15 radome. The Monte Carlo model results are compared to actual demand histories in Figure 2.8, which shows that indeed for these types of components, the simulated demand can be assumed to be accurately modeled with Weibull distributions.

The authors also have findings related to the effects of the aforementioned factors: fleet size, buy period and as-built component life. According to the authors, fleet size is the most significant single factor impacting the lumpiness of demand, with smaller fleets having dramatically higher CV and ADI values than larger fleets. Additionally, it is stated that a fleet size of at least 256 will enable a fleet planner to anticipate that failures will occur every quarter, with minimal variability of demand, thus making the total demand pattern less lumpy and less challenging to forecast.

The final paper to be discussed in this Literature Review section is the most recent research paper also considering underlying demand generating factors, which is applied to improve forecasting methods. The paper is called *"Forecasting spare part demand with Installed Base information: a review"*, written by S. van der Auweraer, R. Boute and A. A. Syntetos [27] and it aims to mainly provide a literature review on installed base forecasting methods.

The authors suggest to work with *Installed Base information*, which does take more factors into account than just historical demand. The benefits of using installed base information are emphasised, with the authors referring to previous work stating that the use of installed base information to forecast spare part demand can lead to cost savings up to 58% [28].

The authors mainly have the objective to present a summary of existing and relevant literature in similar fields, in order to motivate future researchers to consider installed base information as means of forecasting for spare parts. The authors first describe the most prominent part characteristics that

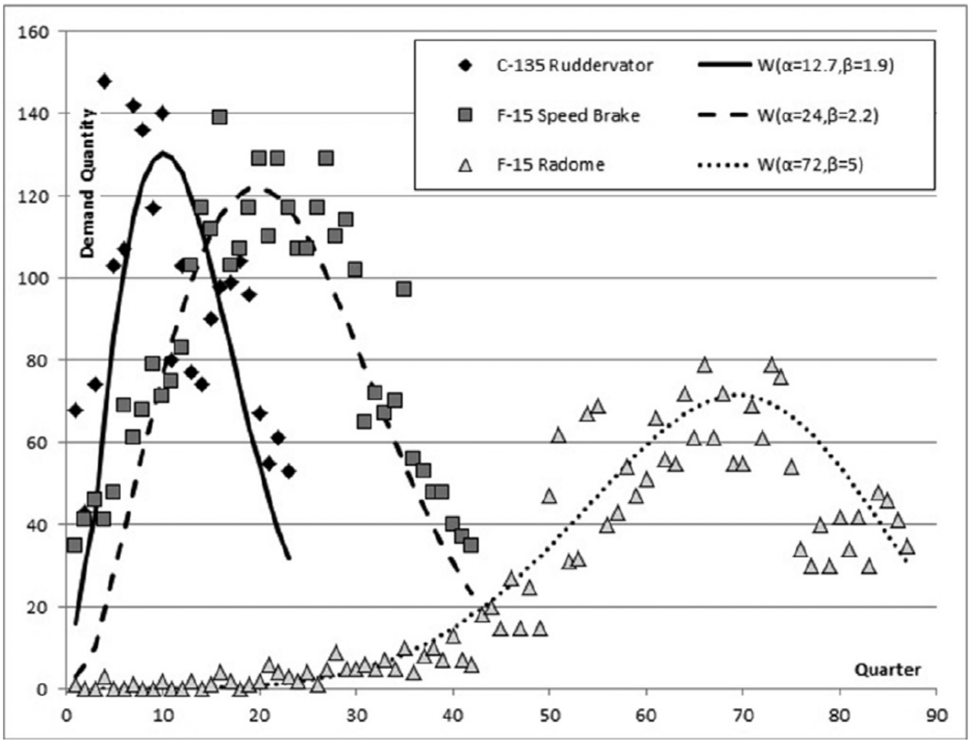

Figure 2.8: Weibull fits of spare parts demands on actual historic data [6]

cause difficulties in forecasting spare parts demand to be as follows:

1. Part demands show very particular patterns
2. They are generated by maintenance policies and part breakdowns
3. Parts tend to have a limited amount of historical demand data available
4. They are subject to obsolescence

From an installed base perspective, the key drivers of spare parts demand are the maintenance activities, which is very different from preventive maintenance spare parts management. With that in mind, the authors proceed to explain how to use installed base information to forecast CM demand. They provide the governing equations used to express the expected demand under four conditions: constant installed base, increased installed base, decreasing installed base and fluctuating installed base. The authors cite other authors that state that all four modifications of using installed base information can be deemed appropriate methods (in some way or another) of forecasting demand and investigating the causal factors.

In comparing corrective maintenance with preventive maintenance, the authors state that CM is characterised by a stochastic arrival of demand, while the demand size is deterministic. In the case of PM the arrival of demand is deterministic, while the demand size can often be stochastic. The authors thus suggest that the use of installed base information might be more suited for unplanned corrective maintenance.

According to the authors, their research shows that rich information can be made available to improve spare parts demand forecasting. They state that the use of causal methods is appealing, but the application of the presented information is not exclusive to causal methods alone. It is for example also possible to use time series models in combination with installed base information.

In the research performed by B. Hellingrath and A. Cordes [29], a time-series method is combined with causal information. The authors implement data generated from an Intelligent Maintenance System, which is a physics-based model that considers physical characteristics of individual components and relates those characteristics to the probability of failure of that component.

In their research, Hellingrath and Cordes focus on integrating the IMS data with the SBA method proposed by Syntetos and Boylan in 2001, due to its proven accuracy under lumpy conditions [18]. They execute this by using the output of the IMS data as input for determining the parameters of the underlying pdf distribution of the SBA method. The authors also estimate the demand values forecast by the SBA method without taking into account IMS data, after which both series of results are compared with each other and with the actual demand data to draw conclusions regarding the accuracy of both methods.

From the obtained findings it was found that when IMS data is included in the forecasting, the estimated demand values are in fact closer to the actual demand values, compared to the forecasting method that did not include IMS data. This is a very interesting finding, as this confirms for this specific data set that considering and implementing underlying causal factors does in fact improve forecasting accuracy. It should be noted however that only five different types of spare parts were forecasted, and these results may not necessarily hold true for all aircraft spare parts in general.

These results do reinforce the fact that the integration of underlying factors and information could benefit the accuracy of existing forecasting methods, thus justifying further research in this specific area. Therefore, the integration of causal methods or underlying factors with existing time-series methods is rightfully so a major focal point of this master thesis.

2.1.5 Main shortcomings in current state of the art

Combining the main takeaways of the reviewed literature of all three categories, it can be said that the proposed thesis research will contribute a novel addition to each of the three discussed categories. Following from the initial statistical analysis, the most statistically significant factors will be implemented in the second phase of model building and adjusting.

This is also an aspect that is rarely performed in existing research. Many of the sources describe methods to improve forecasting accuracy by changing or updating the models themselves, but this is often done without taking into account the underlying causal factors. The research of Hellingrath and Cordes [29] have successfully executed this, although the scope of their research was limited to a small data set of spare parts.

Throughout the literature review, it was found that not many academic articles deal specifically with the subdomain of both investigating inherent causal factors and implementing them to improve spare parts forecasting methods. This imposes some difficulties in defining the current state-of-the-art and how the existing academic knowledge can be used to devise an appropriate research methodology for this specific issue. This does however emphasise the fact that this is actually a very novel research area, and many improvement opportunities still exist in this area.

The few research papers that have touched upon this area have shown promising results with respect to improvement of forecasting methods if additional factors are considered. If the research objectives and questions as described in Section 2.2 can be satisfied properly, significant contributions can be made to the existing academic and industrial state-of-the-art by this thesis project.

2.2 Research scope and research questions

This section details the description of the project scope of the thesis following from the identified gaps in the literature review. It will outline which elements will be of importance during the execution of the research, and which topics will be considered. The project scope will be limited to the proper execution and research of four main pillars, which will be described in the first subsection. After this, the research questions according to the research scope will be presented in the second subsection.

2.2.1 Description of project scope

The first pillar of the thesis research is the extraction of specific data sets from the MRO data base and the identification of the inherent characteristics of the demand patterns of aircraft spare parts. The main issue in forecasting spare parts is not that the existing forecasting methods are inadequate, but that they are very inaccurate for forecasting demand patterns with high variety. Therefore it can be of significant importance to first understand which elements generate demand in an MRO and why the demand size and frequency is so varied. If these insights can be identified, they can provide opportunities to improve the effectivity of existing forecasting methods.

The second pillar of the thesis research concerns itself with the selection of existing forecasting methods to be used as a baseline forecasting method. This baseline forecasting method will be applied to the specific data sets extracted in the initial phase of the research. Furthermore, this baseline method will be altered according to the insights gained in the previous pillar; the causal factors will be incorporated with the adjusted forecasting methods. The altered forecasting methods will then also be applied to the selected data sets.

The third pillar of the project scope is to measure, evaluate and compare the performance of the selected baseline and adjusted forecasting models. This pillar will be where the findings of the previous two pillars come together, and based on the results it will clarify whether or not the incorporation of the identified driving factors has in fact had a positive impact on the forecasting accuracy. This step will yield the main results of the research, and based on these results recommendations can be provided regarding future implementations and development.

The fourth and final pillar will be dedicated to validating the approach through the use of data sets of additional component categories within the MRO database. The applied approach in the first three pillars will be repeated for a selection of validation data sets, which will yield the main general conclusions of the thesis. At this stage, a sensitivity analysis will also be performed to assess how slightly changing the assumed model parameters may impact the general conclusions.

2.2.2 Formulation of research questions

Based on the four main pillars of Project Scope laid out in the previous subsection, the main objective of the thesis research will be *"To demonstrate that aircraft spare parts demand forecasting accuracy will improve when inherent causal factors are taken into account while forecasting with time-series methods"*. To reach this objective, multiple research questions will have to be answered throughout the research. The main research question that the thesis research aims to answer is the following:

- ***Will spare parts demand forecasting accuracy improve if inherent causal factors are taken into account while forecasting with time-series methods?***

These main question in turn also generates multiple secondary research questions, which can be answered subsequently in order to find answers and conclusions for the primary research question. These questions will form the underlying framework of the methodology to be applied, where answering the secondary research question will eventually lead to findings that answer the primary research question, and as such the objective of the thesis research will be achieved. The list of secondary research questions is listed as follows:

1. Which underlying causal factors can be identified to have a significant impact on the endogenous demand patterns?
2. What is the statistical significance of these factors regarding impact on specific component

removals in the data base?

3. How can an existing forecasting model be altered to incorporate the effect of the key causal factors?
4. Which error measure is suitable to be used to evaluate the forecasting accuracy of the chosen forecasting model?
5. What is the forecasting accuracy of the selected model in its baseline conditions, without taking into account the causal factors?
6. What is the forecasting accuracy of the selected model in adjusted conditions, taking into account the causal factors?
7. Can an improvement of accuracy be established when comparing the baseline forecasting method with the adjusted forecasting method?
8. Does using a different data set of aircraft spare parts components result in similar findings, thus validating the suggested approach?

CHAPTER 3

Methodology

With the academic state of the art outlined and the research domain described in the previous chapter, it is now relevant to introduce the main methodology applied to the thesis research. This chapter will therefore outline the main functions, inputs and outputs of the methodology that is applied to satisfy the proposed research questions. Section 3.1 will describe the model that was applied, after which Section 3.2 will detail the selected forecasting methods and error metrics. Finally, Section 3.3 will present the proposed methodology for the altered forecasting method, and Section 3.4 will describe the Verification and Validation strategy that was applied in this research.

3.1 Model description

The model that was built to be applied in this research is quite extensive and includes multiple inputs, outputs and functions to generate the required results. This section will outline the details of this model and will describe the general flow of actions that is applied in this model.

Figure 3.1 shows the flow of operations in the initial phase of the model. This phase mainly deals with the selection of specific datasets within a big database. The main input for this module is the database provided by the MRO, which contains a large number of data entries for component removals since the 1930s. The first step is to filter this data to a more recent timeframe, so that the results are more useful in the operations of the MRO. This timeframe is set between 2008 and 2015, to initiate the model with sufficient data and to use more representable and consistent data in recent years.

After this, the data is selected and split further based on operator type, aircraft type and component type. Finally, since all the component removals in the data base are registered on a specific day of the month, it is necessary to generate monthly quantities for the component removals and the causal factors to be analysed. This entire process will generate as outputs the monthly patterns of component removals and the causal factors in the timeframe between 2008 and 2015 for specific operators, aircraft types and component types. The results of this phase are presented and discussed in Chapter 4.

The patterns generated in the first phase of the model will be used as inputs for the second phase of the model, which is depicted in the flowchart shown in Figure 3.2. The second phase of the model mainly deals with identifying any statistical relations between the component removals and the selected causal factors. First, a data scatter is created for component removals vs. the causal factors, after which the Pearson's correlation coefficient is computed to find out if there exists a statistical relation between

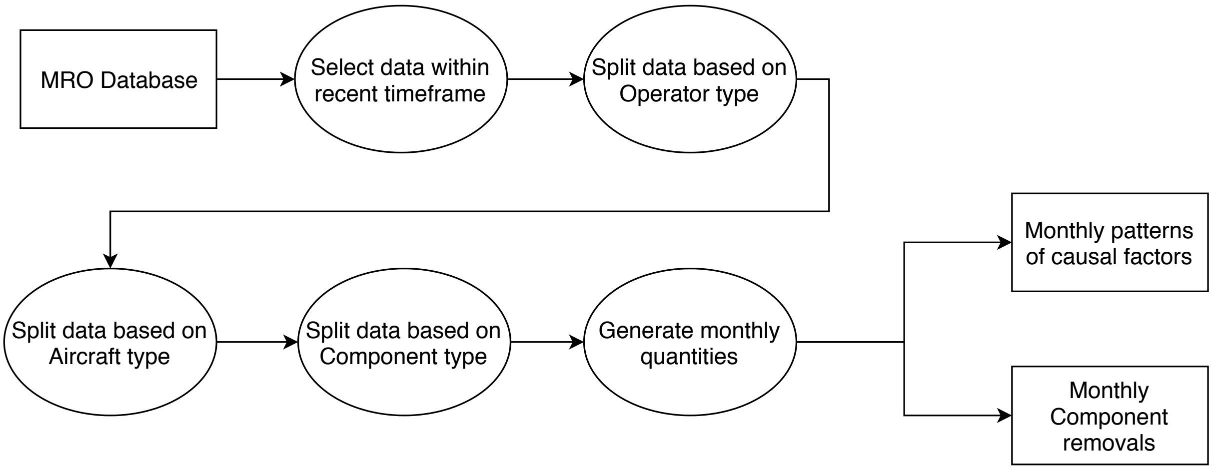

Figure 3.1: Methodology of the first phase of the applied model

the component removals and the causal factors. The main output of the second phase of the model are the values for correlation coefficients, which will be used as inputs the third phase of the model. The results corresponding to the second phase of the model are discussed in Chapter 5.

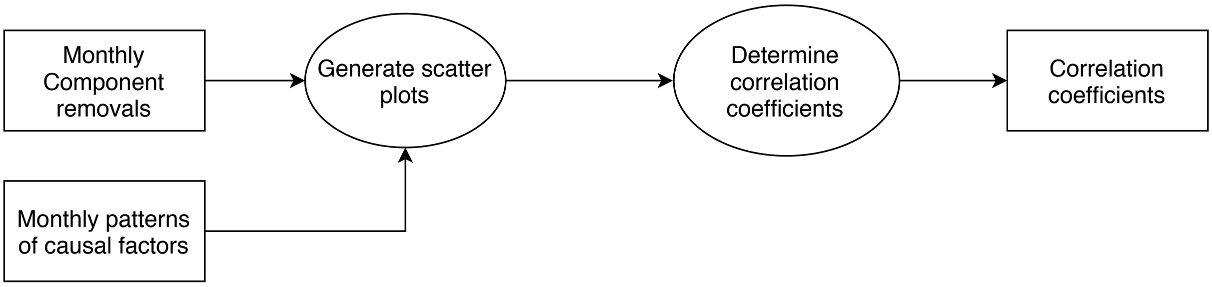

Figure 3.2: Methodology of the second phase of the applied model

Finally, the outputs generated in phases 1 and 2 of the model will be used in the third phase of the model. The process of this phase is shown in Figure 3.3. This phase is the most important aspect of the applied model, since it deals with applying and evaluating the baseline and adjusted forecasting methods. First, the baseline forecasting methods are applied, after which the predicted demand volumes will be compared to the actual demand to compute the forecast error.

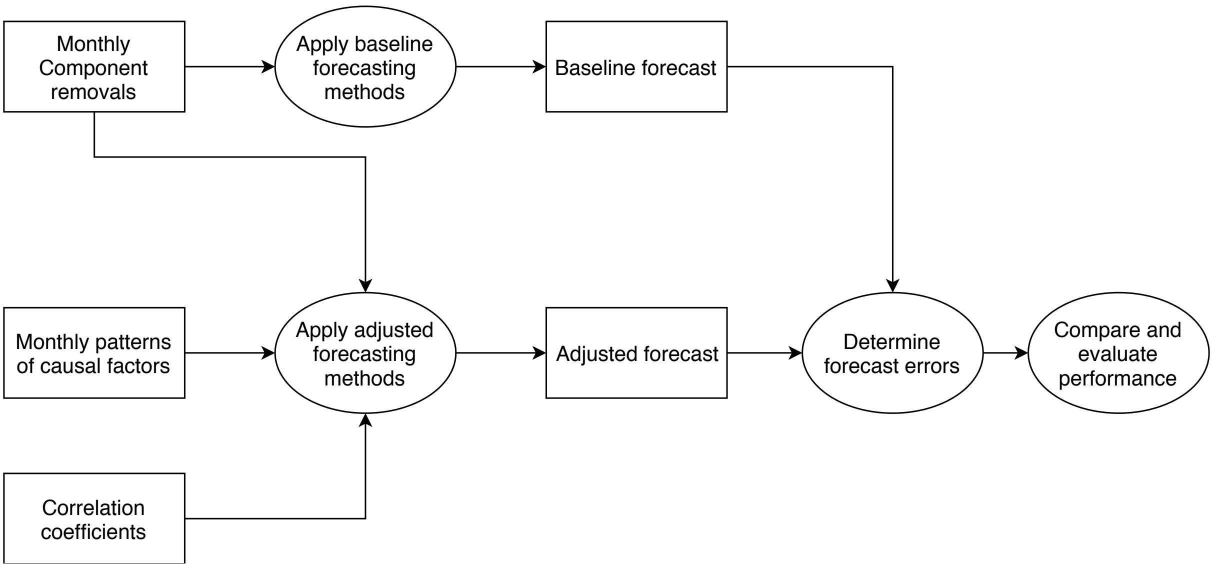

Figure 3.3: Methodology of the third phase of the applied model

Next, using the patterns of component removals, the causal factors and the corresponding correlation coefficients, the adjusted methods are applied to the selected data sets. This will generate an adjusted forecast, which will also be compared to the actual demand volumes to compute the forecast errors. Finally, the performance of the baseline methods will be evaluated and compared to the performance of the adjusted methods, to determine which methods are the most accurate in forecasting the spare parts. The results of this phase of the model will be presented and discussed in Chapter 6.

3.2 Description of baseline forecasting methods, error metrics and causal factors

This section will detail which forecasting methods will be applied to the selected datasets. Two relevant forecasting methods will be selected as baseline methods, which will both be applied to the endogenous demand data sets of the most common and relevant component categories. The baseline methods will be the Moving Averages (MA) method and the Single Exponential Smoothing (SES) method, which are time-series methods that are commonly used to forecast the demand of spare parts in practice. The reason that time-series methods are used for this research, is because time-series methods are more suitable for short term forecasting and are computationally less demanding compared to stochastic models. Even though scientific literature shows that Croston's method is the most applicable method in forecasting lumpy demand patterns, in this case SES will be a suitable alternative since no zero-demand months exist in any of the component removal data subsets.

3.2.1 Moving Averages method

The MA method takes the average of the last m values of a time series to determine a forecast value [30]. Equation 3.1 shows the mathematical relation that governs the Moving Averages forecasting method.

$$f_{t+1} = \frac{1}{m} \sum_{k=0}^{m-1} y_{t-k} \tag{3.1}$$

In this equation, m is the user-set parameter that determines how much historical demand is included in defining the average. A smaller value for m leads to a more reactive forecasting method. For the purpose of initialising the baseline forecasting methods for this specific research, m is set at a value of 3 (months). This means that the MA forecasts presented in Section 6.1 use the (moving) average value of the previous three months to determine the forecast value for the upcoming month.

3.2.2 Single Exponential Smoothing method

The SES method is one of the most accurate forecasting methods when forecasting aircraft spare parts demand data. It takes the forecast error into account and adjusts it with a certain smoothing constant α. Equation 3.2 [31] shows the governing mathematical equation for the SES forecasting method.

$$f_{t+1} = \alpha y_t + (1 - \alpha) f_t \tag{3.2}$$

The smoothing constant is essential in determining how reactive the SES method is to its own forecast errors, with a higher α leading to a higher reactiveness to the forecast error. Usually this value is between 0.1 and 0.3 [31], but for the purpose of applying the baseline methods in this thesis research, an α of 0.3 has been assumed, since the component removal data sets to be forecast are very volatile in their demand size.

3.2.3 Root Mean Square Error metric

To evaluate the forecasting performance of both the MA method and the SES method, the Root Mean Square Error (RMSE) will be measured and compared. The RMSE is an error metric that sums the squared error values of each forecast, and then takes the root of this sum. In doing so, the RMSE shows the magnitude of the overall error that has been made by the forecasting methods. Equation 3.3 shows the mathematical relation that was used to determine the RMSE for each forecast demand data set.

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n} {e_t}^2} \tag{3.3}$$

3.2.4 Mean Absolute Percentage Error metric

In addition to assessing the RMSE values of each forecast, the MAPE will also be determined for each forecast. In contrast to the RMSE metric, the MAPE metric is not scale dependent, so it provides a better estimate of the forecasting performance when comparing multiple methods in various databases, since the overall demand size does not have to be taken into account. The RMSE gives a restricted sense of the overall performance of the forecasting method, if the scale of the demand sizes are not taken into account.

The drawback of using the MAPE metric is that it is only applicable to demand data sets without any zero-demand months, while the RMSE metric is suitable for all demand patterns. The mathematical relation that was used to determine the MAPE value is given by Equation 3.4.

$$MAPE = \frac{1}{n}\sum_{t=1}^{n} |\frac{e_t}{y_t}| \cdot 100 \tag{3.4}$$

3.2.5 Selection of causal factors

The causal factors to be selected are factors that may have an impact on the demand pattern of a specific component. Of course there are many factors in aircraft maintenance that may impact the demand generation of component removals. Possible examples of these factors can be listed as follows;

- Environmental effects
- Flight cycles
- Pilot complaints
- Fleet size
- Time to failure of a component
- Aircraft landings
- Operator type
- Aircraft type
- Characteristics of component type
- Maintenance policy of MRO

For the scope of this research however, the methodology will limit itself to the implementation of the causal factors pilot complaints and aircraft landings only. The main reason for this is the abundance of available data in the MRO database on these specific factors, and because it is feasible to expect a statistical correlation between the number of pilot complaints, the utilitisation rate of the aircraft and the number of removed components. Additionally, the operator type, aircraft type and component type will be used as factors to segregate the data in the initial phase, prior to generating the monthly patterns of the causal factors and component removals. A more elaborate motivation for the selection of these causal factors is provided in Section 5.1.

3.3 Methodology for altered forecasting method

An approach to improving the existing methods with additional insights, is to somehow incorporate the correlation coefficients obtained for the causal factors with the forecast demand output. The correlation coefficients describe how strongly the component removal data would follow a relative change in the causal factors. It is therefore also necessary to include the ratio of pilot complaints and aircraft landings in the current month, over the average value for these factors in the past three months.

Multiplying these ratio's with the correlation coefficients for the causal factors, will tune the forecast value either upwards or downwards. For example, in case a certain month relatively has a lot of pilot complaints and aircraft landings, the forecast demand output obtained from the MA or SES method will be tuned upwards. In case there are relatively very few pilot complaints and aircraft landings, the forecast value will be tuned downwards.

The hypothesis is that this tuning effect will reduce forecasting errors, since additional statistically significant explanatory factors are taken into account. Equation 3.5 shows the governing relation that will be used to improve the baseline forecasting methods. This improved forecasting methodology will be implemented and applied in Chapter 6.

$$F^* = F' \cdot \left(\frac{c_{PC} \cdot [PC_1/PC_0] + c_{LD} \cdot [LD_1/LD_0]}{c_{PC} + c_{LD}} \right) \tag{3.5}$$

Where;

- F' is the MA or SES demand forecast value
- c_{PC} is the correlation coefficient for Pilot Complaints
- c_{LD} is the correlation coefficient for Aircraft Landings
- PC_1 is the number of Pilot Complaints in the current month
- PC_0 is the average number of Pilot Complaints in the past three months
- LD_1 is the number of Aircraft Landings in the current month
- LD_0 is the average number of Aircraft Landings in the past three months

3.4 Verification and Validation methods

In order to ensure that the findings and conclusions are representative of reality, the suggested methodology also needs to be verified and validated throughout the research. This section will therefore briefly explain the applied verification and validation strategies.

3.4.1 Verification strategy

The approach is mainly verified in the initial stages of the research. The main objective of the initial phase of the model is to correctly load and select data from an Excel environment into the MATLAB enivornment. To ensure that this goes without errors, the method is verified by recalculating the results found in MATLAB with Excel. The verification of the model is thereby applied by looking into any discrepancies in the results generated by MATLAB and those in Excel. In case the results are the same, the approach is deemed to be verified successfully.

For example, during the initial phase of the model, most of the data will be selected and imported to generate the data patterns. The results in the data patterns are then verified by confirming that excel yields the same quantities for random months in the time span of eight years. This process is repeated several times for random months, and if the results are equal then it is verified that the model successfully is able to import and handle the data base stored in Excel.

3.4.2 Validation strategy

The applied model also needs to be validated to ensure that the results and conclusions are not applicable for one specific situation and set of requirements only, but that the approach also is able to successfully generate results under other conditions. In case of the described model, the validation will be applied in the latter phase of the analysis. In the last phase, initially the baseline and adjusted forecasting methods are applied to a few specific component types for a specific operator.

After the results and conclusions are generated for these specific component types, the whole process will be repeated for components in other categories as well, thus validating the approach if the additional findings are supporting the initial conclusions obtained. The validation strategy is therefore basically applying the whole model again under different conditions for multiple different component types. If the results of the validation are in line with the conclusions in the initial research, the method can be deemed successfully validated.

CHAPTER 4

Data selection and analysis of historic demand patterns

The first phase of the project concerns itself with the selection of proper databases and the identification of the main demand characteristics of these selected databases. Using an extensive database with endogenous and exogenous spare parts demand data from an existing MRO, the specific demand patterns are visualised and extracted. Section 4.1 will describe the preliminary analysis that was performed on the MRO database, which will be followed by an analysis of the demand size variation as described in Section 4.2.

4.1 Preliminary analysis MRO database

The first step was to actually analyse the available data on a preliminary level, such that an effective and more complex analysis of the demand pattern characteristics and statistically significant factors can be performed using MATLAB. To ensure the generation of results effectively without jeopardising the computational efforts required, the large initial database was split into a smaller data sets with a sample space that could guarantee efficient yet thorough analysis.

First of all, the historic timeframe of the sample space to be analysed was restricted to a timespan of eight years, investigating all components removals occurring between January 2008 and December 2015. This period was mainly chosen due to the fact that it became apparent from the database, that the more recent years contain more relevant, complete and consistent data. The quality of consistent data was the main motivating factor in deciding which database to start the general analysis with. For the remainder of the thesis research, this sample space is also restricted to demand data between January 2008 and December 2015, to ensure more qualitative and consistent results and conclusions.

The monthly removals of all components for Fokker 50, Fokker 70 and Fokker 100 aircraft are combined and represented in Figure 4.1. At a first glance, looking at the removal of all components by Fokker between 2008 and 2015 leaves the impression that the demand patterns could be characterised to be smooth. However, this is mainly explained by the fact that every single component type (CT) is pooled together in this case, which negates the erratic nature of individual components in a bigger collective group of components. In practice, any useful forecasting method would be applied for individual components (categories), rather than for the entire batch of components in the inventory. This emphasises the fact that the demand for spare parts should always be considered on a more detailed level.

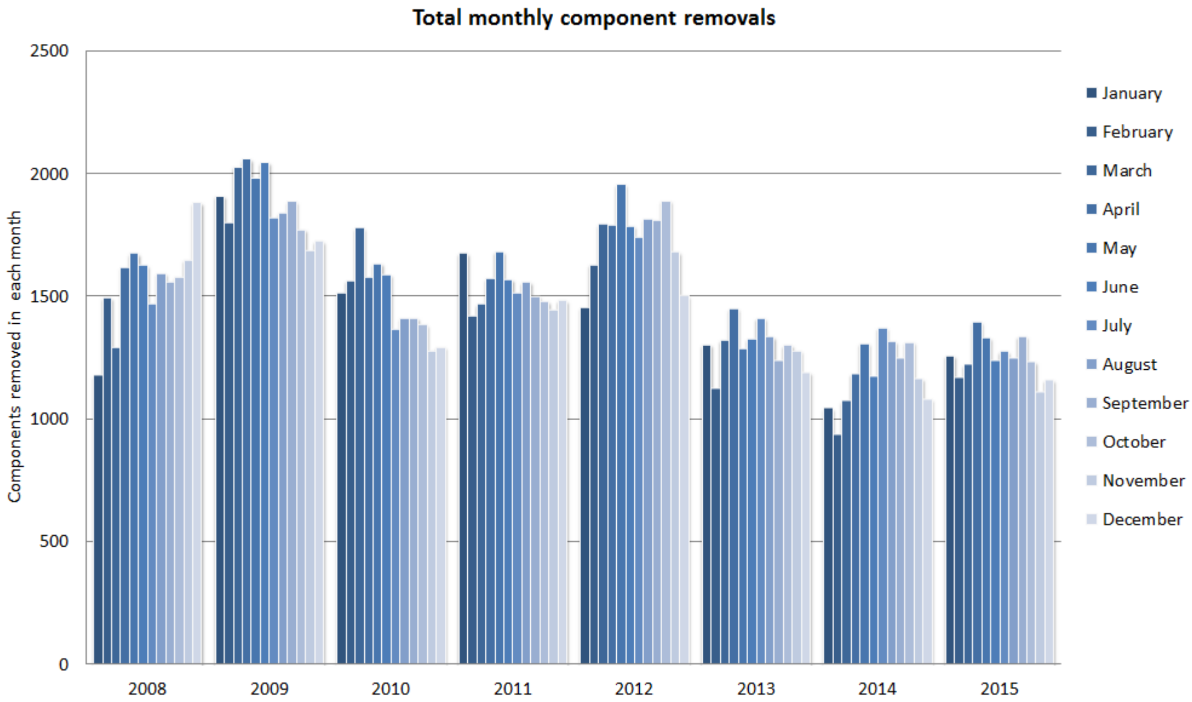

Figure 4.1: Component removals of all categories and aircraft types between 2008 and 2015

It is expected that the demand patterns for more detailed component categories will show a more erratic distribution compared to the demand pattern for all components combined. For this reason, all the components in the sample space are grouped together with their ATA3-chapter code for each aircraft type in the database, to identify the demand variation of certain component categories. To do this, separate data subsets were generated for each ATA3-component for each aircraft type, yielding the component removals per month over the same timespan of eight years.

For illustrative purposes, Figure 4.2 shows the demand pattern for all components in the ATA-342 chapter for the Fokker 100 aircraft type, which represents the *Attitude and Direction* section of the *Navigation* category. Just by looking at the distribution of demand data for this specific set of ATA-342 components, it can already be seen that the demand volumes show a larger variation compared to the demand pattern of all components depicted in Figure 4.1, with monthly demand sizes ranging between 12 and 44. This is of course in line with the expectation that the demand patterns for more detailed components will show a less smooth distribution than for all components combined.

With the developed preliminary analysis method, similar graphs are generated for each possible ATA3-component and aircraft type combination existing in the identified timeframe. From this preliminary analysis, it becomes apparent that the majority of the ATA3-categories of components have a commonality of less than **1%**. To ensure measurable advantages of the newly proposed forecasting method and to minimise the computational efforts, a pre-selection was made of ATA3-categories with a commonality higher than **3%**. This subset of component removal data accounts for almost **40%** of the total removals and it will be used for further analysis, since any forecasting benefits found in this subset could lead to significant impacts on the overall operations, due to the significance of these six ATA3-categories of components.

This subset of data was then imported into **MATLAB** for further analysis of the demand pattern characteristics. The CV^2-values for each of the ATA3-component datasets was determined, to gain insights on the degree of demand size variation for each component category. Table 4.1 shows an overview of the overall CV^2-values for each of the six components, per aircraft type. The CV^2-values are calculated over the same period of eight years, between 2008 and 2015.

Looking only at the CV^2-values presented in Table 4.1, it can already be seen that all of the values

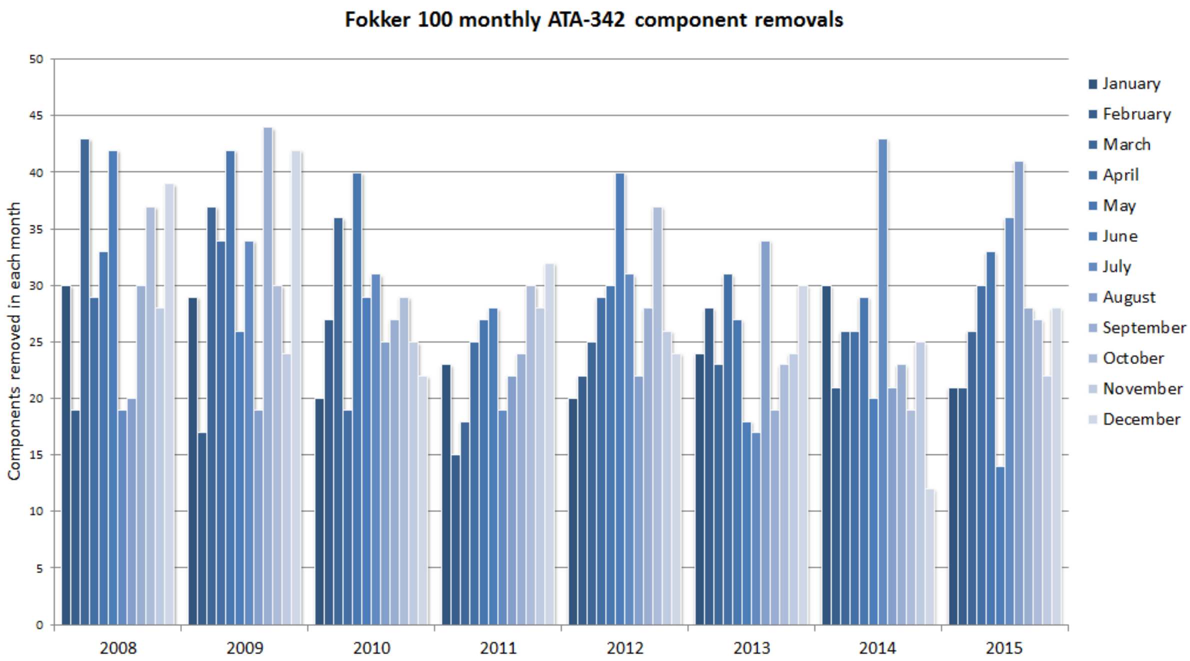

Figure 4.2: Component removals of the ATA-342 category for Fokker 100 aircraft between 2008 and 2015

Table 4.1: Overview of CV-values for the six ATA3-component demand patterns between 2008 and 2015

Component category	Description	Commonality (%)	CV^2-values		
			F-50	F-100	F-70
253	Equipment/Furnishing - Buffet/Gallery	3.2	5.78	1.88	2.55
324	**Landing Gear - Wheels and Brakes**	**20**	**1.76**	**2.32**	**2.14**
334	Lights - Exterior	3.0	2.48	1.71	2.98
342	Navigation - Attitude and Direction	4.4	1.82	2.03	2.8
345	Navigation - Dependent Position Determining	3.1	2.09	2.09	3.60
351	Oxygen - Crew	5.2	2.01	2.45	1.16

are over the threshold CV^2-value of **0.49**, implying that all of the demand patterns in this subset show irregularity in the demand size. The ADI-values on the other hand are 1.00 for all datasets, since there are no zero-demand months and therefore every month demand is to be expected for each of the six presented component categories. This is due to the fact that still many different types of unique and more detailed components are included under the ATA3-chapter level.

Looking at the 324 category for example, Landing Gear - Wheels and Brakes includes components ranging from nose and main wheel tires to the brake valves and sensors. For this reason, a more detailed analysis is performed for one of these component categories on the ATA6-level. This approach separates the parts based on their ATA6-codes, and thus this separation is based on more detailed characteristics/categories of the individual components. As such, a more detailed analysis is performed on the ATA-324 component category (Landing Gear - Wheels and Brakes), mainly because this category contains the most components by far, and therefore the achieved results in this subset of data could lead to more substantiated implications in the total database.

4.2 CV^2-analysis of Wheels and Brakes components

The 324-category in the ATA3-chapter description represents the Wheels and Brakes group of components within Landing Gear. This section will present the results gained from a more detailed analysis of demand patterns for the Wheels and Brakes category specifically. It should be noted however that many subcategories of the 324-category (Wheels and Brakes) are represented within this particular data subset with a commonality of less than **1%**.

It is for that reason that all components at the ATA6-chapter description with a commonality of less than **10%** are initially omitted from the data subset. The following ATA6-chapter subcategories which remain in the data subset for a more detailed analysis of CV values are presented in Table 4.2.

Table 4.2: Overview of ATA6-categories for detailed CV analysis

ATA6-category	Description	Commonality
324-101	Main Wheel Tire (MWT)	43%
324-103	Nose Wheel Tire (NWT)	24%
324-201	Main Wheel Brake Unit (MWBU)	16%

As can be seen from this overview, the vast majority of components in the ATA3-324 category belong to the subcategories 324-101, 324-103 or 324-201. It should also be noted that the sum of these three subcategories accounts for **83%** of all the data within the 324 category, which is in accordance with the 80/20 Pareto-rule. All other component subcategories in the Wheels and Brakes category amount to a volume size of less than **20%** of the total.

With the data subsets now defined, the more detailed CV-analysis can be performed for the three different Aircraft Types (ACT) between the period January 2008 until December 2015. It is assumed that the **ADI values will always be equal to 1**, since in each month there will be a quantity of at least one component. Therefore during the detailed analysis, only the CV^2-values will be of interest and the degree of demand size variation will be assessed purely on whether or not the CV^2 will be above (thus erratic) or below (thus smooth) the threshold value of 0.49.

For each data subset, the standard deviation is divided by the mean of that subset, to retrieve a value for the CV^2. This process is applied to the three subcategories (324-101,324-103 and 324-201) for three aircraft types (Fokker 50, Fokker 100 and Fokker 70), yielding nine different CV^2-values for all of the analysed data subsets. An overview of these results is presented in Table 4.3.

Table 4.3: CV^2-values for selected ATA6-categories

ATA6-category	Subcategory description	CV^2-values		
		Fokker 50	Fokker 100	Fokker 70
324-101	Main Wheel Tire (MWT)	0.985	1.63	0.995
324-103	Nose Wheel Tire (NWT)	0.00	1.70	1.23
324-201	Main Wheel Brake Unit (MWBU)	1.08	0.592	0.490

Looking at the preliminary results in Table 4.3, it becomes clear that almost all demand patterns in the data subsets can be classified to be erratic. The most erratic spare parts data is seen in the Main Wheel Tire (324-101) and Nose Wheel Tire (324-103) components for the Fokker 100 aircraft, followed by the Nose Wheel Tire (324-103) components for the Fokker 70 aircraft and the Main Wheel Brake Unit (324-201) components for the Fokker 50 aircraft.

Interestingly, the demand pattern for MWBU components for the Fokker 70 has a CV^2-value that is exactly equal to the threshold value of 0.49. Furthermore, the NWT components for the Fokker 50 aircraft result in a CV^2-value of exactly 0.00, which is an unlikely value for a demand pattern. Upon further investigation of the data subset, it appears that this particular data subset contains insufficient data to perform the CV^2-values computations and to verify the pattern characteristics. This subset contains 30 entries equal to 1, which is far too few data points over a time span 8 years (and thus 96 months). Considering this fact, the data subset of 324-103 for the Fokker 50 will also be omitted from the remainder of the preliminary analysis.

With the remaining 8 subsets of Wheels and Brakes component demand data, it can be stated that the demand pattern for MWBU-Fokker70 will show the least erratic characteristics, while the demand pattern for NWT-Fokker100 will show the most erratic characteristics. To gain insights on the demand patterns themselves, an overview and discussion of the 8 subsets of data is presented in Section 4.3

4.3 Demand patterns of components within Wheels and Brakes

In the previous section, it was determined through a detailed CV^2-analysis of certain Wheels and Brakes subcategories that all of the remaining data subsets showed erratic demand characteristics. This confirms the hypotheses in relevant literature studies that the majority of aircraft spare parts demand data show large variation in demand size.

In order to gain a better understanding of the underlying inherent demand-generating factors, it is important to visualise the data subsets analysed in the Section 4.2. Figures 4.3 and 4.4 show the results of the preliminary analysis of the most erratic and least erratic data subsets, respectively. The demand patterns for the remaining data subsets are represented by Figures A.1 through A.6, which can be found in Appendix A. Again, all the data subsets are from January 2008 until (and including) December 2015.

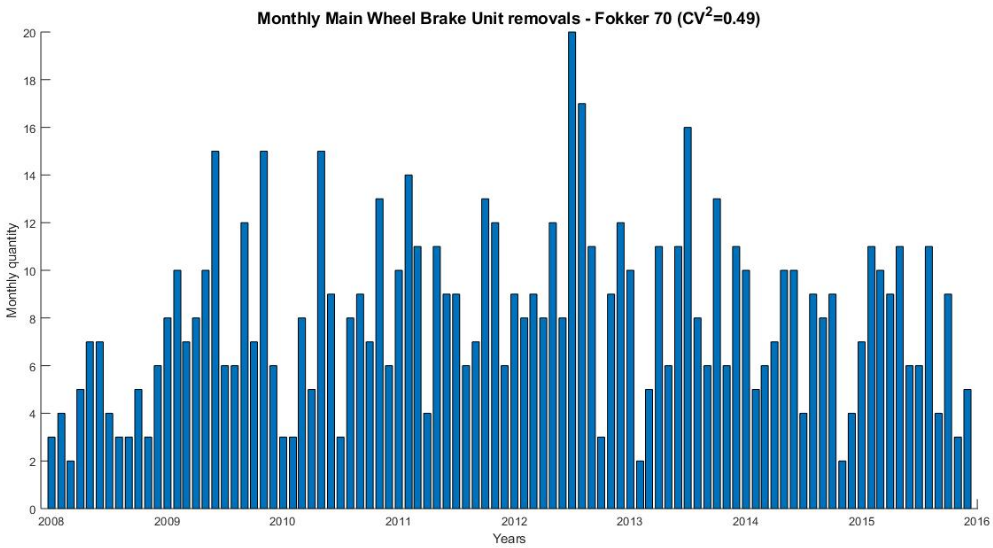

Figure 4.3: Monthly component removals of 324-201 category for Fokker 70 between 2008 and 2016

Looking at Figure 4.3, it becomes clear why this specific demand pattern is the least erratic compared to the other seven demand patterns. Almost every month in the eight year time frame has a monthly removal quantity between **5** and **15**, with outlying demand sizes being **2** at the least, or **20** at the most. For some reason, the demand size for the MWBU component of the Fokker 70 has been rather consistent, and could therefore perhaps be forecast with greater accuracy in the (near) future.

In contrast, Figure 4.4 shows the most erratic demand pattern of the eight analysed data subsets, with many large outliers of demand size in many different months. The monthly demand size of the NWT component of the Fokker 100 varies between **5** and almost **100**. This large variation in monthly demand size is of course coupled with the fact that this data subset has the highest CV^2-value of the eight investigated patterns.

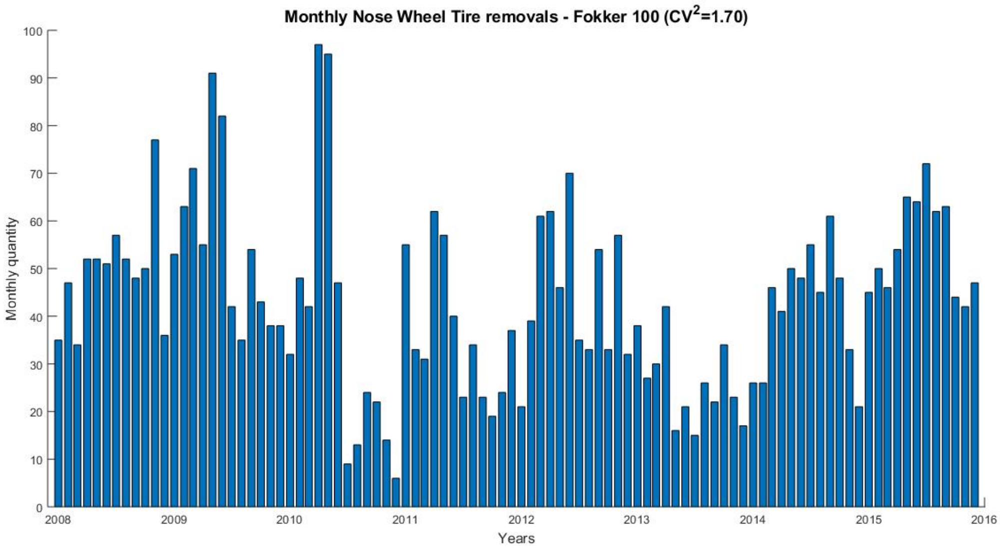

Figure 4.4: Monthly component removals of 324-103 category for Fokker 100 between 2008 and 2016

With these newly gained insights, it can be stated that the inherent spare parts demand characteristics differ per aircraft type, but also per component type. Overall, the MWBU components show less erratic demand patterns compared to the main/nose wheel tire components. In addition, the detailed CV-analysis reaches the same conclusion as the preliminary analysis; the F-100 aircraft type has the most erratic demand pattern within the Wheels and Brakes category. Also it is of interest to note that in general, the pooling and categorisation of spare part components also affects the found CV^2-values.

CHAPTER 5

Impact of Causal Factors on Component Removal data

With the most relevant components identified and their pattern characteristics analysed in the previous chapter, it is now of importance to generate data sets for specific component removals and to relate these statistically to other possible factors. Section 5.1 describes which demand-driving factors can be identified, after which Sections 5.2 and 5.3 present the results of the generated data sets for the component removals and the causal factors. Finally, Section 5.4 will describe how the generated data sets are statistically correlated to the data sets of the causal factors.

5.1 Identifying inherent characteristics of aircraft spare parts demand patterns

To investigate which demand-generating causal factors impact the main findings and observations in Chapter 4, it is important to have an understanding of how Component Removal (CR) demand is generated in the first place. The following two subsections will outline the general flow of demand generation in a typical MRO practice, and the corresponding causal factors that could potentially be incorporated to improve forecasting methods.

5.1.1 The general flow of demand data in a typical MRO

This subsection will outline the general flow of demand data and any possible key factors that may impact the monthly demand size of component removals. The general flow of demand data in the current MRO industry is depicted in Figure 5.1.

The red arrows represent the flow of actual and available data, which is present due to historic data generated in the past. The blue arrows represent demand data generated after the occurrence of component removals, which is used as input for Forecasting methods. Finally, the green arrows represent the main research gap to be investigated throughout this specific thesis work, focusing mainly on the integration of inherent demand pattern characteristics with existing high-performing time-series models.

Furthermore, the underlying factors impacting the generation of demand data is represented by green or yellow blocks. The green blocks represent internal factors that either contribute to a component removal

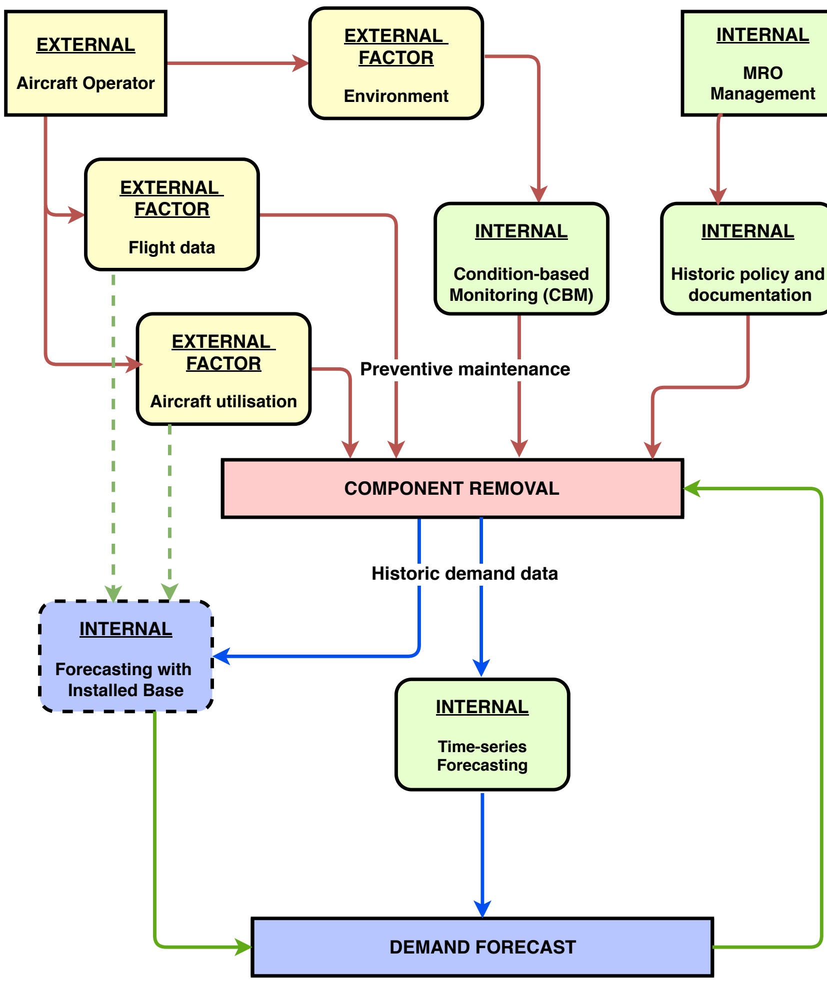

Figure 5.1: Typical flow of demand data generation in MRO industry

or a forecast of demand, based on the current state-of-the-art of scientific and industry knowledge (as summarised in Chapter 2). Specific examples of typical internal factors could include:

- The finance department of the MRO
- Upper-management of the MRO
- The specific maintenance strategy of the MRO
- Component removal guidelines in MRO documentation based on historic policy
- Forecasting methods based on historic data (Time-series forecasting)

The yellow blocks represent external factors that may have a direct or indirect impact on the generation of demand within the particular MRO, but which are outside the control of the MRO itself. Mapping out these factors is essential in understanding the inherent characteristics of aircraft spare parts demand patterns. Factors may include (but are not limited to) the following list:

- Environmental factors like weather conditions and geographical limitations.
- Flight data generated by each flight (pilot complaints, CNS data)
- Installed base of the operator (fleet size, aircraft types)
- Utilisation rate of the aircraft

Finally, the blue block represents a field of expertise that can be regarded as having the most scientific and practical contributions to the current standard forecasting models (which are often time-series forecasting methods).

5.1.2 Inherent key demand-driving factors to be analysed

The specific research field represented by the blue block in Figure 5.1 is generally directly related to improving existing forecasting models and the demand size these models forecast, and it is also the main research focus for this specific thesis project. The main objective of this research is to find patterns and correlations in historic data and take them into account while forecasting with existing methods. In doing so, the hypothesis is that the forecasting errors will be reduced measurably.

Of course many factors may directly or indirectly impact the actual number of component removals in a given time period, but for the purpose of this thesis the scope is initially limited to consider four factors. These specific factors are selected based on the available data related to these factors. Looking back at the overview of data generation represented in Figure 5.1, the initial factors of interest will be related to Pilot Complaints and the Utilisation Rate of the aircraft, due to the significant amount of available data related to these causal factors in the database. Furthermore, the data set will first be split based on Aircraft type and external Operator. As such, the key demand-driving causal factors that will be explored in this research will be:

- **Indirect factor:** Aircraft Operator (ACO)
- **Indirect factor:** Aircraft Type (ACT)
- **Direct factor:** Pilot Complaints (PIC)
- **Direct factor:** Aircraft Landings (LND)

The first two key factors (ACO and ACT) are indirect factors that inherently determine the quantity of the remaining two key factors (PIC and LND), which are in fact factors that may be directly correlated to component removals. Therefore, the indirect key factors Aircraft Operator and Aircraft Type will be used to categorise the datasets into smaller data subsets, which will then be used to analyse any correlation between the remaining two direct key factors (PIC and LND) and component removal quantities. The analysis of the aforementioned causal factors will be explored in further detail in the following sections.

5.2 Indirect factors: Aircraft Operator and Aircraft Type

The available MRO database includes many different types of operators which all have an individual demand pattern regarding Wheels & Brakes (W&B) component removals. It is therefore interesting to select a few operators before conducting a detailed analysis on their demand patterns to identify any impacts of the stated causal factors. This section will describe which operators were selected for the analysis of their causal factors, and it will present the specific demand patterns to be analysed.

5.2.1 Selection of Aircraft Operators

During Phase I of the research, there was not yet a distinction made between operators. The bulk of data was divided only based on Aircraft types and specific component categories. This means that all the subsets of demand data discussed in Phase I, represent Wheels & Brakes component removal data for Aircraft types 1, 2 and 3, for **all operators** within the eight year time span. In the more detailed second Phase of the research, inherent characteristics and correlations will need to be identified. Therefore it is of importance to select appropriate operators in which these key factors can actually be signified. The selection of these operators is based on the following two selection criteria:

1. Data continuity for Wheels & Brakes component removals between 2008 and 2015
2. Geographic diversity in selected operators

The main selection criterion of these operators is that they need to have data continuity for component removals, in order to draw solid and valid conclusions at a later stage in the research. A secondary selection criterion was that the operators need to operate in geographically diverse regions, such that any environmental differences can be associated to the presented data. With these criteria in mind, three operators were selected for further analysis.

1. **Operator 1:** Regional airline in NW Europe
2. **Operator 2:** Domestic airline in Oceania
3. **Operator 3:** Regional airline in Nordic country

All three of these operators have Wheels & Brakes component removal data continuously present between 2008 and 2015. The specific demand patterns for each of these operators for specific Aircraft Types are presented in the following subsection.

5.2.2 Monthly component removal quantities per Operator and per Aircraft type

Figures 5.2 through 5.4 represent monthly Wheels & Brakes component removal data for operators 1, 2 and 3, respectively. The presented demand patterns are separated based on aircraft type, component category and aircraft operator. It should be noted however that the component type categorisation is based on ATA3-level chapter descriptions. This means that the monthly quantities represent all components related to the Wheels & Brakes category, so no further distinction of subcategories is made at this stage of the analysis.

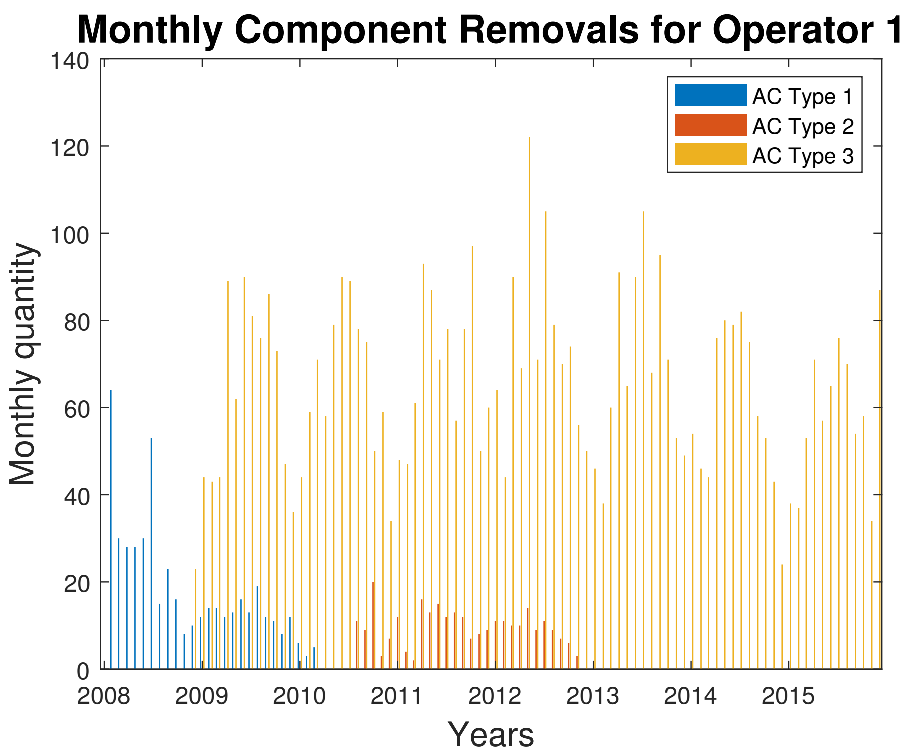

Figure 5.2: Monthly Component removal data for Operator 1

Figure 5.2 shows the monthly component removals for Operator 1, between 2008 and 2015. It is interesting to note that this operator used all three aircraft types in this time frame, but stopped using

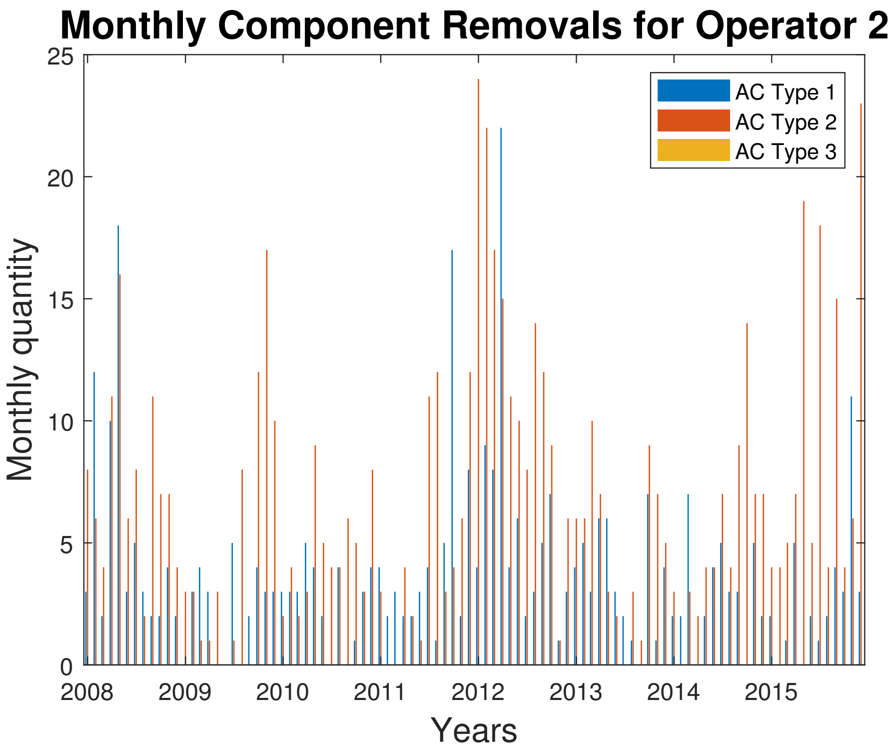

Figure 5.3: Monthly Component removal data for Operator 2

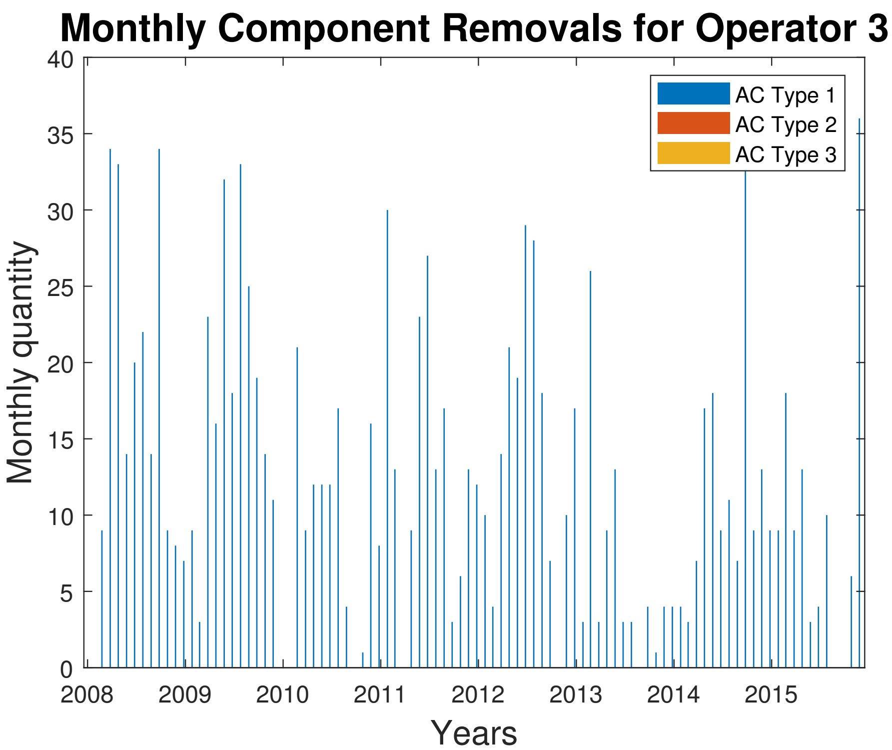

Figure 5.4: Monthly Component removal data for Operator 3

Aircraft types 1 and 2 in 2010 and 2012, respectively. It can also be clearly seen that the vast majority of Wheels & Brakes component removals is related to aircraft type 3, with monthly removal quantities ranging between approximately 30 and 120 units.

Looking at Figure 5.3, the monthly component removal distribution for Operator 2, it can already be seen that this operator has only used aircraft types 1 and 2 in the specific time span of eight years.

The data shown appears to be very erratic, with monthly Wheels & Brakes component removals ranging between one and 24 units.

Finally, Figure 5.4 shows the component removal demand pattern for Operator 3, which apparently has only utilised aircraft type 1 in the eight year time span. Even though that the other aircraft types are not present, this is still an interesting data set given its continuity. It can also be seen that this demand pattern is very erratic, with a few zero demand months, and monthly quantities ranging between zero and 35 units.

5.3 Direct causal factors: Pilot Complaints and Aircraft Landings

Before the impact of the two causal factors on the component removal rates can be assessed, it is necessary to gain an understanding of the historic patterns of these factors themselves. For that purpose, the number of Pilot Complaints and the number of Aircraft Landings between 2008 and 2015 will be visualised in this section. The subsets of data have been retrieved in a similar fashion as the component removal data presented in Section 5.2. The data is first split based on operators and then based on aircraft types, after which monthly quantities of PIC and LND have been found for each operator, per aircraft type.

5.3.1 Historic data patterns of Pilot Complaints

The data subset of Pilot Complaints is selected for the three operators, for each aircraft type. It is important to note that these specific Pilot Complaints relate to the Landing Gear category (ATA 32) only. This distinction was made because the component removals to be analysed are related to the Wheels & Brakes category, which was the major category in the dataset. Therefore it is assumed that most of the PIC regarding Landing Gear, are related to the Wheels & Brakes category as well.

With that, the number of monthly Pilot Complaints related to Landing Gear for each Operator are represented in Figures 5.5 through 5.7. At first glance, the three figures already show a very sporadic behaviour in the Pilot Complaints data patterns, similar to the component removal patterns. Furthermore it becomes clear that Operator 1 used three different aircraft types, while Operator 2 used Aircraft Types 1 and 2 and Operator 3 only used Aircraft Type 1.

5.3.2 Historic data patterns of Aircraft Landings

Since the scope of the research is initially focused on Wheels & Brakes component removal data, it is also interesting to look at a factor that can likely be directly linked to the demand size of these specific components. Therefore the number of Aircraft Landings will be a direct key factor which will be researched in this thesis project as well. It is likely that there is a correlation between the number of Aircraft Landings and the demand size for Wheels & Brakes component removals.

Before this hypothesis can be confirmed it is necessary to generate the data subsets that contain the monthly Aircraft Landings for each operator, per aircraft type. These data patterns are generated using the same methodology as described in the previous subsection. Figures 5.8 through 5.10 represent the monthly Aircraft Landings for each operator and aircraft type.

From Figure 5.8 it becomes obvious that Operator 1 utilised Aircraft Type 3 most intensely, with monthly aircraft landings ranging between approximately 3000 and 5000 times. Also it can be seen that the utilisation of all three aircraft show rather smooth patterns, compared to the data patterns of Component Removals and Pilot Complaints.

Figure 5.9 shows that over the eight year time span, Operator 2 slightly increased the utilisation of

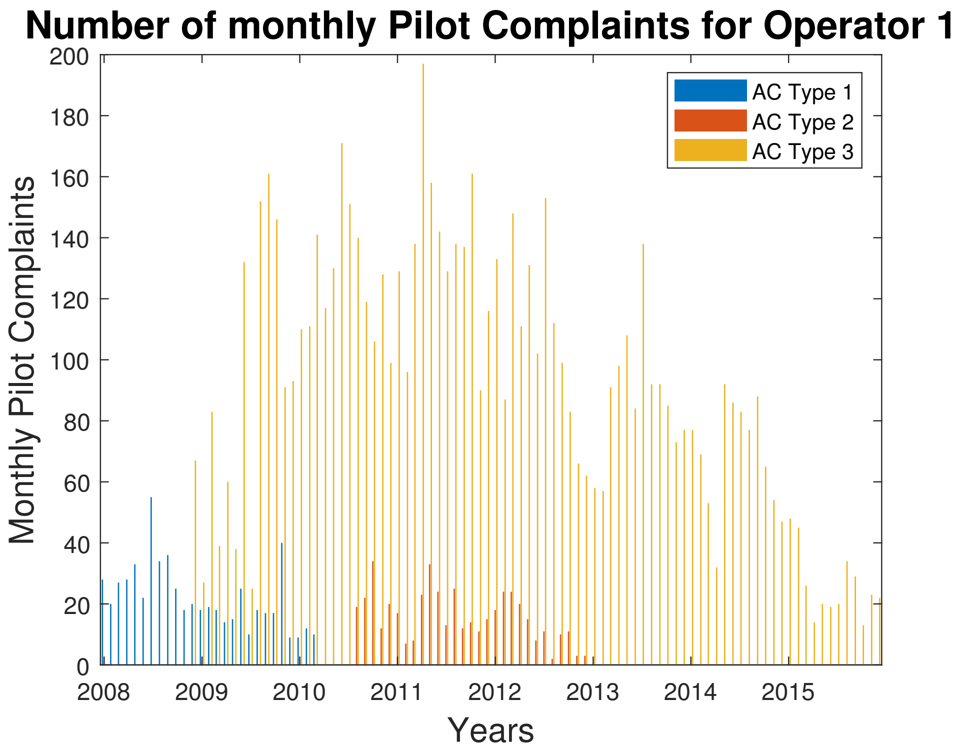

Figure 5.5: Monthly number of Pilot Complaints for Operator 1

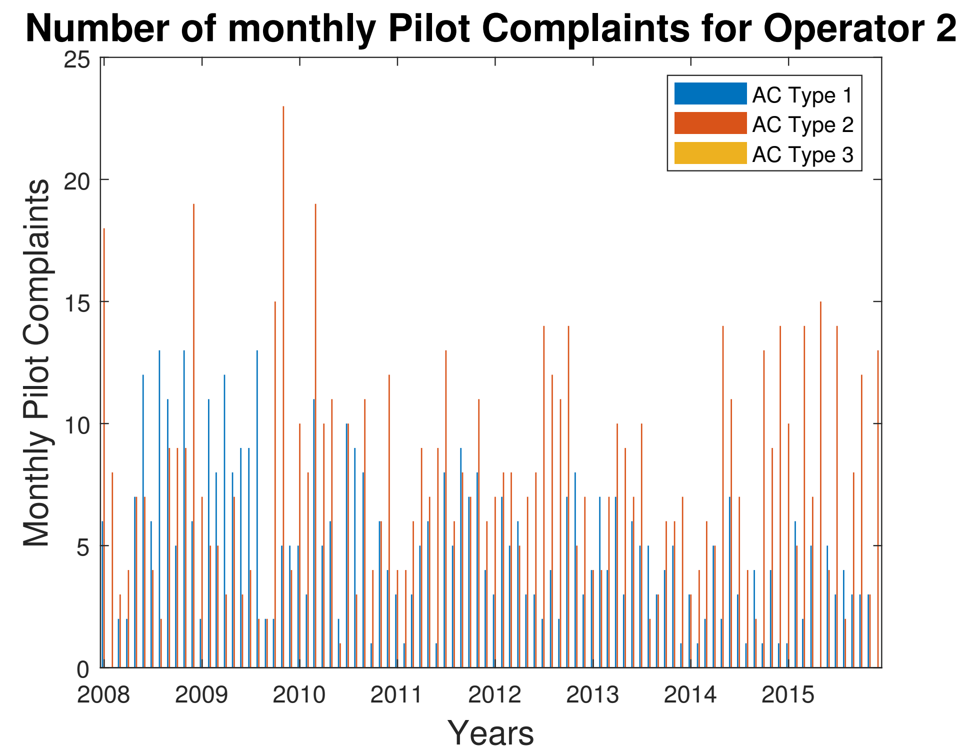

Figure 5.6: Monthly number of Pilot Complaints for Operator 2

Aircraft Type 2, while the utilisation of Aircraft Type 2 slightly decreased. Secondly, it can be seen that in general, the Aircraft Landings data patterns for Operator 2 are very smooth compared to the other operators.

Finally, Figure 5.10 shows that the number of Aircraft Landings have been rather sporadic for Operator 3, with values ranging between approximately 400 and 1200 landings. There is also a general trend of

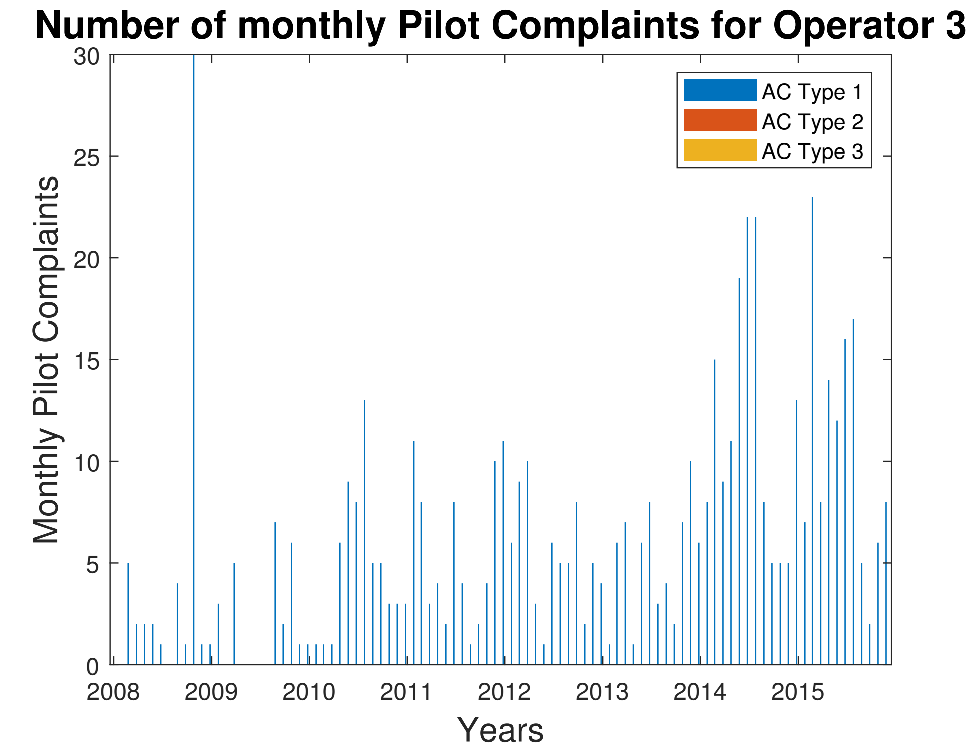

Figure 5.7: Monthly number of Pilot Complaints for Operator 3

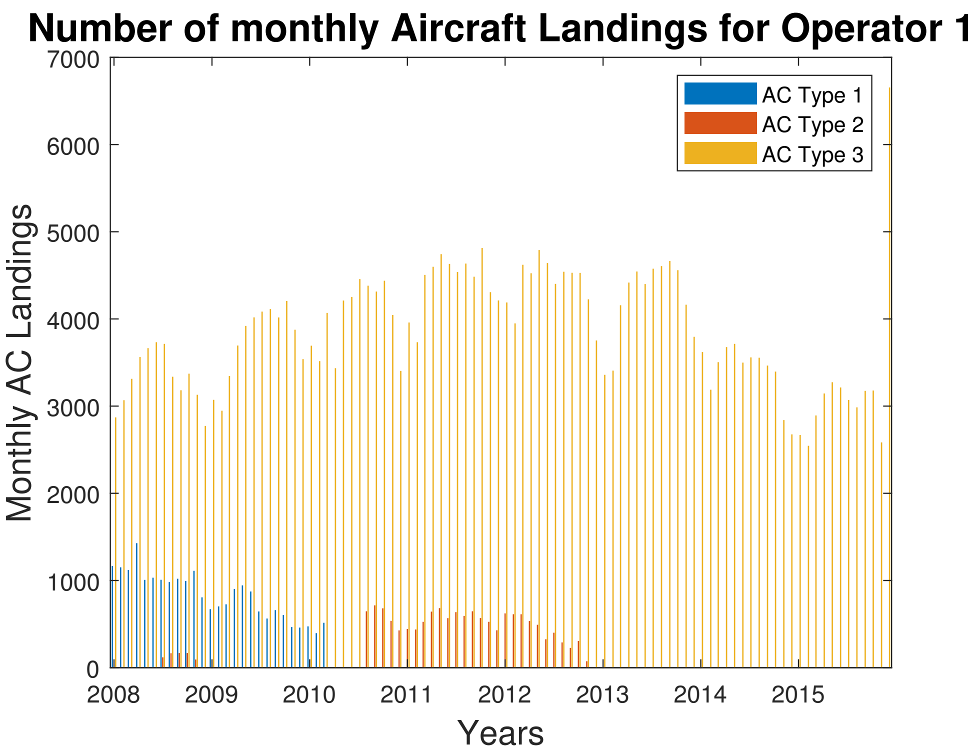

Figure 5.8: Monthly number of Landings for Operator 1

decreasing utilisation by this operator. With the data patterns of the direct key factors now generated and visualised for each operator and aircraft type, the next section will detail the performed statistical analysis to find any correlation between the causal factors and the component removal data.

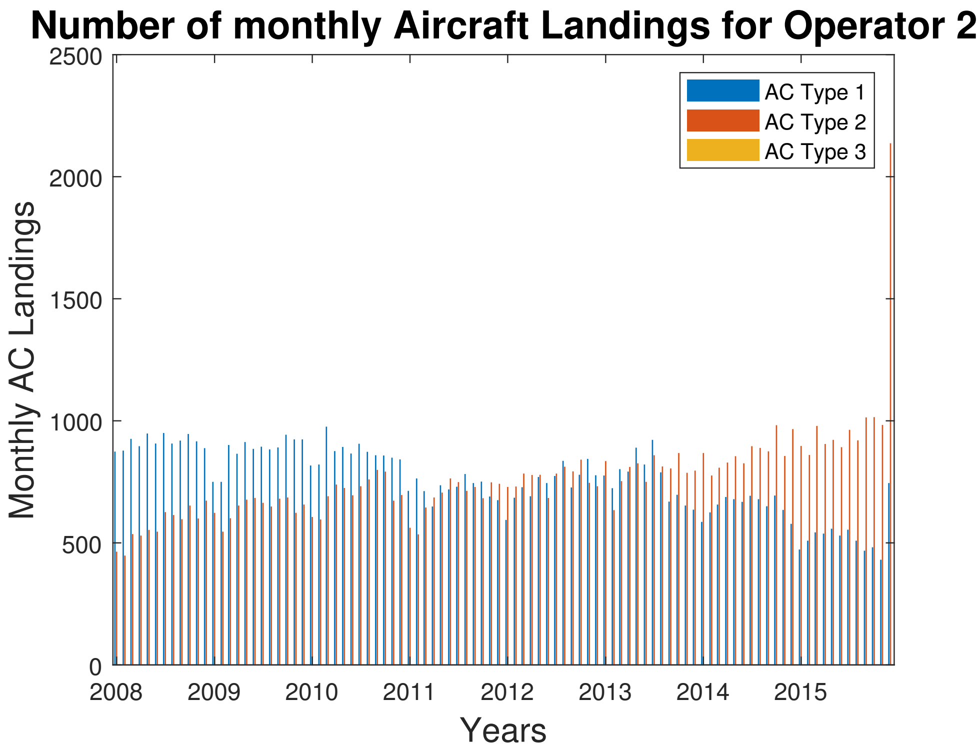

Figure 5.9: Monthly number of Landings for Operator 2

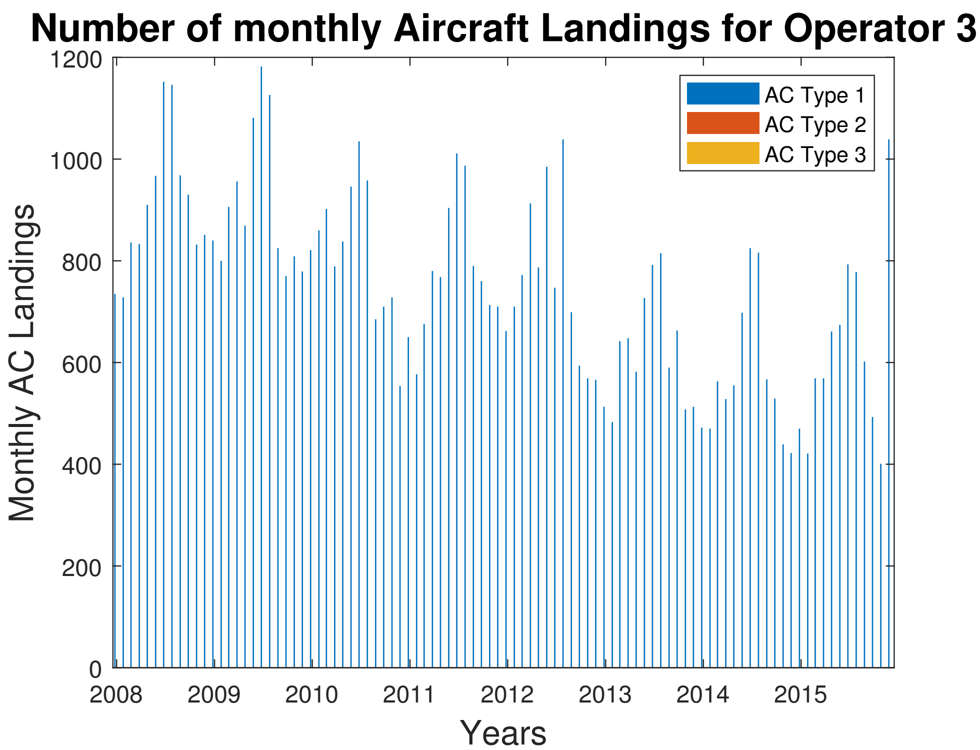

Figure 5.10: Monthly number of Landings for Operator 3

5.4 Statistical analysis of correlation between direct factors and component removals

The available MRO data is now split into subsets which are suitable for the statistical analysis. The statistical analysis will show which of the direct and indirect factors have a statistically significant impact on the component removal demand patterns. Specifically, this section will detail if there are statistical correlations between the removal of Wheels & Brakes components and the discussed direct factors, for each operator and aircraft type. Once those values are known, the baseline forecasting methods can be initiated, after which the causal factors can be incorporated while forecasting. This process will be detailed in Chapter 6.

5.4.1 Correlation between Pilot Complaints and Component Removals

The correlation between the Pilot Complaint data sets and W&B Component Removal data sets has been determined by finding the Pearson's correlation coefficients. The coefficients resulting from this analysis represent any correlation between PIC and CR with a 95% confidence interval, where a value of 1.0 can be interpreted as a perfect linear correlation and a value of 0.0 would mean that there is no correlation whatsoever. It is important to note that the analysis of correlation is applied to only the first 6 years of the dataset, since the last two years will be used to validate the improved forecasting method by generating forecast demand. Table 5.1 shows an overview of the obtained correlation coefficients between PIC and CR, for the three operators, per Aircraft Type. Note that since Operators 2 and 3 have not utilised all three aircraft types, the correlation coefficients for some aircraft types are not available.

Table 5.1: Overview of correlation coefficients between Pilot Complaints and Component Removal quantities

	Fokker 50	Fokker 100	Fokker 70
Operator 1	0.779	0.923	0.813
Operator 2	-0.067	0.269	n/a
Operator 3	-0.034	n/a	n/a

The results in Table 5.1 show that there is a significant correlation between the PIC and W&B CR databases for Operator 1, across all three aircraft types. The strongest correlation exists for the Fokker 100, which has a correlation coefficient of **0.923**. Unfortunately Operator 1 does not operate the Fokker 50 and Fokker 100 anymore, but for the aircraft that they still have in operation (Fokker 70) there also exists a positive significant correlation of **0.813**. The data scatter of pilot complaints versus component removals for all three aircraft types of Operator 1 are shown in Figures 5.11 through 5.13.

Regarding the databases for Operator 2, there appears to not be a correlation between PIC and CR for the F50 aircraft, but there seems to be a slight correlation of **0.269** for the F100 aircraft. Finally, there is also no positive or negative correlation between the PIC and CR databases for Operator 3, since the correlation coefficient is almost zero. These insights show that the Operator 1 data sets can offer the most potential to incorporate the effects of the key causal factors when forecasting W&B component demand sizes.

5.4.2 Correlation between Aircraft Landings and Component Removals

The correlation between the LND data base and CR data base is also obtained by determining Pearson's correlation coefficient, in order to estimate the potential impact of Aircraft Landings on the monthly demand size of Wheels & Brakes component removals. Again, the results indicate if there is a correlation between LND and CR with a 95% confidence interval, with 1.0 being a perfect correlation and 0.0 being no correlation at all. Again, it should be noted that the statistical analysis is only applied to the first 6 years of the dataset, to ensure efficient validation of the last two years of demand data. Table 5.2 shows the overview of correlation coefficients resulting from the statistical analysis, for each operator and aircraft type. Again, three coefficients are missing due to the fact that their operators have not utilised all three aircraft types.

Looking at Table 5.2, it can be seen that again the data sets of Operator 1 show the most significant

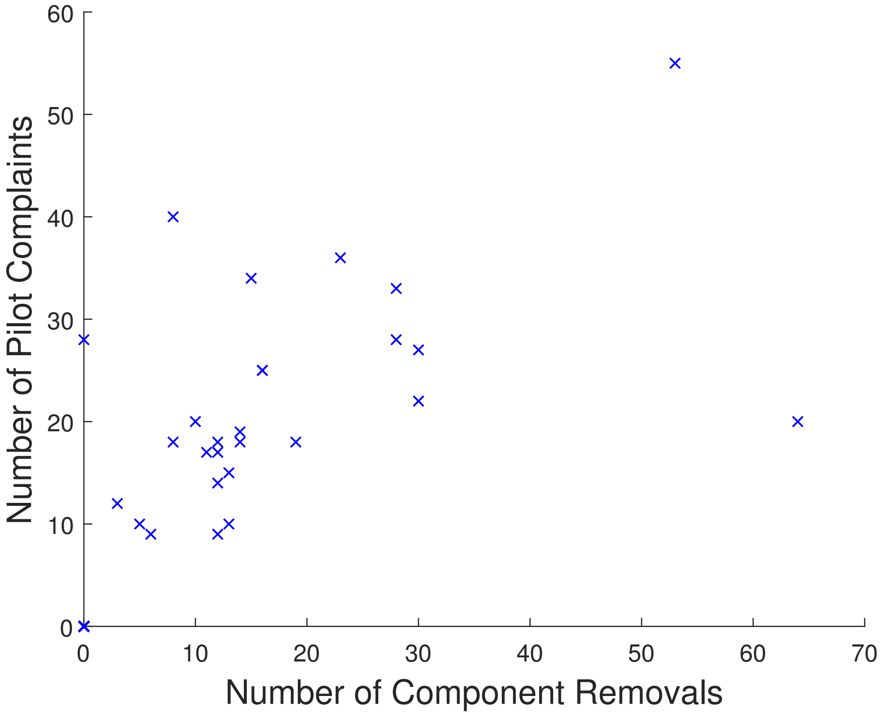

Figure 5.11: Scatter plot of PIC vs. CR for F50 of Operator 1

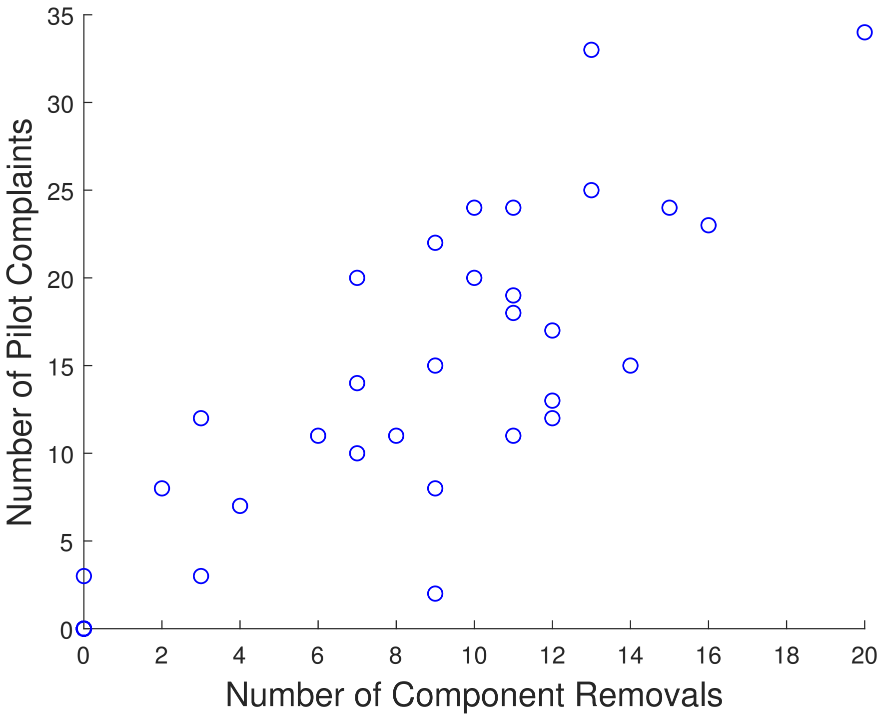

Figure 5.12: Scatter plot of PIC vs. CR for F100 of Operator 1

correlation coefficients, with values of **0.790**, **0.905** and **0.772** for the F50, F100 and F70 respectively. For Operator 2, no apparent correlation can be seen for any of their aircraft. Interestingly, for Operator 3 there appears to be a correlation of **0.463** between their F50 landings and component removals, while there was no apparent correlation between pilot complaints and component removals. Again, to visualise the data behind the most relevant correlation coefficients, Figures 5.14 through 5.16 show the Aircraft Landings versus Component Removals scatter plots of all three aircraft types of Operator 1.

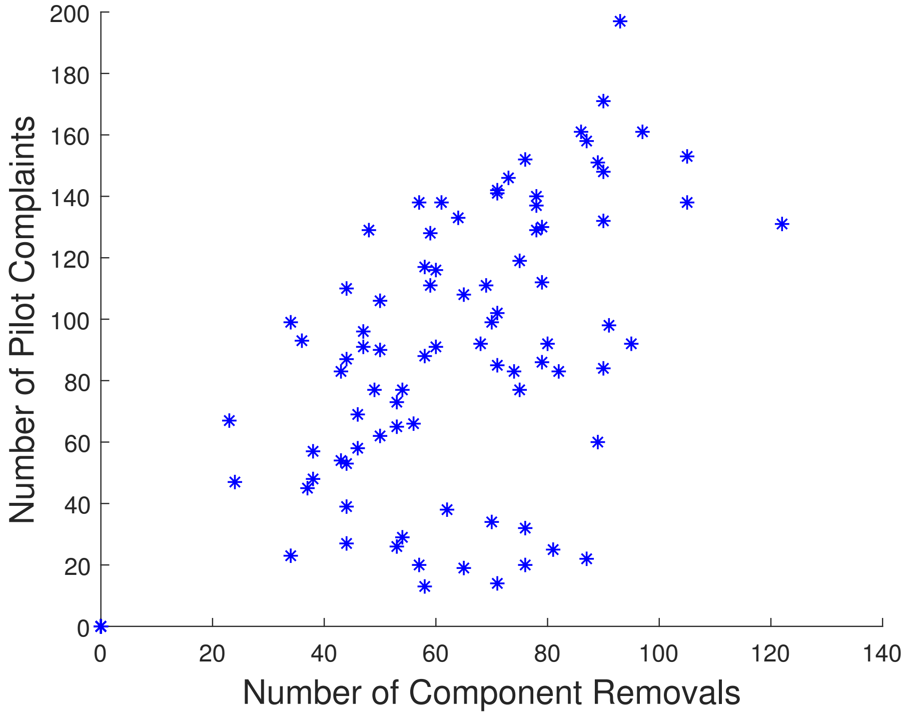

Figure 5.13: Scatter plot of PIC vs. CR for F70 of Operator 1

Table 5.2: Overview of correlation coefficients between Aircraft Landings and Component Removal quantities

	Fokker 50	Fokker 100	Fokker 70
Operator 1	0.790	0.905	0.772
Operator 2	-0.124	0.100	n/a
Operator 3	0.463	n/a	n/a

These results indicate that the Fokker 70 data sets of Operator 1 can have the most promising and relevant improvements when forecasting demand sizes, which is why the baseline forecasting methods will be applied to this specific data set initially. Furthermore, the W&B component removals for the Fokker 70 are for more than 75% accounted for by Operator 1, thus further emphasizing the significance of this specific data set and the potential forecasting improvements that can be achieved.

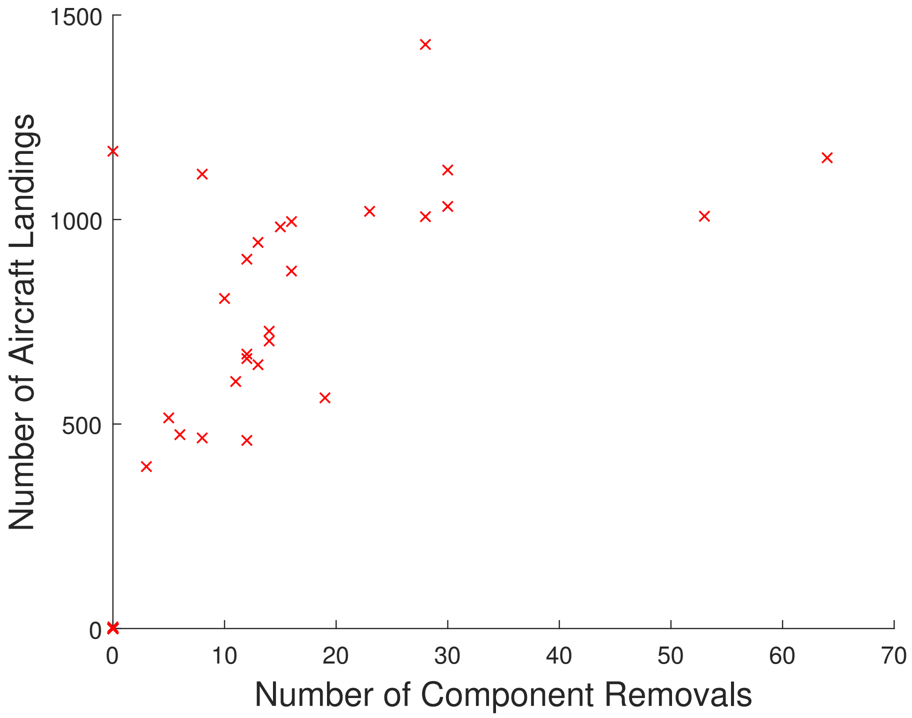

Figure 5.14: Scatter plot of LND vs. CR for F50 of Operator 1

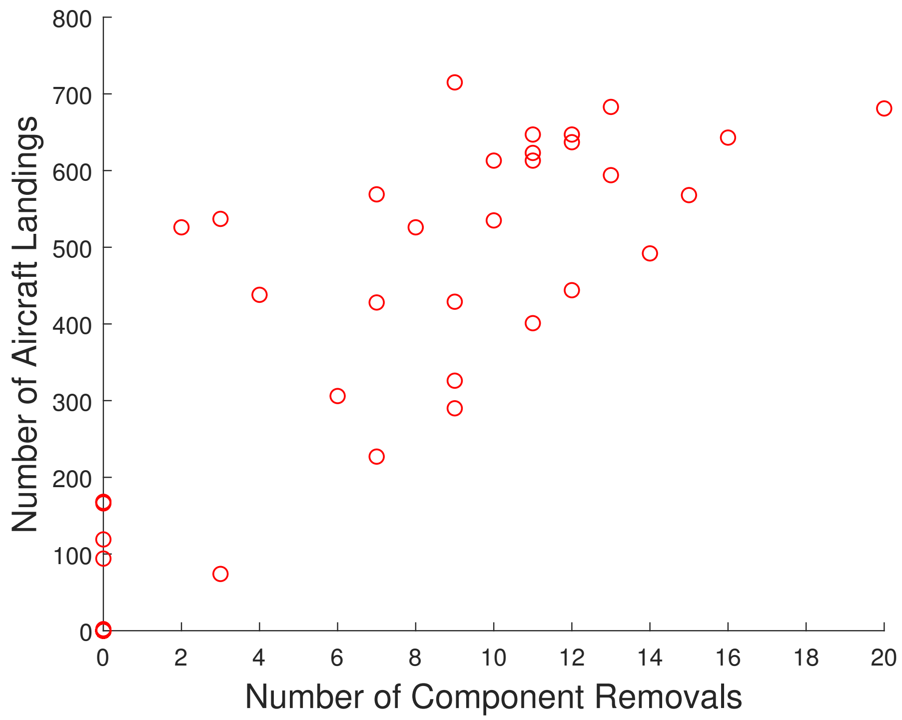

Figure 5.15: Scatter plot of LND vs. CR for F100 of Operator 1

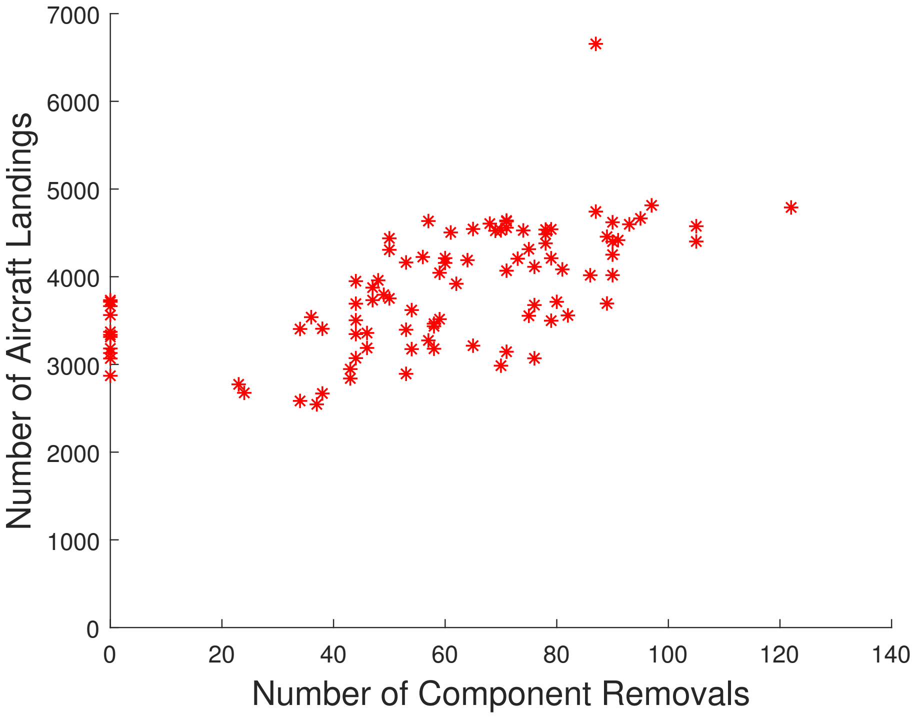

Figure 5.16: Scatter plot of LND vs. CR for F70 of Operator 1

CHAPTER 6

Application of forecasting methods

The relevant correlation coefficients are found in the statistical analysis performed in the previous chapter, so now the forecasting methods described in Chapter 3 can be applied. This chapter will therefore mainly focus on the application of the baseline and adjusted forecasting methods, starting with the application of the baseline methods described in Section 6.1. Following from this, Section 6.2 will describe how the effect of seasonality can be factored in while forecasting, after which Section 6.3 will describe how the adjusted forecasting methods are applied to the Wheels & Brakes data sets.

6.1 Applying baseline forecasting methods

The two baseline forecasting methods described in Section 3.2 are now used to forecast the three Wheels & Brakes components Main Wheel Tire, Nose Wheel Tire and Main Brake Wheel Unit, specifically for the F70 aircraft of Operator 1. In the following subsections, the actual and forecast demand data are visualised, after which the forecast errors are displayed. Finally, in order to assess the performance of both forecasting methods, the overall forecast error is determined with the RMSE and MAPE metrics. It should be noted that both forecasting methods will forecast demand between 2009 and 2015, using the actual demand in the same time frame as a comparison to determine the forecast error.

6.1.1 Forecasting Main Wheel Tire components

The Main Wheel Tire components are the most significant segment within the Wheels & Brakes category. The actual demand of the Main Wheel Tire component removals for the Operator 1 F70 are depicted in Figure 6.1. In the same figure, the MA and SES forecast demand patterns are presented with a blue dashed line and a red line with dots, respectively. Following from that, Figures 6.2 and 6.3 depict the monthly forecast error produced by the MA and SES methods when forecasting the Main Wheel Tire demand size.

From Figure 6.1 it becomes clear that both the forecasting methods follow the general trend of ups and downs in forecasting the demand data. Both methods fail to capture the truly outlying months, when there is suddenly a month with a very high or very low demand size. While the actual demand data size varies between 10 and almost 70 units, the forecast demand in general varies between 20 and 50

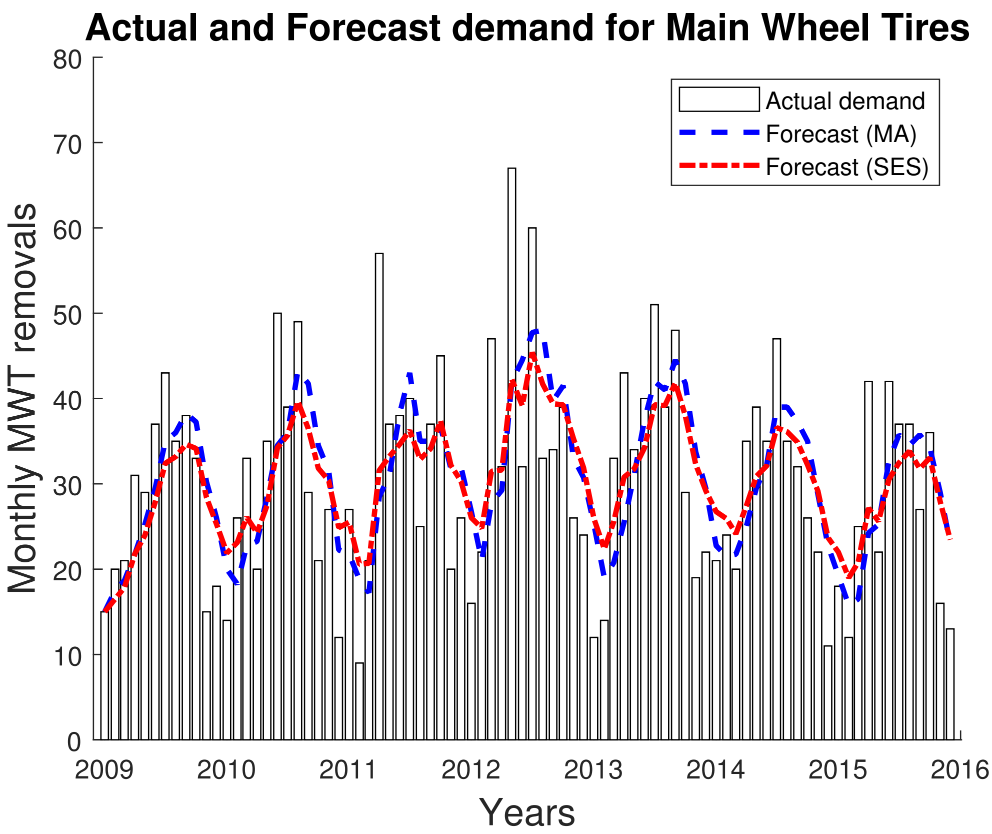

Figure 6.1: Actual and forecast monthly demand volumes for Main Wheel Tire components

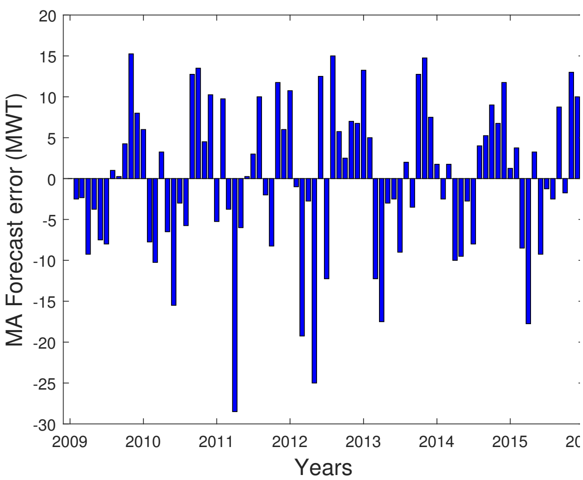

Figure 6.2: MA forecast error for Main Wheel Tires

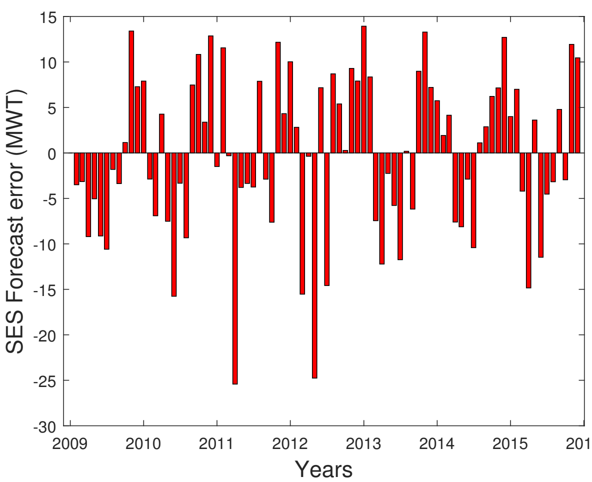

Figure 6.3: SES forecast error for Main Wheel Tires

units. This indicates that the forecast methods are less volatile than the actual demand pattern.

When looking at the forecast performance of both methods, as depicted in Figures 6.2 and 6.3, both methods show a similar error pattern. Many months are underestimated and overestimated, with some months being underestimated by more than 25 units, which is quite significant. To assess which method performed best, the error values of the baseline MA method and SES method have been determined to be as follows:

- RMSE for Moving Averages method equals: **9.34**
- RMSE for Single Exponential Smoothing method equals: **8.71**

- MAPE for Moving Averages method equals: **29.3%**
- MAPE for Single Exponential Smoothing method equals: **28.8%**

These RMSE and MAPE values indicate that the general average magnitude of error for the MA method is higher compared to the SES method. Based on these facts, it can be stated that for this specific data set, the SES method yields a more accurate forecast demand.

6.1.2 Forecasting Nose Wheel Tire components

The Nose Wheel Tire components are also forecast with the MA method and SES method. The results of these forecasts is presented in Figure 6.4. It can already be seen that the overall demand size of NWT components is less than MWT components, but the variability is still very high. Again, both forecast methods fail to accurately predict the months with extremely outlying demand sizes. Both of the methods also follow the same general trend, while the MA method is slightly more volatile compared to the SES method.

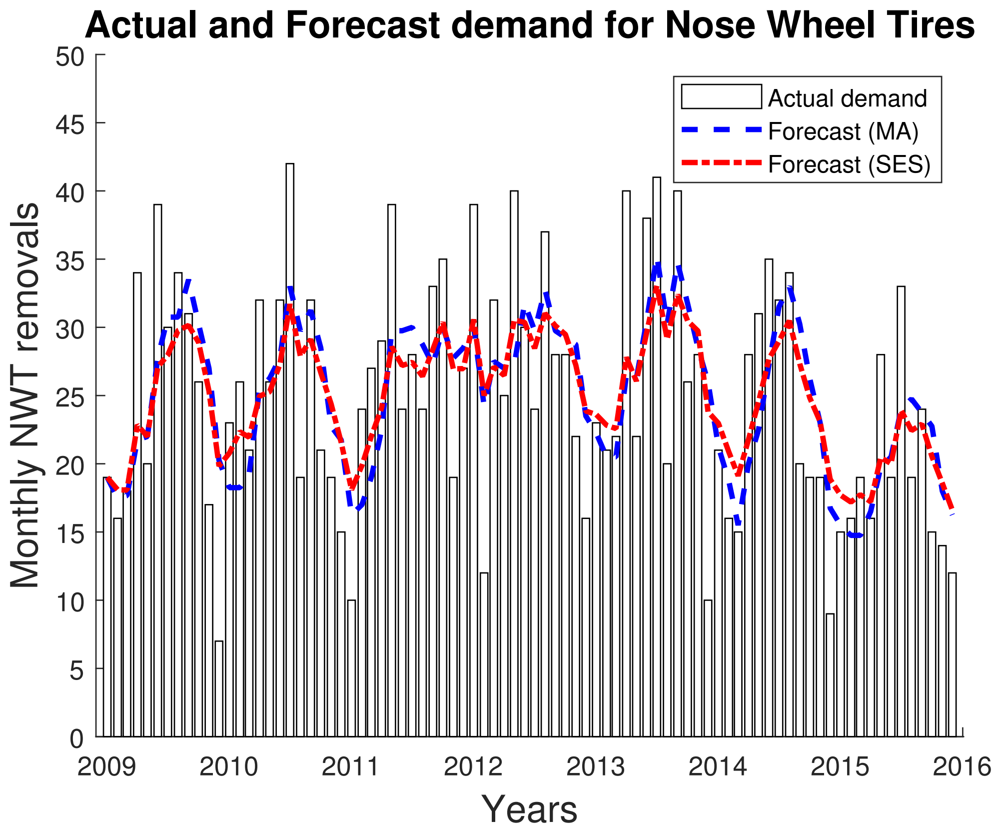

Figure 6.4: Actual and forecast monthly demand volumes for Nose Wheel Tire components

Figures 6.5 and 6.6 show the monthly forecast error graphs for both of the forecasting methods. The graphs represent the fact that both methods are rather inaccurate in general, as many months are either underestimated or overestimated. In the extreme cases, both methods often underestimate or overestimate the demand size with more than 10 units, which is a fairly significant amount. To determine the forecasting accuracy of MA and SES for the Nose Wheel Tire data subset, the RMSE values have been determined to be:

- RMSE for Moving Averages method equals: **6.51**
- RMSE for Single Exponential Smoothing method equals: **6.07**
- MAPE for Moving Averages method equals: **25.6%**

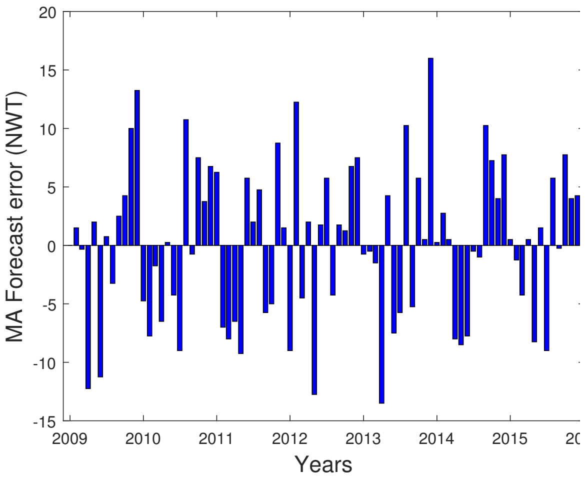

Figure 6.5: MA forecast error for Nose Wheel Tires

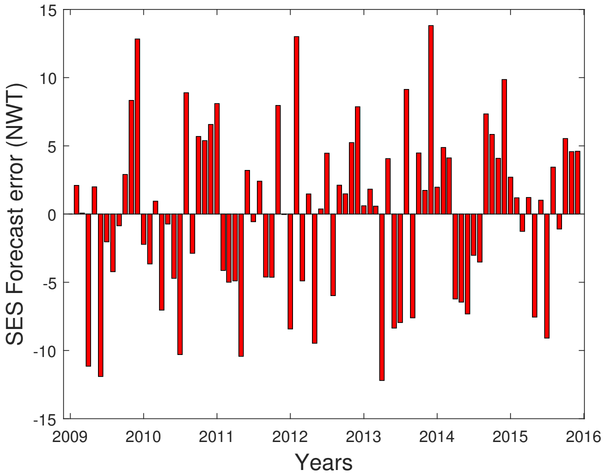

Figure 6.6: SES forecast error for Nose Wheel Tires

- MAPE for Single Exponential Smoothing method equals: **24.5%**

Looking at both the RMSE and MAPE values, again it can be said that the SES method slightly outperformed the MA method. The obtained RMSE values are also lower than the RMSE values of the Main Wheel Tire data set, but this is due to the fact that the magnitude of overall demand sizes (and therefore the magnitude of forecast errors) is lower for Nose Wheel Tires than for Main Wheel Tires. When looking at the MAPE values, it can be seen that they are in roughly the same range as the MAPE values for MWT components, which further demonstrates the fact that the MAPE metric is less susceptible to the magnitude of overall demand sizes, in comparison to the RMSE metric.

6.1.3 Forecasting Main Wheel Brake Unit components

Finally, the Main Wheel Brake Units are a significant portion of the Wheels & Brakes data set, which is why the demand volumes of these specific components are also forecast with the MA and SES method. Figure 6.7 shows the actual and forecast demand volumes for Main Wheel Brake Unit components between 2009 and 2015, and Figures 6.8 and 6.9 show the corresponding forecast errors produced by the MA and SES method, respectively.

Figure 6.7 shows that the overall demand sizes for MWBU components are lower compared to MWT and NWT components. This leads to the forecast values for MWBU components being less volatile, in comparison to the forecast values for MWT and NWT components. As can be seen from Figure 6.7, the forecast values range between 4 approximately 10, while the actual demand size varies between 1 and 17 units. Once more, the MA method seems to be more sensitive to extremes compared to the SES method.

Looking at Figures 6.8 and 6.9, the performance of the MA method and SES method can be determined. Again, both methods have underestimated and overestimated the demand sizes in most of the months, but the overall forecast error is less compared to the MWT and NWT component data sets. Finally, to compare the forecasting accuracy of the two forecasting methods for the Main Wheel Brake Unit data set, the RMSE and MAPE values have been determined as follows:

- RMSE for Moving Averages method equals: **2.68**
- RMSE for Single Exponential Smoothing method equals: **2.28**
- MAPE for Moving Averages method equals: **49.8%**

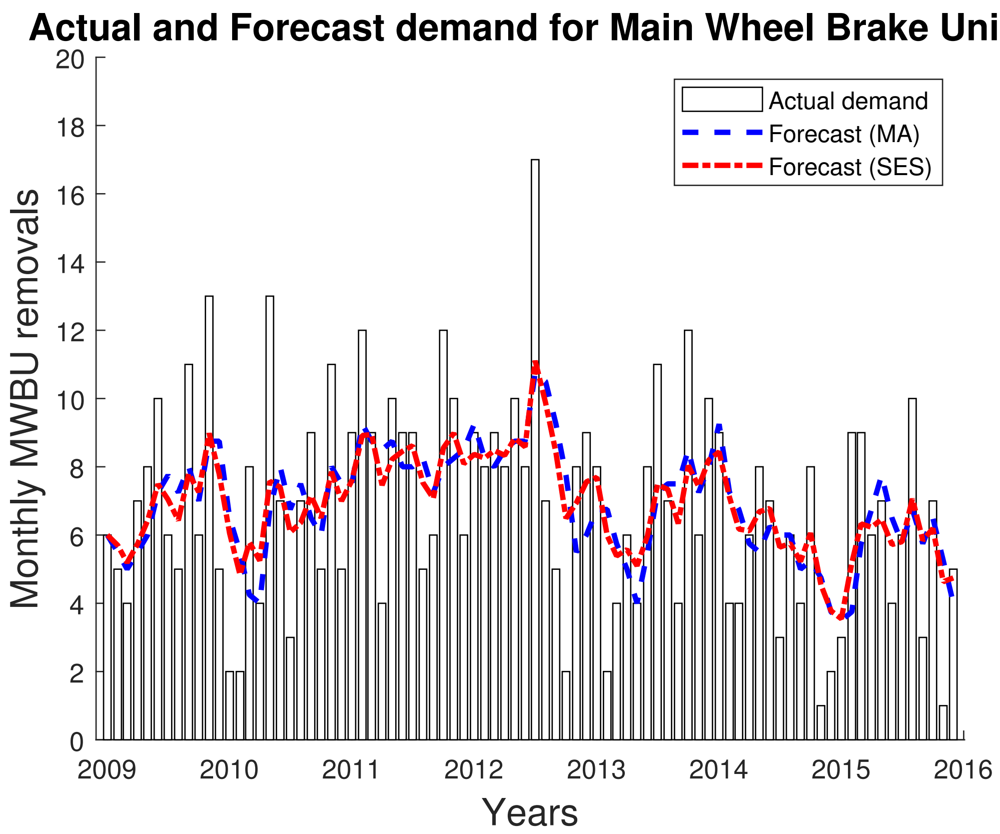

Figure 6.7: Actual and forecast monthly demand volumes for Main Wheel Brake Unit components

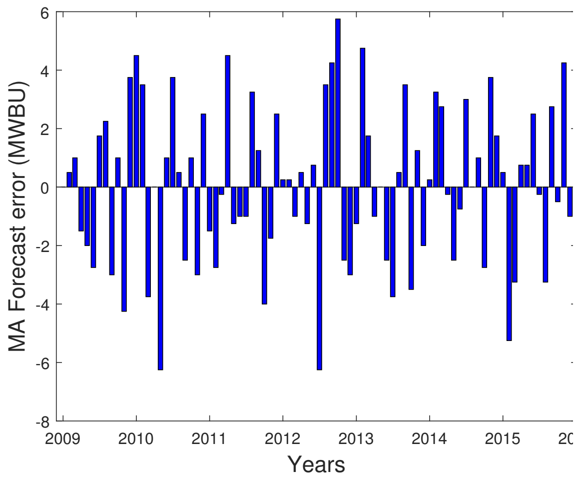

Figure 6.8: MA forecast error for MWBU parts

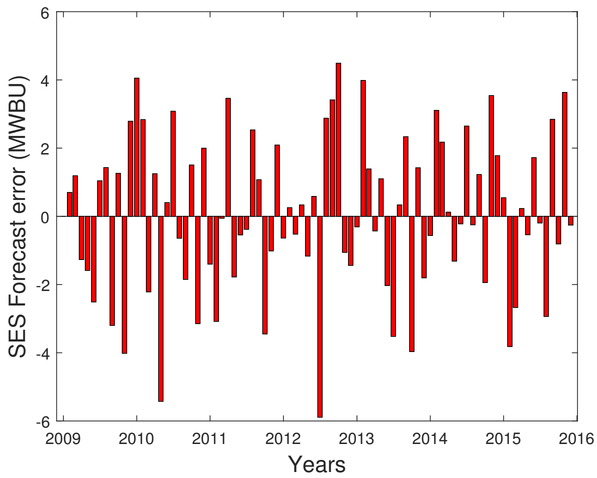

Figure 6.9: SES forecast error for MWBU parts

- MAPE for Single Exponential Smoothing method equals: **42.6%**

Interestingly, the RMSE values for the MA method and SES method are very similar, with SES outperforming MA with a marginal RMSE error magnitude of only 0.4. While the RMSE results appear to be the smallest for the MWBU component type, it should be taken into account that the monthly demand size for MWBU components is also significantly less than the demand sizes for MWT and NWT components. In this case, the MAPE values provide a more complete sense of the forecasting accuracy, as they are significantly higher than the MAPE values determined for MWT and NWT components. This indicates that the baseline forecast methods have shown the least forecasting accuracy when forecasting Main Wheel Brake Unit components.

6.2 Adjusting forecasting with seasonality

Looking at the actual demand data patterns as shown in the previous section, a seasonal trend can be recognised in the data patterns of Main Wheel Tires and Nose Wheel Tires. This trend can be explained by the fact that the quality of tires is directly related to the operating environment. Thus, in hot summer months, the necessity to replace these specific components is increased under warm environmental conditions. Similarly, one can expect less tire removals during the cold winter months.

The seasonal trends can also be taken into account when forecasting, in order to improve the accuracy of the forecast. Figures 6.10 and 6.11 show the demand patterns with a sinusoidal function fitted to the data, to highlight the seasonal trend.

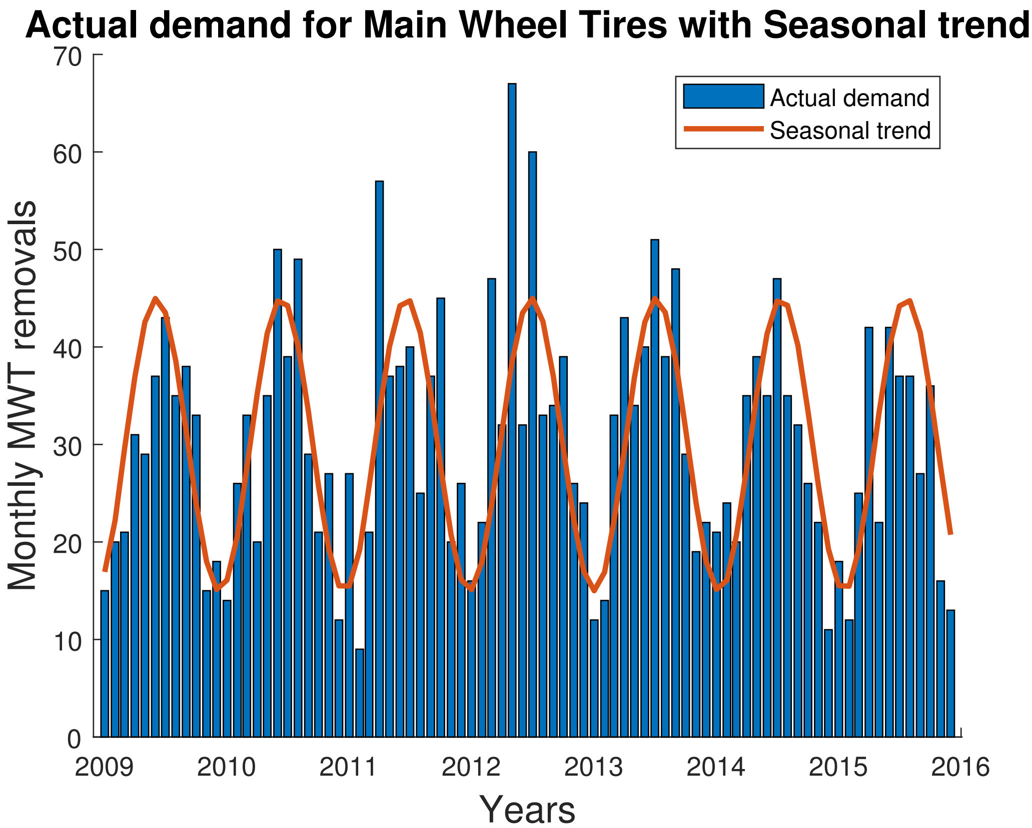

Figure 6.10: Monthly demand volumes for MWT components, with Seasonal trend

One method to take this seasonal effect into account is to tune the forecast values accordingly. With this method, the forecast demand in a month with high seasonality will be tuned upwards, while the forecast demand will be tuned downwards if the forecast is for a month with low seasonality. Equation 6.1 shows the mathematical relation that was implemented to determine if adjusting for seasonality will reduce forecast errors. In this equation, the parameter s determines how strongly the forecast should follow the seasonal trend.

$$F^* = F' \cdot s \cdot [Y1/Y0] \tag{6.1}$$

Where;

- F' is the MA or SES demand forecast value
- s is the seasonal correction factor
- Y_1 is the seasonal trend in the current month, given by the output of the fitted sinusoidal function

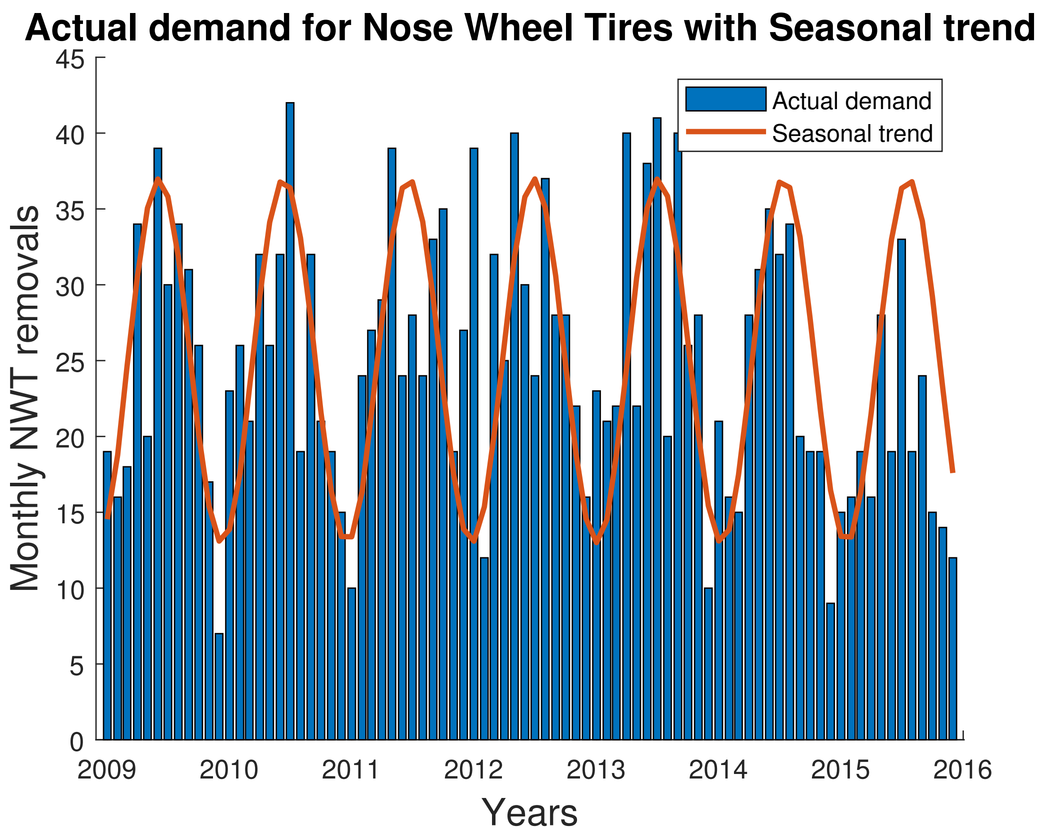

Figure 6.11: Monthly demand volumes for NWT components, with Seasonal trend

- Y_0 is the mean seasonal equilibrium in the data set

The effects of seasonality have been incorporated for both the MA and SES forecasting methods. Figures 6.12 through 6.15 show the resulting graphs depicting the actual demand, the baseline MA or SES forecast, and the forecasts adjusted for seasonality. For this analysis, the seasonal correction factor s was assumed to be equal to **0.8**, which yielded the least forecasting errors.

Looking at Figures 6.12 and 6.13, at first glance it can already be seen that the forecast including Seasonality more accurately follows the actual demand pattern, compared to the baseline MA and SES methods. To quantitatively determine the forecasting improvements achieved by adjusting for Seasonality, the forecasting error has been computed for each forecasting method, using a MAPE error metric. Table 6.1 provides an overview of the MAPE values for the baseline MA and SES methods, the forecasting methods adjusted for Seasonality, for MWT and NWT components.

Table 6.1: Overview of MAPE values for baseline and adjusted forecasting methods

	MAPE	
	Main Wheel tires	Nose Wheel tires
MA (baseline)	29.3%	25.6%
MA (adj. for Seasonality)	28.5%	30.3%
SES (baseline)	28.8%	24.5%
SES (adj. for Seasonality)	25.5%	27.4%

Based on the MAPE values presented in Table 6.1 it can be concluded that tuning the MA and SES forecast with Seasonality resulted in a better forecasting performance for Main Wheel Tires only. Adjusting for Seasonality when forecasting Main Wheel Tires has reduced the MAPE from 29.3% to 28.5% and from 28.8% to 25.5%, for the MA and SES methods respectively. Unfortunately, accounting for Seasonality when forecasting Nose Wheel Tire components leads to increased forecasting errors.

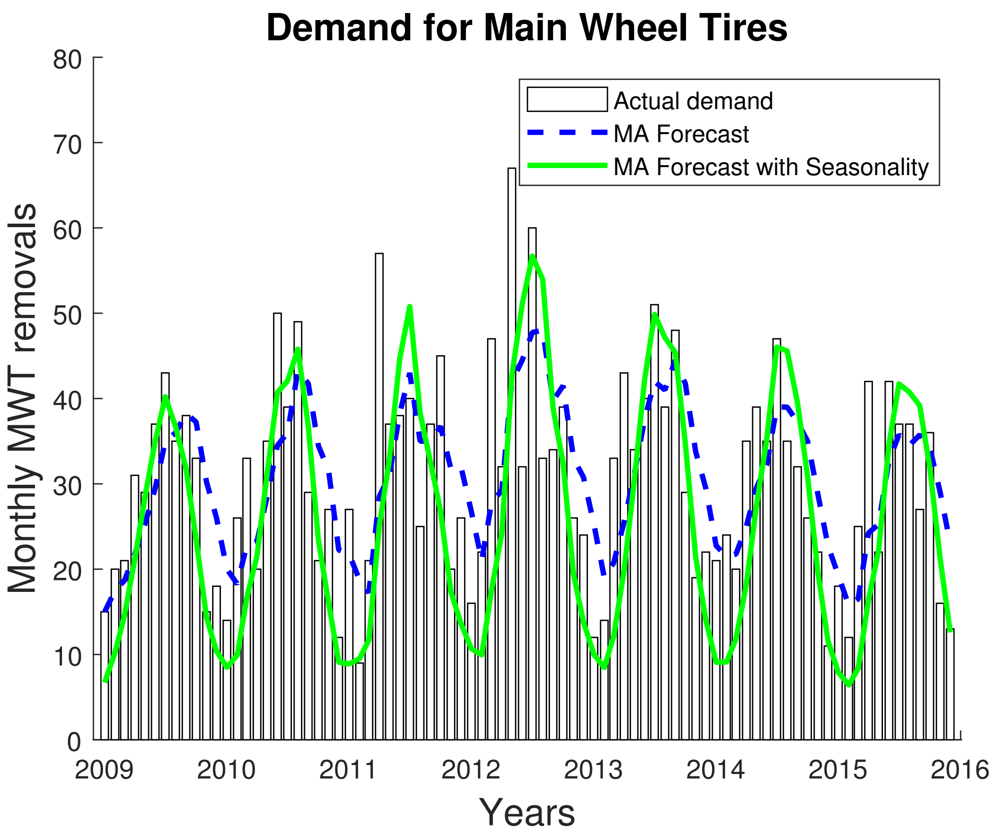

Figure 6.12: Actual and MA forecast demand volumes for MWT components, including Seasonality

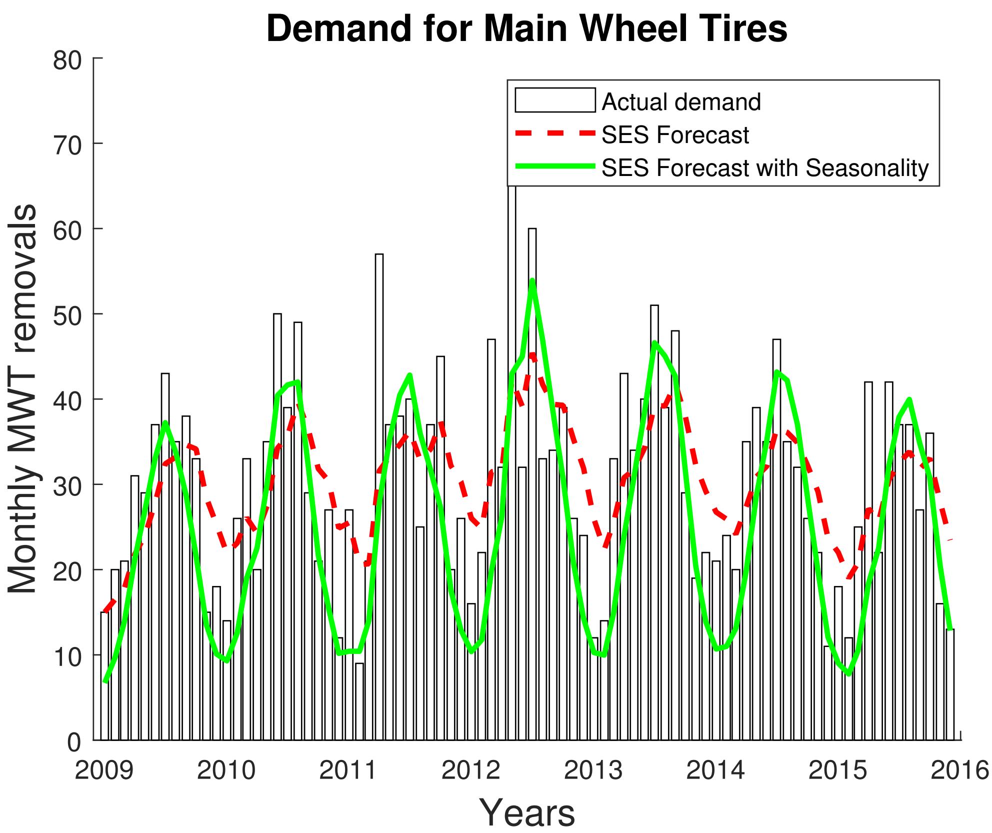

Figure 6.13: Actual and SES forecast demand volumes for MWT components, including Seasonality

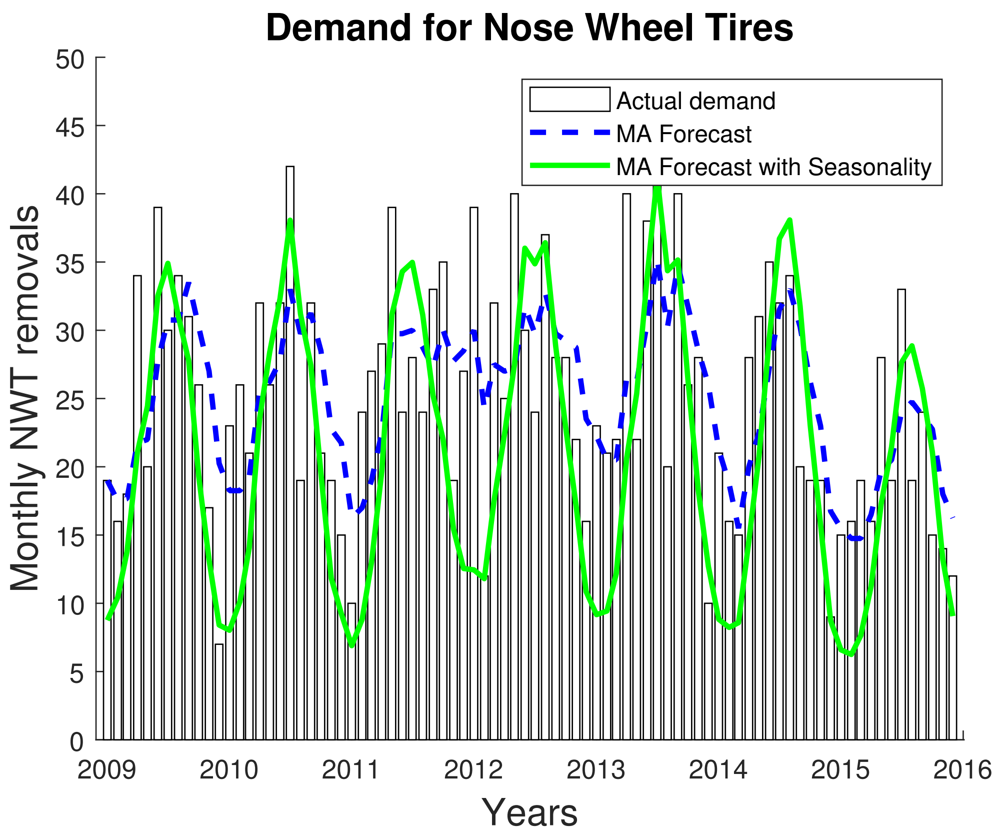

Figure 6.14: Actual and MA forecast demand volumes for NWT components, including Seasonality

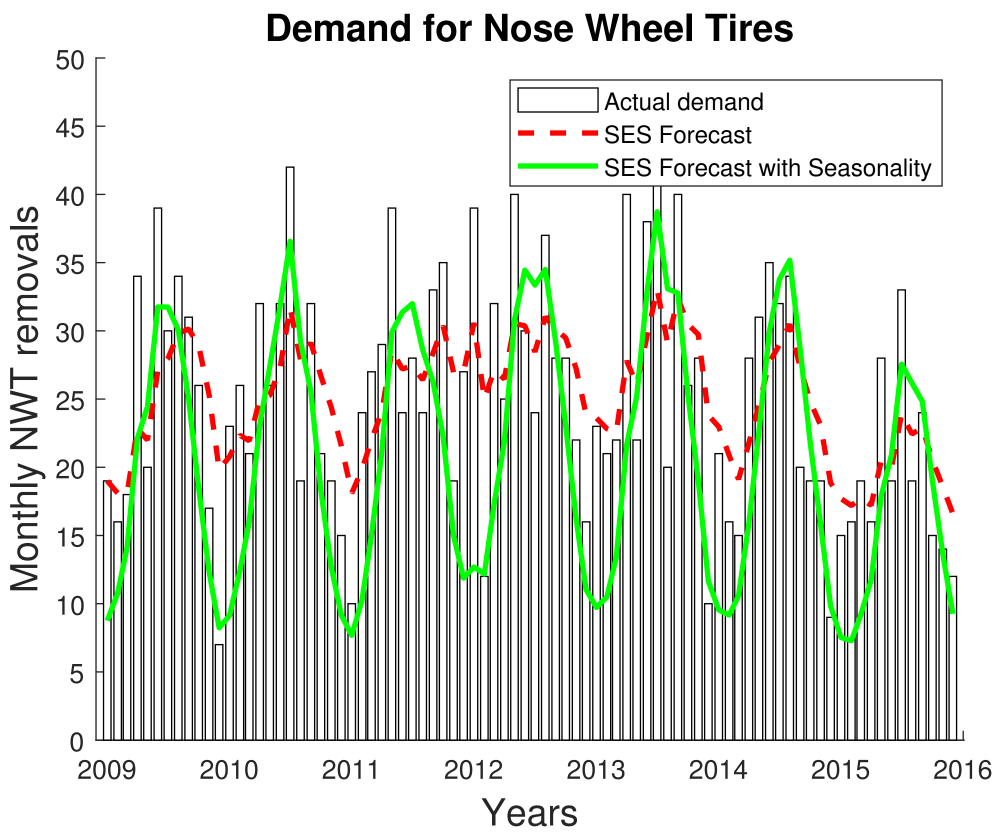

Figure 6.15: Actual and SES forecast demand volumes for NWT components, including Seasonality

These findings indicate that the sinusoidal function fitted to the Main Wheel Tire demand data is a better representation of the seasonality, compared to the seasonality function fitted to the Nose Wheel Tire demand data. Therefore it can be recommended to include the Seasonality trend for Main Wheel Tires, and to not take Seasonality into account when forecasting Nose Wheel Tires.

6.3 Applying adjusted forecasting method

This section will deal with the application of the altered forecasting method as described in Section 3.3 in the previous chapter. This is the method that will help answer the main research questions by determining the forecasting improvements that can be realised when additional factors are incorporated. These additional factors are the number of Pilot Complaints and Aircraft Landings in a recent time frame.

As explained in Section 3.3, the time-series forecast that results from historic data will be tuned depending on the statistical correlation between the two additional factors and the historic component removal data. The results of the baseline and adjusted Moving Averages forecasting methods are presented in Figures 6.16 through 6.18. In a similar fashion, the baseline SES forecasting method has been altered, the results of which are found in Figures B.1 through B.3 in Appendix B.

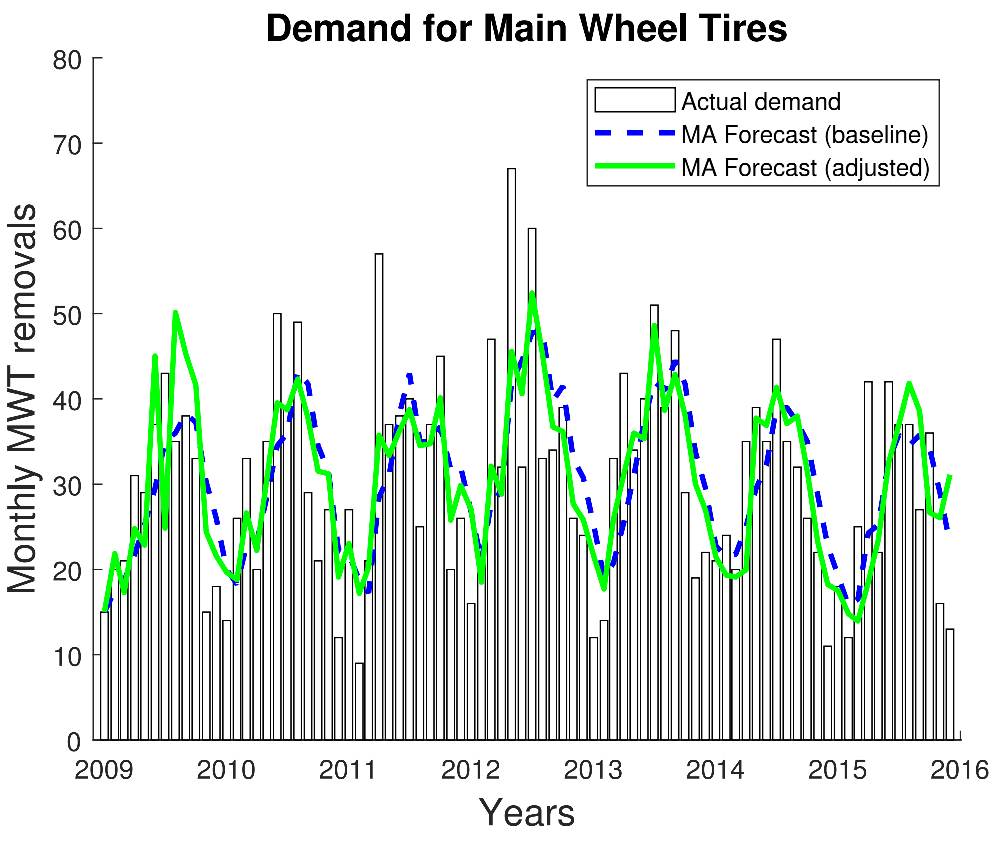

Figure 6.16: Actual and MA Forecast demand for Main Wheel Tires

Figures 6.16 through 6.18 show that the adjusted forecasting method in general roughly follows the same trend as the baseline forecasting method. The differences between the two methods are less extreme compared to the difference between the baseline forecasting method and the forecasting method adjusted for seasonality, as depicted in Figure 6.12 for instance.

However, by just looking at the forecasting results it is hard to determine whether the adjusted MA method has outperformed the baseline MA method for the three components. That is why the MAPE of each forecast is computed and presented in Table 6.2.

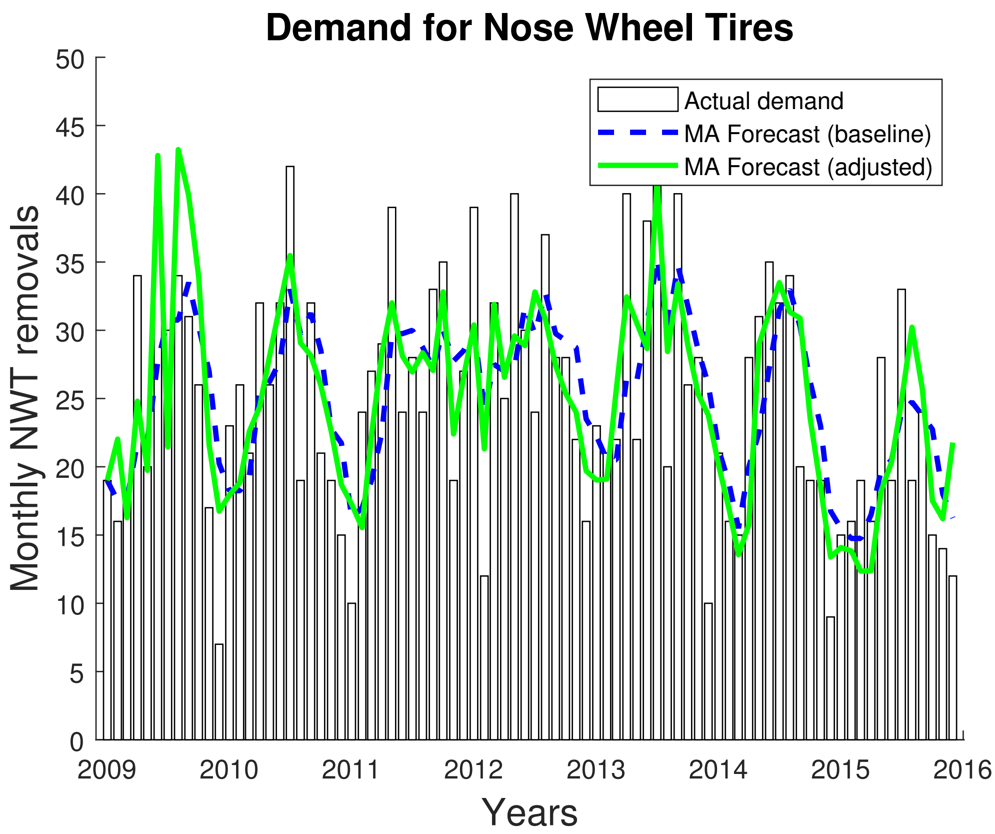

Figure 6.17: Actual and MA Forecast demand for Nose Wheel Tires

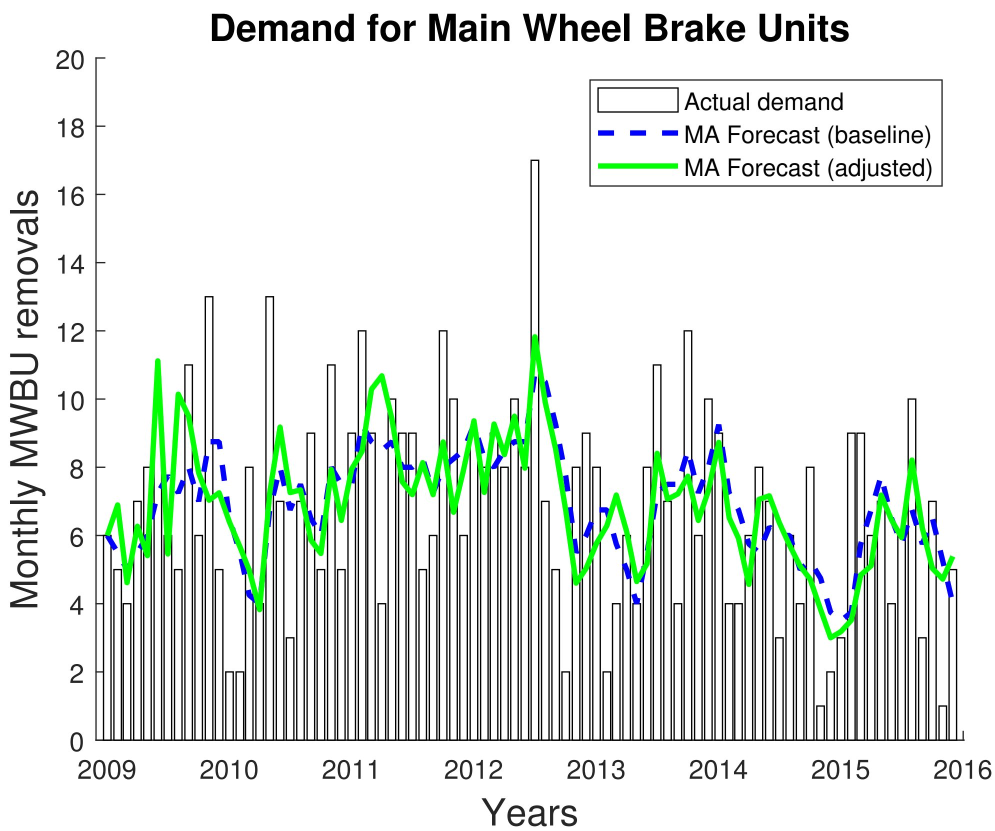

Figure 6.18: Actual and MA Forecast demand for Main Wheel Brake Units

Table 6.2: Overview of MAPE values for baseline forecasting methods and forecasting methods adjusted for Key factors

	MAPE		
	Main Wheel Tires	Nose Wheel Tires	Main Wheel Brake Units
MA (baseline)	29.3%	25.6%	49.8%
MA (adj. with Key Factors)	24.8%	23.2%	47.6%
SES (baseline)	28.8%	24.5%	42.6%
SES (adj. with Key Factors)	23.6%	21.5%	42.6%

Based on the results presented in Table 6.2, it can be observed that almost all MAPE values have decreased in magnitude after taking into account the two key factors. Only the SES forecast of MWBU components has shown no improvement in forecast accuracy, which has stayed the same. With that, it can be effectively concluded that tuning the time series forecast with the actual behaviour of the two key factors generally will reduce the forecasting error and therefore improve the forecasting accuracy. The extent of the improvement is still dependent on the component type as well, as the forecasts of Main Wheel Tires have shown a more significant reduction in MAPE, compared to Nose Wheel Tires and Main Wheel Brake Units components.

Since the impact of including Seasonality was also positive for Main Wheel Tires, it is also interesting to determine how a combined method would perform, which takes into account the two key factors as well as Seasonality. Again, this analysis is only performed for the MWT and NWT components, since there was no seasonal trend to be observed for the MWBU components. The final results of this combined forecasting approach are presented in Table 6.3.

Table 6.3: Overview of MAPE values for baseline FC methods and FC methods adjusted for Key factors and Seasonality

	MAPE	
	Main Wheel tires	Nose Wheel tires
MA (baseline)	29.3%	25.6%
MA (adj. with Key Factors)	24.8%	23.2%
MA (adj. with Key Factors and Seasonality)	23.7%	23.7%
SES (baseline)	28.8%	24.5%
SES (adj. with Key Factors)	23.6%	21.5%
SES (adj. with Key Factors and Seasonality)	**20.7%**	**20.4%**

As can be seen from the results in Table 6.3, altering the baseline forecasting methods with both the Key Factors and Seasonality leads to an even further reduction of MAPE values. Compared to the baseline SES method, the altered methodology yields a reduction of forecasting error of 8.1 and 4.1 percent point for the Main Wheel Tires and Nose Wheel Tires, respectively.

This approach has shown that the forecasting accuracy can be improved to almost 20%, thus outperforming the baseline methods significantly. Based on these results, it can therefore be concluded that taking into account Pilot Complaints, Aircraft Landings and Seasonality has a significant positive effect on the forecasting accuracy for both the MA method and the SES method when forecasting MWT and NWT components. Adjusting the MA method for MWBU components has improved the MAPE values, while no significant impact could be determined for the performance of the SES method. To ensure that the improved forecasting methods also yield favorable results with other component categories in the MRO data base, the approach will be validated in the next chapter.

CHAPTER 7

Validation of altered forecasting methods

In Chapter 6 it became clear that the incorporation of additional factors while forecasting with the MA and SES method, will lead to forecasting accuracy improvements for components in the Wheels & Brakes category. In order to ensure that the method is validated and the same results can be expected with multiple types of components, this chapter will deal with the application of the adjusted forecasting method to additional datasets of components that do not necessarily belong to the Wheels & Brakes category. Section 7.1 will describe the component categories that will be used for the validation, followed by Section 7.2 which will describe the main results of applying the adjusted methods to the validation datasets. Section 7.3 will present an evaluation of the performance of the adjusted method and finally Section 7.4 will describe the sensitivity analysis that was applied to conclude the research.

7.1 Description of validation datasets

To validate that the adjusted forecasting method will also lead to forecasting improvements for component types other than Wheels & Brakes, a total of twelve additional datasets from several other component categories have been selected for the validation phase of the analysis. These datasets are mainly selected based on the largest commonality (i.e. the percentage of how much a specific component category is represented within the total database) after Wheels % Brakes, all of which have a commonality between **1% and 5%**. The selected component categories are described in an overview in Table 7.1. The actual demand patterns themselves will be presented in Section 7.2, together with the forecast demand. Also it should be noted that the selected component categories all belong to the F70 aircraft type of operator 1.

Looking at the overview presented in Table 7.1, it can be seen that the majority of the validation datasets can be classified as erratic, as many of the CV^2-values are over the threshold value of **0.49**. Furthermore, even though the data subsets with the highest commonality are selected, the majority of the validation data still show a commonality of less than 3%. This is due to the vast number of component categories that are available in the MRO dataset, with the majority of component categories having a commonality of less than 1%. With the additional component categories of the validation datasets amounting up to approximately 30%, a more generally valid sense of the performance of the adjusted forecasting method can be determined.

Table 7.1: Overview of component demand data subsets to be used for validation

	ATA3-code	Component category	Commonality	CV^2
Component type 1	212	Air conditioning - Distribution	1.5%	0.75
Component type 2	215	Air conditioning - Cooling	1.1%	0.84
Component type 3	221	Auto flight - Autopilot	2.8%	0.54
Component type 4	231	Communications - Speech	1.3%	0.78
Component type 5	235	Communications - Audio integrating	2.8%	0.58
Component type 6	253	Equipment/furnishing - Buffet/galley	3.2%	0.34
Component type 7	291	Hydraulic power - main system	1.2%	1.30
Component type 8	323	Landing gear - Extension and Retraction	1.1%	0.84
Component type 9	334	Lights - Exterior	3.0%	0.46
Component type 10	342	Navigation - Attitude and Direction	4.4%	0.43
Component type 11	345	Navigation - Dependent Position Determining	3.1%	0.52
Component type 12	351	Oxygen - Crew	5.2%	0.19

7.2 Application of baseline and adjusted FC methods on validation datasets

With the validation datasets selected as described in Section 7.1, it is now possible to apply the baseline and adjusted forecasting methods to the validation datasets. The results of the demand forecast will be presented in this section, after which the forecast errors will be evaluated in Section 7.3. The MA, SES and adjusted forecasts of component type 1, 7 and 12 are represented in Figures 7.1 through 7.6. These three component types are selected to depict how the forecasting methods deal with the most smooth dataset (CT12), the most erratic dataset (CT7) and a dataset with a high CV-value (CT1). The remainder of the results for each component type can be found in Appendix C.

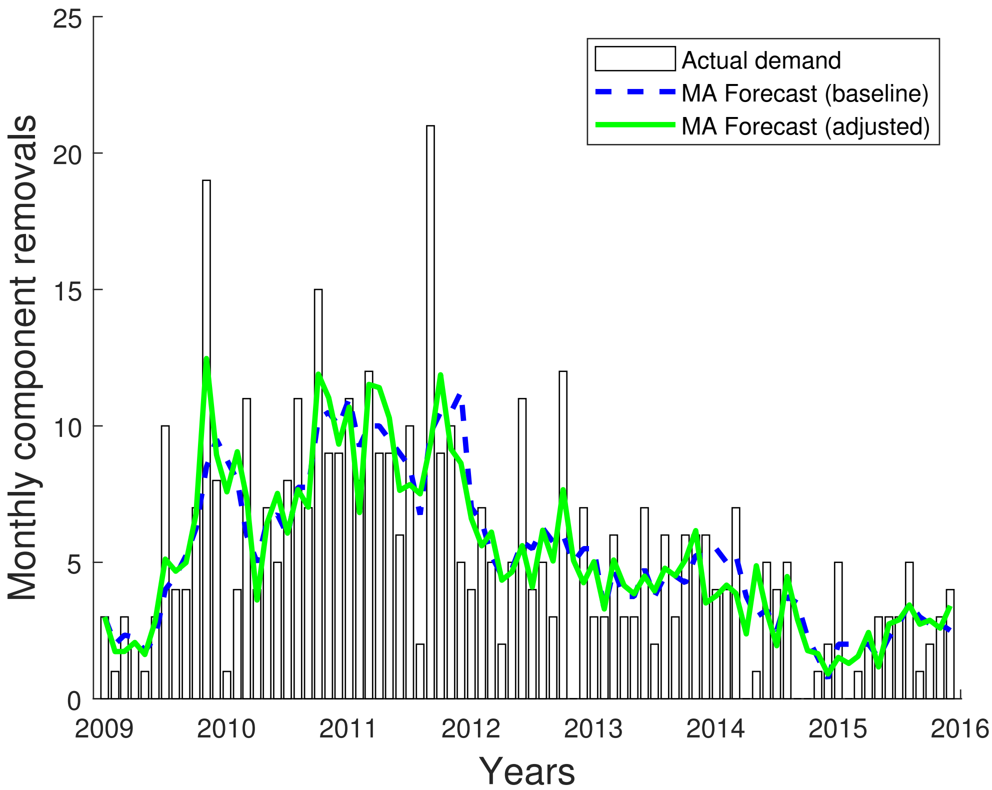

Figure 7.1: Actual demand, MA forecast and adjusted forecast for component type 1

Based on Figures 7.1 and 7.2 it is not directly clear from visual inspection which method was the most

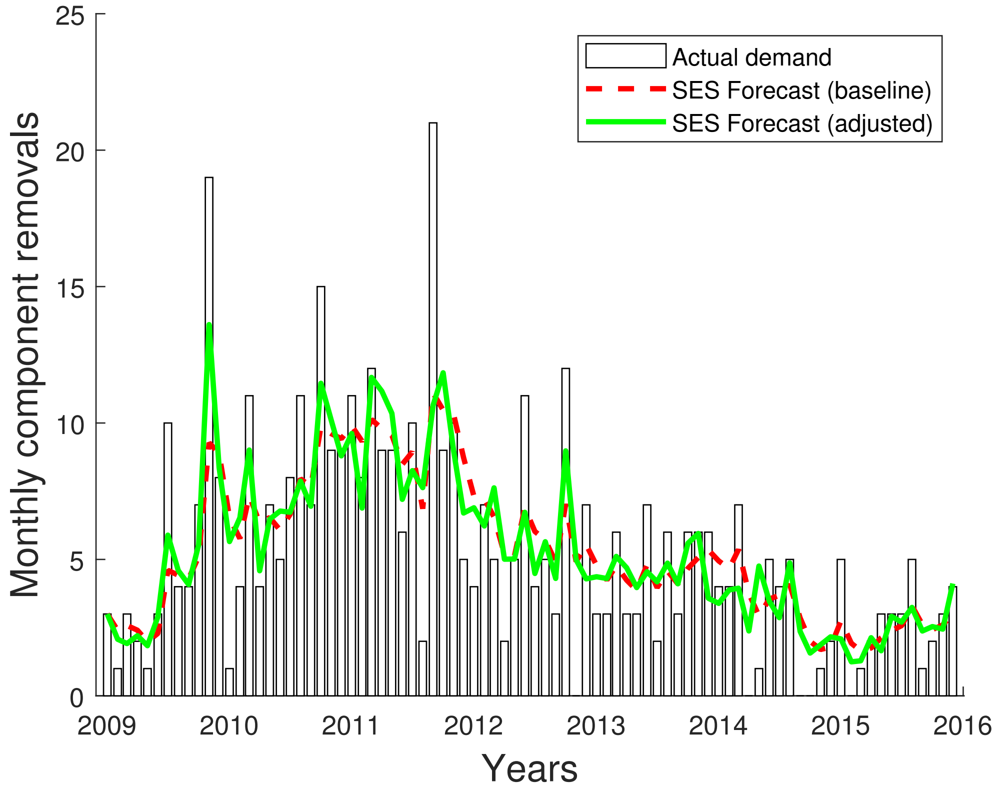

Figure 7.2: Actual demand, SES forecast and adjusted forecast for component type 1

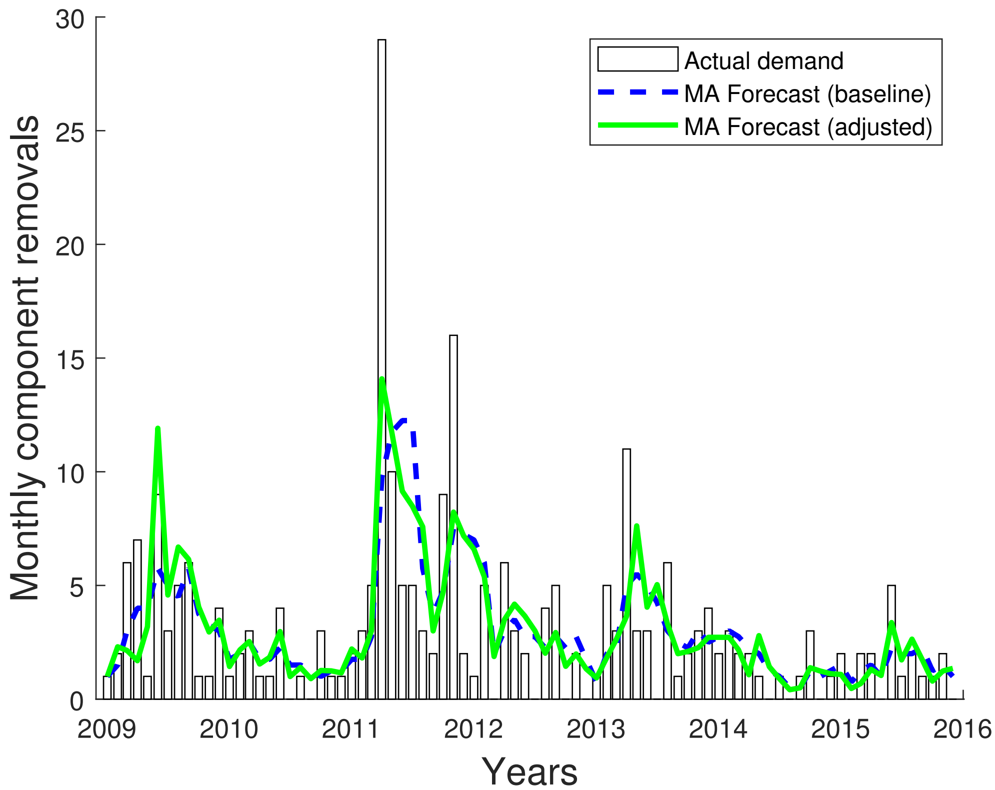

Figure 7.3: Actual demand, MA forecast and adjusted forecast for component type 7

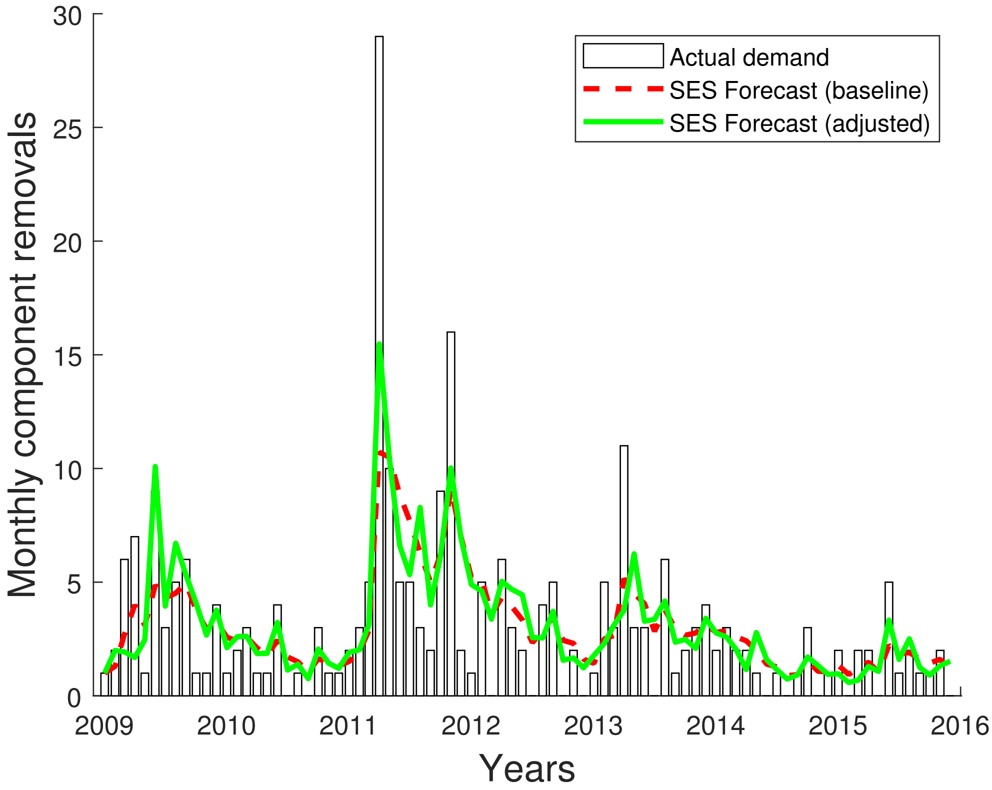

Figure 7.4: Actual demand, SES forecast and adjusted forecast for component type 7

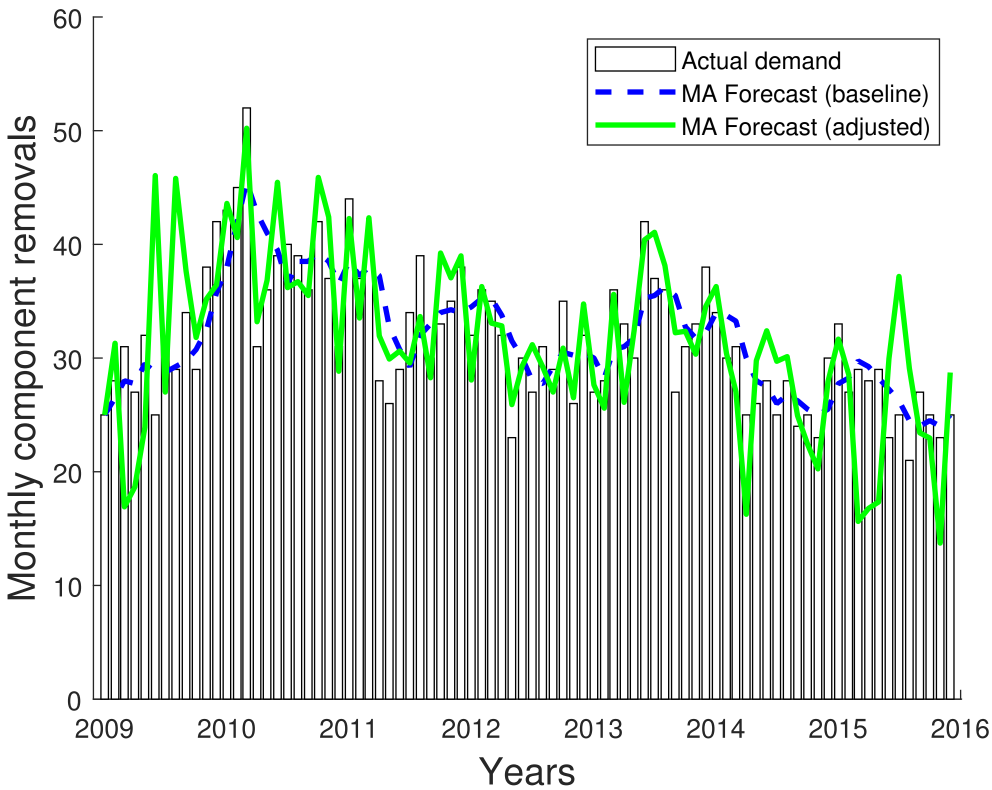

Figure 7.5: Actual demand, MA forecast and adjusted forecast for component type 12

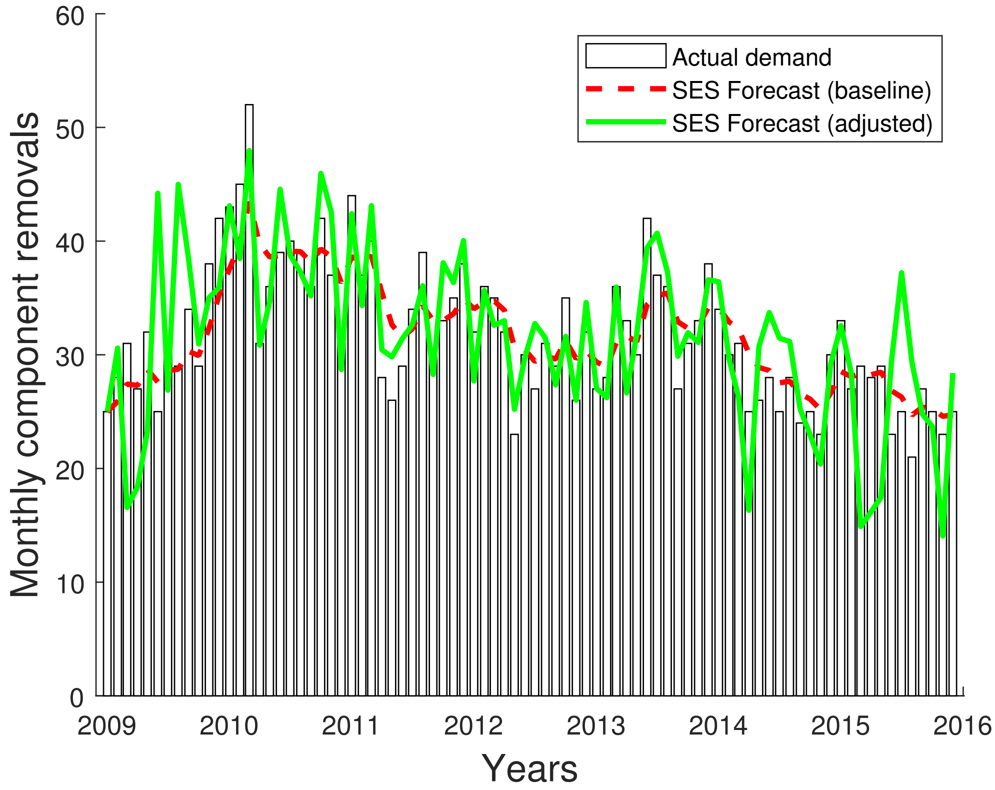

Figure 7.6: Actual demand, SES forecast and adjusted forecast for component type 12

accurate in forecasting the spare parts of CT1, but it does seem that the adjusted method follows the outlier months more accurately. Judging from Figures 7.3 and 7.4, the adjusted MA and SES methods are more suitable to forecast the spare parts of CT7, which is the most erratic data set of the validation data sets. This can especially be seen from the fact that the adjusted MA and SES methods are more accurate in forecasting the most outlying month in 2011.

Finally, looking at Figures 7.5 and 7.6, it appears that the adjusted methods are less accurate than the baseline methods in forecasting the spare parts of the smooth CT12 data set. This can be observed from the fact that the demand volumes predicted by the adjusted methods appear to be more volatile compared to the actual demand pattern and the demand forecast by the baseline methods. These observations initially imply that the adjusted methods are more accurate in forecasting erratic demand patterns and less accurate in forecasting smooth demand patterns. To confirm these conclusions, the performance of the baseline and adjusted methods will be evaluated in Section 7.3.

7.3 Evaluation of forecast performance of the applied methods

This section will evaluate the forecast errors of the baseline and adjusted forecasting method on the validation datasets. However since some of the validation datasets contain few zero-demand months, the MAPE error metric is not suitable to be used for the evaluation, as the MAPE produces computation errors when dealing with zero-demand months. For the evaluation, the RMSE error metric which was introduced in Chapter 3 would be more suitable, but a major drawback of this metric is its scale dependency. This fact makes it challenging to effectively compare the forecast of multiple datasets solely using the RMSE metric, since the average demand volumes and scale of the data also has an impact on the value of the RMSE.

For this reason, a normalised measure of the total error will be introduced first, since this metric is most suitable for comparing the forecast performance on multiple datasets. This metric is called Theil's U-statistic and is mathematically represented by Equation 7.1. Theil's U-statistic yields a value between 0 and 1, with smaller values of the U-statistic indicating a better forecasting performance, where a U-statistic equal to 0 represents a perfect fit [32].

$$U = \frac{\sqrt{\frac{1}{n}\sum_{t=1}^{n} e_t^2}}{\sqrt{\frac{1}{n}\sum_{t=1}^{n} f_t^2}\sqrt{\frac{1}{n}\sum_{t=1}^{n} y_t^2}} \tag{7.1}$$

To confirm whether the suggested adjusted forecasting methods improve the accuracy of the validation data sets, the Theil's U-statistic has been determined for the baseline methods and the adjusted methods, for each component type dataset. The statistical relation to Pilot Complaints and Aircraft Landings has also been determined for each data set, in a similar fashion as described in Section 5.4. Table 7.2 shows an overview of the U-statistic values for the baseline and adjusted methods, the difference between these values (δU), the CV-values and the correlation coefficients with Pilot Complaints (CC_{PIC}) and Aircraft Landings (CC_{LND}) corresponding to each validation data set. Note that in this Table, a positive δU should be interpreted as a reduction of the U-statistic (thus an improvement in forecasting accuracy), while a negative δU indicates an increase of the U-statistic (and therefore a decrease in forecasting accuracy). Furthermore, Figures 7.7 and 7.8 present a graphical representation of the U-statistic values for the MA method and SES method, respectively.

Table 7.2: Overview of U-statistic values for validation data sets

CT	CV^2	CC_{PIC}	CC_{LND}	U_{MA}	$U_{MA,adj}$	δU_{MA}	U_{SES}	$U_{SES,adj}$	δU_{SES}
CT1	0.75	0.628	0.373	0.0727	0.0655	0.0071	0.0640	0.0569	0.0071
CT2	0.84	0.531	0.285	0.1515	0.1276	0.0239	0.1332	0.1081	0.0251
CT3	0.54	0.368	0.00	0.0382	0.0347	0.0035	0.0341	0.0307	0.0034
CT4	0.78	0.379	0.275	0.1731	0.1666	0.0065	0.1540	0.1486	0.0054
CT5	0.58	0.468	0.381	0.0571	0.0570	0.0002	0.0518	0.0527	-0.0010
CT6	0.34	0.386	0.186	0.0164	0.0168	-0.0003	0.0150	0.0155	-0.0004
CT7	1.3	0.391	0.233	0.1612	0.1370	0.0242	0.1504	0.1229	0.0275
CT8	0.84	0.022	0.110	0.1560	0.1536	0.0024	0.1434	0.1413	0.0021
CT9	0.46	0.426	0.322	0.0494	0.0499	-0.0005	0.0442	0.0449	-0.0008
CT10	0.43	0.309	0.397	0.0342	0.0310	0.0032	0.0306	0.0276	0.0030
CT11	0.52	0.292	0.126	0.0552	0.0534	0.0018	0.0487	0.0477	0.0010
CT12	0.19	0.731	0.264	0.0036	0.0054	-0.0017	0.0033	0.0052	-0.0020

Looking at the results presented in Table 7.2 and Figures 7.7 and 7.8, several conclusions can be drawn regarding the performance of the applied forecasting methods. First of all, it can be stated that the suggested methodology to incorporate the key factors has generally led to improvements in forecasting accuracy, thus successfully validating the approach. The adjusted method has led to a minor increase of the U-statistic in only three cases for the MA method and four cases for the SES method, thereby decreasing the accuracy of the baseline methods slightly. For the majority of component types, a decrease in U-statistic was found and therefore an improvement in forecasting accuracy is established. These results show that including additional statistically correlated factors when forecasting improves the forecasting accuracy not only for Wheels & Brakes components, but also for additional component categories in the MRO database.

Furthermore it can be stated that overall, the baseline and adjusted SES method are more accurate in forecasting the demand volumes compared to the baseline and adjusted MA method, respectively. This can be concluded from the fact that all values for the U-statistic are smaller for the SES method than for the MA method, with the smallest U-statistic values obtained for the adjusted SES method. Based

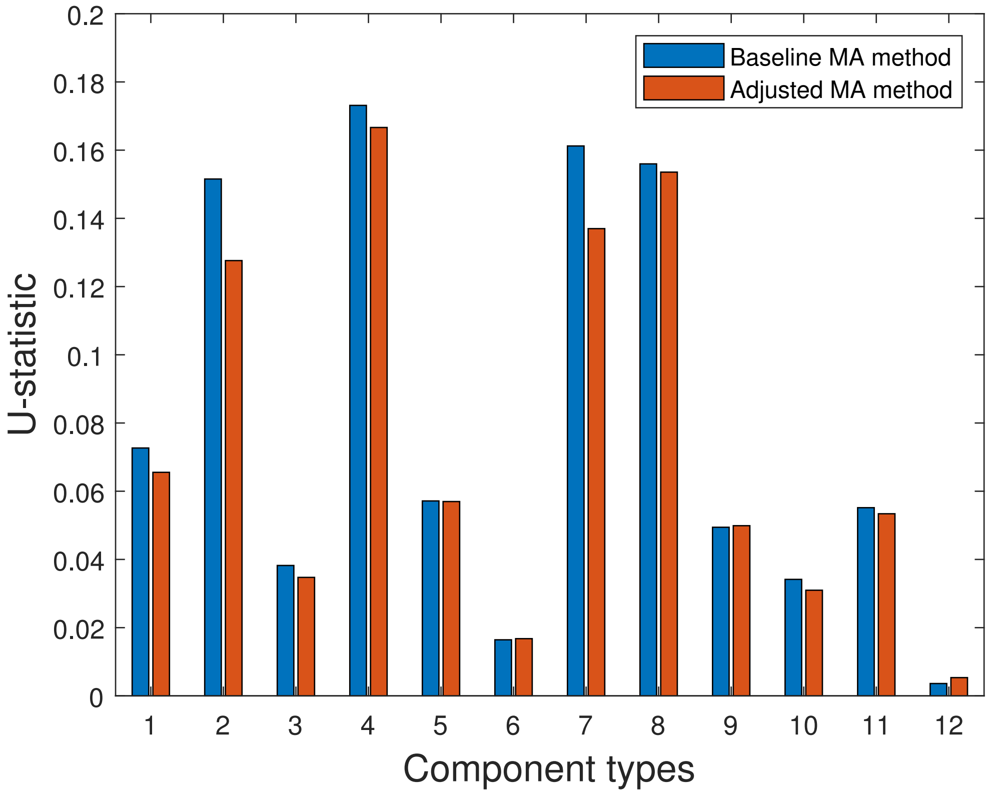

Figure 7.7: U-statistic values for Baseline MA and Adjusted MA methods

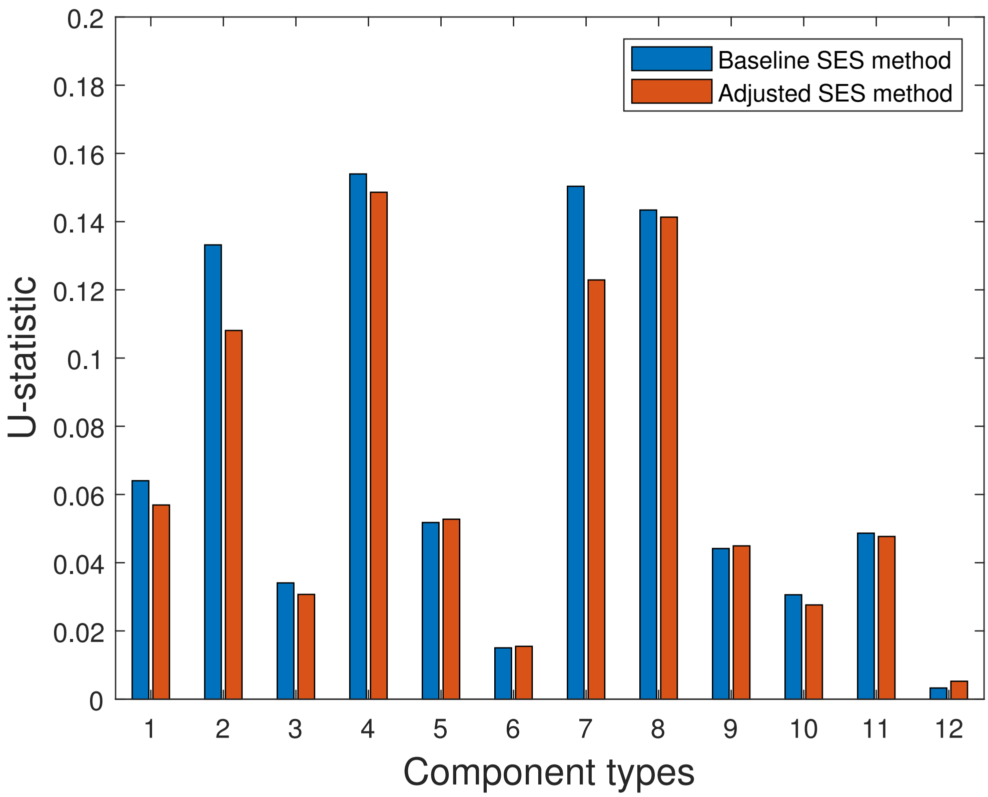

Figure 7.8: U-statistic values for Baseline SES and Adjusted SES methods

on this fact, it can be concluded that the adjusted SES method is the most accurate forecasting method for all component types, and therefore should be recommended as the most suitable method to forecast the demand volumes of components within the specific MRO database.

Finally, by observing the values in the δU_{MA} and δU_{SES} columns of Table 7.2 and comparing them to the values in the CV^2 column, there appears to be a relation between the CV^2-value of a data set and the extent to which the U-statistic of that data set can be decreased by incorporating the statistically relevant factors. Based on the presented results, it appears as if the data sets that inherently have larger variation in demand size (and thus a higher CV^2-value) would benefit the most from following the trend of the causal factors, since these data sets generally show higher values for δU. This can be especially seen in the data sets with the highest and lowest CV^2-value, where the data set of CT7 (with the highest CV^2-value of **1.3**) has shown the largest improvement in the U-statistic, whereas the data set of CT12 (with the lowest CV^2-value of **0.19**) shows the largest deterioration of Theil's U-statistic.

To confirm if there exists a positive correlation between the CV^2-value of a data set and the δU that can be obtained, a linear regression has been applied on the scatter plots between these units. This regression is depicted in Figures 7.9 and 7.10 for the MA methods and for the SES methods, respectively. In these figures, the improvement in Theil's U-statistic after adjusting both the MA method and SES method is shown on the y-axis, versus the CV^2-value of the data sets on the x-axis. Also, the data points in these scatter plots are numbered to indicate which component type data set they represent. Furthermore, Table 7.3 provides the R^2 and p-values for both linear regressions. Based on the p-values provided in this table, it can be stated with a 95% confidence interval that there exists a linear relation between the δU and CV^2 for both of the adjusted MA and SES methods, since the p-values are lower than **0.05**.

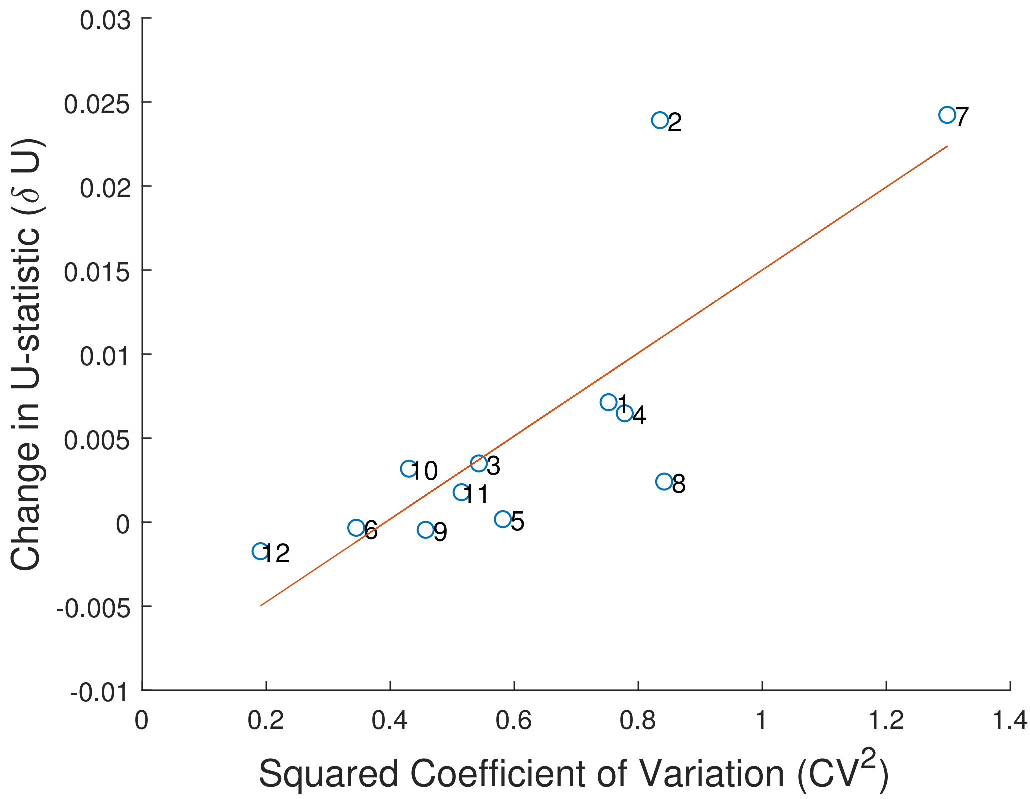

Figure 7.9: Scatter plot with linear regression for δU_{MA} versus CV^2

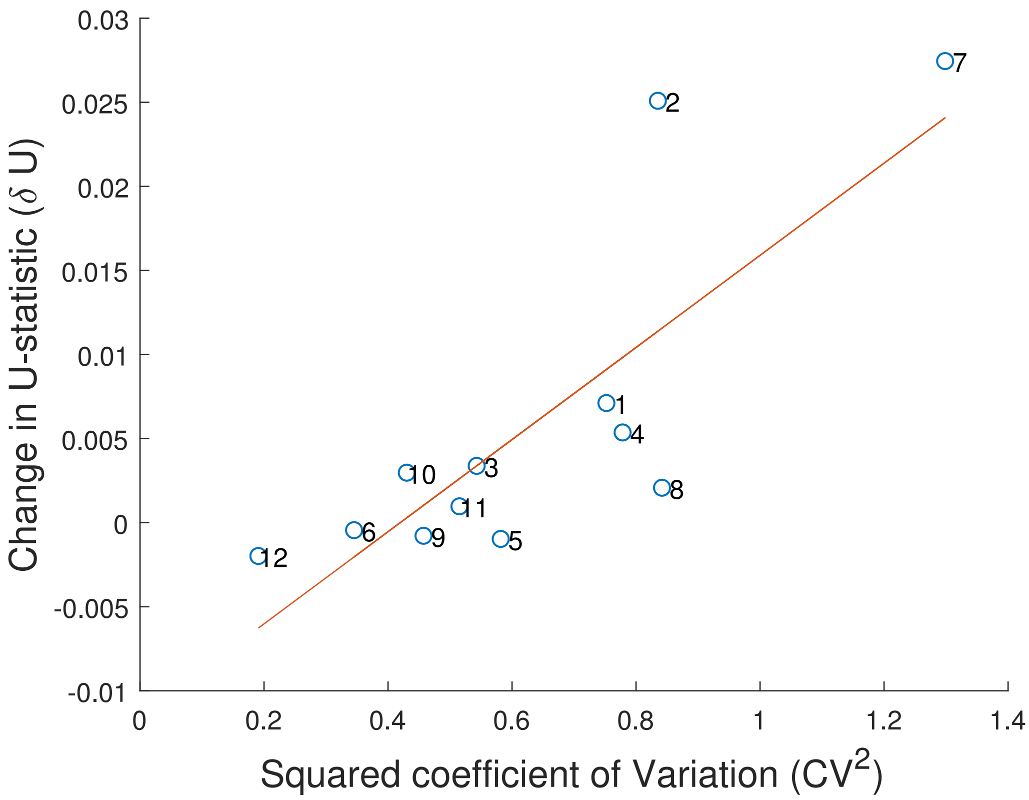

Figure 7.10: Scatter plot with linear regression for δU_{SES} versus CV^2

Table 7.3: Linear regression statistics corresponding to the linear fits between δU and CV^2

	R^2	p-value
Adjusted MA-method	0.6547	0.0014
Adjusted SES-method	0.6523	0.0015

7.4 Sensitivity analysis of baseline and adjusted methods

The results obtained in the research so far have indicated that generally, the existing MA and SES methods can improve their forecasting accuracy and performance if they are adjusted to follow the trend of the statistically significant causal factors Pilot Complaints and Aircraft Landings. However while applying the baseline and adjusted forecasting methods, some inherent user-set parameters are assumed, which have a direct impact on the results of the forecast. This section will therefore describe the sensitivity analysis that was performed to determine if the main conclusions and insights found in the initial analysis are sensitive to change if there are (minor) alterations in these user-set parameters or not.

For both the baseline and adjusted methods several parameters and units were assumed in the analysis. These parameters can be listed as follows:

- $\alpha = 0.3$, the smoothing constant used in the baseline SES method. It determines how reactive the forecasting method is to its forecasting errors.
- $m = 3$, the moving time-window used in the baseline MA method. It determines the number of past months that are considered when calculating the average value of these months.
- PC_0, the average number of Pilot Complaints in the past **three months** used in the adjusted MA and SES methods. This average value can be adjusted to be taken from bigger or smaller time frames.

- LD_0, the average number of Aircraft Landings in the past **three months** used in the adjusted MA and SES methods. This average value can be adjusted to be taken from bigger or smaller time frames.

The following subsections will describe if and how slightly altering the parameters described in the previous list will impact the main conclusions of the analysis. First, a sensitivity analysis will be applied to the user-set parameters of the baseline forecasting methods, after which the sensitivity analysis will be applied to the parameters of the adjusted forecasting methods.

7.4.1 Sensitivity analysis of baseline FC method parameters

First, a sensitivity analysis is performed by varying the value used as the smoothing constant α in the SES method. Initially this constant was set equal to **0.3** when applying the baseline SES method. For the purpose of this sensitivity analysis, this constant is varied from **0.1** to **0.5**, and the impact on the δU is assessed. Table 7.4 presents the δU values for each of the validation data sets, for the five smoothing constant cases.

Table 7.4: Overview of change in U-statistic values for multiple values of α

	δU_{SES}				
CT	$\alpha = 0.1$	$\alpha = 0.2$	$\boldsymbol{\alpha = 0.3}$	$\alpha = 0.4$	$\alpha = 0.5$
CT1	0.0048	0.0064	**0.0071**	0.0070	0.0060
CT2	0.020	0.0234	**0.0251**	0.0252	0.0236
CT3	0.0031	0.0035	**0.0034**	0.0027	0.0016
CT4	0.0037	0.0046	**0.0054**	0.0056	0.0051
CT5	0.0003	-0.0003	**-0.0010**	-0.0018	-0.0027
CT6	-0.0003	-0.0003	**-0.0004**	-0.0006	-0.0008
CT7	0.0183	0.0241	**0.0275**	0.0279	0.0250
CT8	0.0022	0.0023	**0.0021**	0.0015	0.0007
CT9	-0.0002	-0.0004	**-0.0008**	-0.0013	-0.0021
CT10	0.0030	0.0032	**0.0030**	0.0025	0.0018
CT11	0.0009	0.0010	**0.0010**	0.0008	0.0004
CT12	-0.0013	-0.0016	**-0.0020**	-0.0023	-0.0027

Based on the results presented in Table 7.4 it can be stated that varying the assumed value for α in the baseline SES method leads to the same conclusions that were found in the initial research; adjusting the baseline SES method to follow the trend of the key factors will generally lead to a more accurate forecast of the demand size of components. This can be seen from the fact that in the majority of cases, a positive δU is observed for the same data sets as in the initial research. Varying the value of α does slightly impact the extent to which the accuracy of the forecast is improved (or in few cases, deteriorates). With that, it can be said that the main conclusions of the research are not sensitive to changes in the assumed α-parameter.

Next, the sensitivity analysis is performed by varying the time frame $\boldsymbol{m}$ used to compute the forecast in the MA method. In the initial research, this parameter was set equal to **3** when applying the baseline MA method. This constant is varied from **1** to **5** for the sensitivity analysis, and again the impact on the $\boldsymbol{\delta U}$ is assessed. Table 7.5 shows the $\boldsymbol{\delta U}$ values for each of the validation data sets, for the five different cases.

Based on the results presented in Table 7.5, it can be seen that varying the assumed value for $\boldsymbol{m}$ has no impact on the general conclusions found in the initial research; the majority of the data sets would benefit from using the adjusted forecasting methods, as the most cases show a reduction of Theil's U-statistic (a positive $\boldsymbol{\delta U_{MA}}$). As such, the main findings in the initial research are not sensitive to changes in the assumed parameter $\boldsymbol{m}$. However, the assumed value for $\boldsymbol{m}$ does seem to impact the

Table 7.5: Overview of change in U-statistic values for multiple values of m

	δU_{MA}				
CT	$m = 1$	$m = 2$	$m = 3$	$m = 4$	$m = 5$
CT1	0.0050	0.0079	**0.0071**	0.0069	0.0065
CT2	0.0204	0.0262	**0.0239**	0.0220	0.0217
CT3	0.0004	0.0019	**0.0035**	0.0039	0.0045
CT4	0.0031	0.0058	**0.0065**	0.0053	0.0048
CT5	-0.0032	-0.0013	**0.0002**	0.0001	-0.0004
CT6	-0.0007	-0.0003	**-0.0003**	-0.0003	-0.0003
CT7	0.0211	0.0260	**0.0242**	0.0201	0.0188
CT8	0.0004	0.0016	**0.0024**	0.0023	0.0021
CT9	-0.0017	-0.0013	**-0.0005**	0.0002	-0.0003
CT10	0.0001	0.0025	**0.0032**	0.0037	0.0035
CT11	0.0006	0.0020	**0.0018**	0.0011	0.0008
CT12	-0.0025	-0.0019	**-0.0017**	-0.0016	-0.0016

extent to which Theil's U-statistic is reduced.

In some cases, altering the value of m even leads to the baseline MA method being more accurate than the adjusted MA method. This can be seen in the results for the CT5 and CT9 data sets, where it seems that for lower values of m, the baseline MA method outperforms the adjusted MA method. This can be concluded from the negative δU_{MA}-values that are observed for these component types specifically. Although for these specific data sets the baseline MA method is more accurate than the adjusted MA method, for two-thirds of the data sets a positive δU_{MA} is observed, thus supporting the conclusion that the adjusted MA method is more accurate in forecasting than the baseline MA method.

7.4.2 Sensitivity analysis of adjusted FC method parameters

Finally, a sensitivity analysis will be applied to the adjusted forecasting methods. This will be done by varying the time window used in determining the average value for PC_0 and LD_0, which are respectively the average number of pilot complaints and the average number of aircraft landings, in the past 'k' months. For the purpose of this analysis, k is varied from 1 to 5 to determine if the conclusions of the research are sensitive to changes in the assumed parameters in the adjusted FC methods. Table 7.6 shows an overview of the δU_{MA} and δU_{SES} resulting from the sensitivity analysis. In this analysis, the value of the past 'k' months has been varied when computing the average number of pilot complaints (PC_0).

Judging from the results presented in Table 7.6, it can be stated that altering the value of k when computing PC_0 has no major impacts on the general conclusions obtained in the initial research. Still, for the vast majority of data sets improvements in forecasting accuracy could be determined by adjusting the baseline MA and SES methods, regardless of the assumed value for k. The assumed value for k does impact the extent to which an improvement can be realised, with higher values for k often leading to larger values of δU. This can especially be seen for component types 2, 4, 7 and 11, where the δU seems to increase for increasing values of k. As such, it can be recommended to set this parameter equal to at least three months, in order to increase the forecasting accuracy improvements that can be realised.

Similarly, the same sensitivity analysis has been applied to the computation of the average number of aircraft landings (LD_0), where the value of the past k months again ranges from 1 to 5. An overview of δU-statistic values resulting from these changes is presented in Table 7.7.

Based on the results presented in Table 7.7, it can be stated that altering the time frame in determining the average number of Aircraft Landings barely has an impact on the obtained δU-values. The

Table 7.6: Overview of change in U-statistic values for multiple values of k months when computing PC_0

	δU_{MA}					δU_{SES}				
CT	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$
CT1	0.0021	0.0045	**0.0071**	0.0062	0.0060	0.0019	0.0043	**0.0071**	0.0063	0.0064
CT2	0.0108	0.0204	**0.0239**	0.0242	0.0262	0.0126	0.0218	**0.0251**	0.0256	0.0278
CT3	0.0008	0.0028	**0.0035**	0.0038	0.0031	0.0005	0.0025	**0.0034**	0.0040	0.0036
CT4	0.0027	0.0039	**0.0065**	0.0071	0.0101	0.0019	0.0027	**0.0054**	0.0061	0.0095
CT5	0.0005	0.0000	**0.0002**	-0.0004	-0.0005	-0.0004	-0.0011	**-0.0010**	-0.0013	-0.0014
CT6	0.0002	-0.0002	**-0.0003**	-0.0006	-0.0004	0.0001	-0.0003	**-0.0004**	-0.0006	-0.0005
CT7	0.0110	0.0235	**0.0242**	0.0299	0.0348	0.0143	0.0272	**0.0275**	0.0346	0.0395
CT8	0.0022	0.0021	**0.0024**	0.0020	0.0022	0.0020	0.0019	**0.0021**	0.0017	0.0018
CT9	0.0010	0.0009	**-0.0005**	-0.0006	-0.0005	0.0009	0.0007	**-0.0008**	-0.0009	-0.0009
CT10	0.0025	0.0035	**0.0032**	0.0027	0.0022	0.0023	0.0033	**0.0030**	0.0027	0.0024
CT11	-0.0004	0.0006	**0.0018**	0.0024	0.0025	-0.0006	-0.0001	**0.0010**	0.0016	0.0017
CT12	-0.0010	-0.0012	**-0.0017**	-0.0018	-0.0019	-0.0011	-0.0014	**-0.0020**	-0.0020	-0.0022

Table 7.7: Overview of change in U-statistic values for multiple values of k months when computing LD_0

	δU_{MA}					δU_{SES}				
CT	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$
CT1	0.0070	0.0070	**0.0071**	0.0072	0.0073	0.0072	0.0071	**0.0071**	0.0071	0.0072
CT2	0.0229	0.0234	**0.0239**	0.0242	0.0243	0.0240	0.0245	**0.0251**	0.0253	0.0255
CT3	0.0035	0.0035	**0.0035**	0.0035	0.0035	0.0034	0.0038	**0.0034**	0.0034	0.0034
CT4	0.0064	0.0064	**0.0065**	0.0068	0.0073	0.0055	0.0053	**0.0054**	0.0056	0.0061
CT5	0.0001	-0.0001	**0.0002**	0.0001	0.0002	-0.0009	-0.0012	**-0.0010**	-0.0010	-0.0009
CT6	-0.0001	-0.0002	**-0.0003**	-0.0004	-0.0004	-0.0002	-0.0003	**-0.0004**	-0.0005	-0.0006
CT7	0.0212	0.0236	**0.0242**	0.0250	0.0250	0.0249	0.0267	**0.0275**	0.0283	0.0284
CT8	0.0022	0.0029	**0.0024**	0.0018	0.0011	0.0021	0.0027	**0.0021**	0.0014	0.0007
CT9	-0.0003	-0.0004	**-0.0005**	-0.0005	-0.0006	-0.0006	-0.0007	**-0.0008**	-0.0009	-0.0009
CT10	0.0026	0.0029	**0.0032**	0.0032	0.0032	0.0026	0.0028	**0.0030**	0.0030	0.0031
CT11	0.0018	0.0018	**0.0018**	0.0017	0.0017	0.0009	0.0009	**0.0010**	0.0009	0.0009
CT12	-0.0017	-0.0017	**-0.0017**	-0.0018	-0.0018	-0.0020	-0.0020	**-0.0020**	-0.0020	-0.0021

magnitude of both δU_{MA} and δU_{SES} only slightly changes when a higher or lower value for k is assumed. With that, it can be concluded that average number of Aircraft Landings (LD_O) has a less dominant role in the adjusted forecasting methods, and changing the assumed value for k has no major impacts on the main results found in the research. This less significant role of LD_0 in the adjusted methods can be explained by the fact that in general, the correlation coefficients between the component removals and the Aircraft Landings are lower compared to the correlation coefficients obtained for Pilot Complaints.

CHAPTER 8

Conclusions and Recommendations

With the thesis research completed and the main result obtained, it is now possible to formulate the main conclusions and recommendations based on the findings of the research. This final chapter will therefore first present the most important conclusions in Section 8.1, after which the main recommendations will be described in Section 8.2.

8.1 Conclusions

The main problem of the research was introduced to be as follows; *"The uncertain nature of spare parts demand makes it very challenging for MRO's to accurately forecast the need for spare parts, often leading to sub-optimal operations"*. This problem description then resulted in the following main research question: ***"Will spare parts demand forecasting accuracy improve if inherent causal factors are taken into account while forecasting?"***. With the performed research and the obtained results, this main research question can be answered. The obtained results confirm that it is indeed possible to improve demand forecasting accuracy if the patterns of statistically correlated key factors are taken into account when using conventional time-series forecasting methods. Adjusting the baseline methods to incorporate the behaviour of these key factors has generally lead to reduced forecasting errors for multiple component type demand patterns analysed in this research. An initial analysis of the MRO data base showed that the demand size variability varies per component category, and also varies per aircraft type.

Following from this analysis, it was found that the Wheels & Brakes component category was the most extensive category in the data base, which is why the remainder of the thesis methodology was applied to components belonging to the Wheels & Brakes category. Subsequently, the indirect key factors aircraft operators and aircraft type were used to select and generate specific data patterns for component removals. At the same time, the patterns for the direct causal factors pilot complaints and aircraft landings were generated, after which a statistical analysis was performed to show that there exists a strong positive correlation between the causal factors and component removals for operator 1.

As a result, the baseline forecasting methods were applied to forecast the demand for Wheels & Brakes components of operator 1 specifically. This showed that in general the SES method is more accurate than the MA method, since smaller values of RMSE and MAPE were obtained for the baseline SES method. Following from this, an adjusted forecasting method was suggested, in which the correlation coefficients and the data patterns of the causal factors were also included in determining the forecast value. Furthermore, a seasonality effect was recognised in the patterns for Wheels & Brakes, and accounting for both this seasonality factor and the causal factors resulted in minimal MAPE values for the analysed component patterns, compared to the baseline forecasting methods.

To ensure that these findings do not only hold true for specific Wheels & Brakes components, other component categories of operator 1 were considered as validation data sets. The baseline and adjusted methods were applied to twelve additional validation data sets, and for almost all data sets an improvement of forecasting accuracy could be found when comparing the performance of the adjusted methods to the baseline methods. Only for the most smooth data set, a significant increase of the forecast error was found, which suggests that the adjusted forecasting method are the most suitable for data patterns that are high in demand size variability. To confirm this, a positive relation between the improvement in Theil's U-statistic and the coefficient of variation of the data sets was found, indicating that demand patterns with a higher variability in demand size would benefit the most from using the adjusted method over the baseline method.

Finally, a sensitivity analysis was successfully conducted to confirm that the main findings are not sensitive to minor changes in the assumed model parameters. With that, this research has successfully shown that the adjusted forecasting methods outperform the conventional baseline SES and MA methods for a broad scope of components and parameter settings. These findings are another confirmation that only looking at historic demand when using time-series models is a restricted approach, and that the conventional time-series methods can benefit from including factors that are statistically correlated to the number of component removals.

8.2 Recommendations

Following from the main conclusions described in the previous section, some main recommendations can be identified, which will be described in this section. Since this thesis research has shown positive results for including statistically related factors, the first recommendation would be to include additional causal factors and additional data to incorporate in future analysis. This could be environmental data or physics-based information, which could likely be related to component removals as well.

Furthermore it would be recommended to include more elaborate forecasting methods as the baseline models to compare the adjusted methods with. In this research, the scope was limited to two time-series methods only, but this could be expanded as well. Also, the data sets in this research were mainly erratic, with low variability in demand frequency. This approach could therefore be repeated for more lumpy demand patterns, for which Croston's method would be recommended as the baseline forecasting method to compare the adjusted methods with.

Additionally, the research could be applied to substantially more data sets, to confirm for which types of data sets the forecasting errors can be reduced the most. With an extended sample size of data sets, a more in-depth analysis can be performed to find out if demand patterns with the highest demand size variability do also in fact benefit the most from an adjusted forecasting method. Following from this, the main recommendation would be to extend the applied model with Artificial Intelligence elements to automate the process of finding the most suitable forecasting method which will yield the least forecasting errors. Ideally, this developed model will identify the main characteristics of all data sets to be analysed, and based on these characteristics will apply the most suitable forecasting method with the most suitable user-set parameters.

Bibliography

[1] Gerasimov, V., `http://www.vladstudio.com/wallpaper/?airlines2`, 2008, Cover image, edited.

[2] Frepple, "Demand classification," https://frepple.com/blog/demand-classification/, January 2017.

[3] Zorgdrager, M., Curran, R., Verhagen, W. J. C., Boesten, B. H. L., and Water, C. N., *A predictive method for the estimation of material demand for aircraft non-routine maintenance [online].*, 20th ISPE International Conference on Concurrent Engineering: Proceedings, Amsterdam, Netherlands, 2013.

[4] Wallström, P. and Segerstedt, A., *Evaluation of forecasting error measurements and techniques for intermittent demand*, Vol. 128, International Journal of Production Economics, July 2010, 625-636.

[5] Regattieri, A., Gamberi, M., Gamberini, R., and Manzini, R., *Managing lumpy demand for aircraft spare parts*, Vol. 11, Journal of Air Transport Management, 2005, 426-431.

[6] III, A. F. L. and Ciarallo, F. W., *Reliability and operations: Keys to lumpy aircraft spare parts demands*, Vol. 50, Journal of Air Transport Management, 2015, 30-40.

[7] Donnelly, J. M., *The Case for Managing MRO Inventory*, Vol. 17, Supply Chain Management Review, 2013, 18-25.

[8] Altay, N. and Rudisill, F., *Adapting Wright's modification of Holt's method to forecasting intermittent demand*, Vol. 111, International Journal of Production Economics, 2008, 389-408.

[9] Air Transport World, `http://docslide.us/documents/air-transport-world-january-2013.html`, 2013.

[10] Wang, W. and Syntetos, A. A., *Spare parts demand: Linking forecasting to equipment maintenance*, Vol. 47, Elsevier, Transportation Research Part E Logistics and Transportation Review, March 2011, 1194-1209.

[11] Christer, A. H. and Waller, W. M., *Reducing production downtime using delay time analysis*, Vol. 35, Journal of the Operational Research Society, 1984, 499-512.

[12] Christer, A. H., *Developments in delay time analysis for modelling plant maintenance.*, Vol. 50, Journal of the Operational Research Society, 1999, 1120-1137.

[13] Jones, B., Jenkinson, I., and Wang, J., *Methodology of using delay-time analysis for a manufacturing industry.*, Vol. 94, Reliability Engineering and System Safety, 2009, 111-124.

[14] Ghobbar, A. A. and Friend, C. H., *Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model.*, Vol. 30, Computers & Operations Research, 2003, 2097-2114.

[15] Syntetos, A. A. and Boylan, J. E., *On the variance of intermittent demand estimates*, Vol. 128, International Journal of Production Economics, 2010, 546-555.

[16] Croston, J. D., *Forecasting and stock control for intermittent demand.*, Vol. 23, Operational Research Quarterly, 1972, 289-303.

[17] Syntetos, A. A. and Boylan, J. E., *The accuracy of intermittent demand estimates*, Vol. 21, International Journal of Forecasting, 2005, 303-314.

[18] Syntetos, A. A. and Boylan, J. E., *On the bias of intermittent demand estimates*, Vol. 71, International Journal of Production Economics, 2001, 457-466.

[19] Ghobbar, A. A. and Friend, C. H., *The material requirements planning system for aircraft maintenance and inventory control: a note*, Vol. 10, Journal of Air Transport Management, 2004, 217-221.

[20] Kostenko, A. and Hyndman, R., *A note on the categorization of demand patterns.*, Vol. 57, Journal of the Operational Research Society, 10th ed., October 2006, 1256–1257.

[21] Petropoulos, F., Makridakis, S., Assimakopoulos, V., and Nikolopoulos, K., *Horses for courses in demand forecasting*, Vol. 237, European Journal of Operational Research, 1st ed., October 2014, 152–163.

[22] Willemain, T. R., Smart, C. N., and Schwarz, H. F., *A new approach to forecasting intermittent demand for service parts inventories.*, Vol. 20, International Journal of Forecasting, 2004, 375-387.

[23] Porras, E. and Dekker, R., *An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods.*, Vol. 184, European Journal of Operational Research, 2008, 101-132.

[24] Syntetos, A. A., Babai, M. Z., and Gardner, E. S., *Forecasting intermittent inventory demands: Simple parametric methods vs. bootstrapping.*, Vol. 68, Journal of Business Research Forthcoming, 8th ed., August 2015, 1746-1752.

[25] Ghobbar, A. A. and Friend, C. H., *Sources of intermittent demand for aircraft spare parts within airline operations*, Vol. 8, Journal of Air Transport Management, 2002, 221-231.

[26] Smith, A. M. and Hinchcliffe, G. R., *Reliability-centered Mainteance: Gateway to World Class Maintenance*, Elsevier, 2004.

[27] van der Auweraer, S., Boute, R., and Syntetos, A. A., *Forecasting spare part demand with Installed Base information: a review*, KU Leuven - Faculty of Economics and Business, January 2018.

[28] Jalil, M. N., Zuidwijk, R., Fleischmann, M., and a. E. E. van Nunen, J., *Spare parts logistics and installed base information*, Vol. 62, Journal of the Operational Research Society, 2011, 442-457.

[29] Hellingrath, B. and Cordes, A., *Conceptual approach for integrating condition monitoring information and spare parts forecasting methods*, Vol. 2:1, Production & Manufacturing Research, 2014, 725-737.

[30] Verhagen, W., *AE4-465: Maintenance Modeling & Analysis, Lecture 7 - Supportability engineering: basic concepts & assumptions; forecasting spare parts demand*, Delft University of Technology, 2018, Slide 33.

[31] Verhagen, W., *AE4-465: Maintenance Modeling & Analysis, Lecture 7 - Supportability engineering: basic concepts & assumptions; forecasting spare parts demand*, Delft University of Technology, 2018, Slide 35.

[32] Verhagen, W., *AE4-465: Maintenance Modeling & Analysis, Lecture 7 - Supportability engineering: basic concepts & assumptions; forecasting spare parts demand*, Delft University of Technology, 2018, Slide 41.

Appendices

APPENDIX A

Analysis of Wheels and Brakes data sets

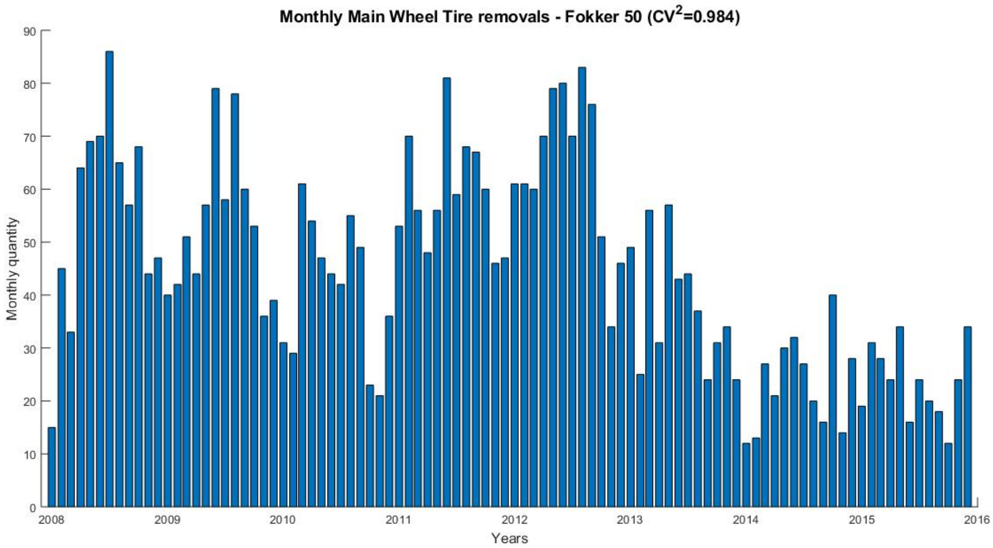

Figure A.1: Monthly component removals of 324-101 category for Fokker 50 between 2008 and 2016

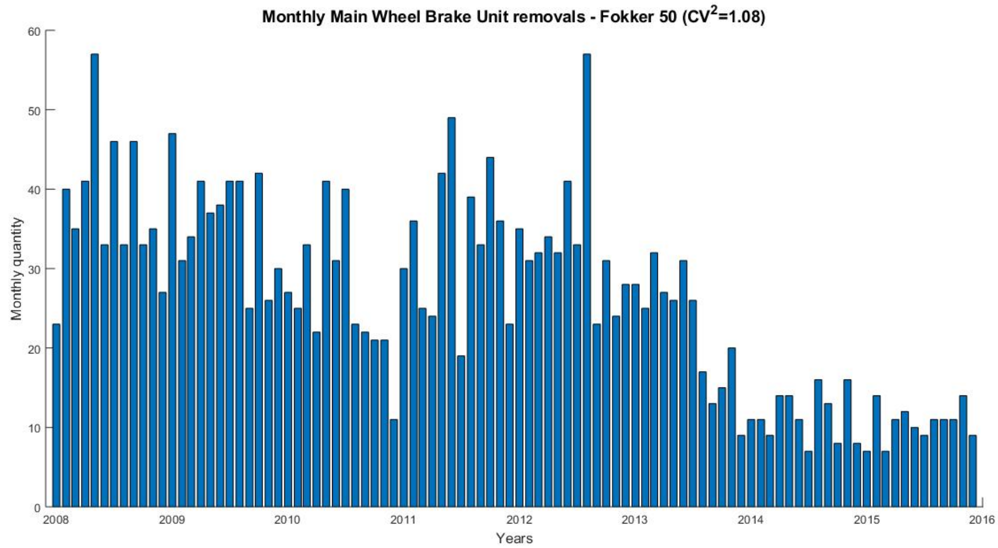

Figure A.2: Monthly component removals of 324-201 category for Fokker 50 between 2008 and 2016

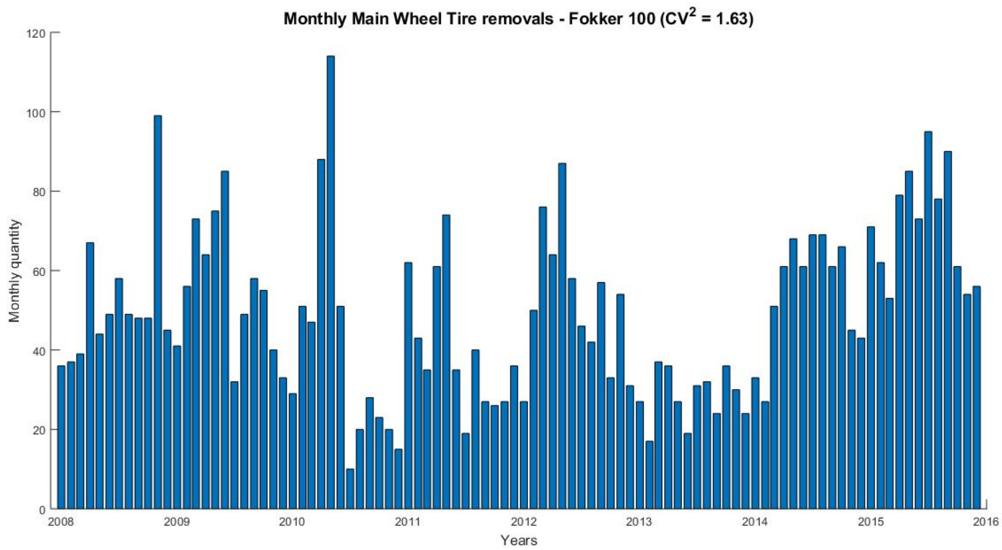

Figure A.3: Monthly component removals of 324-101 category for Fokker 100 between 2008 and 2016

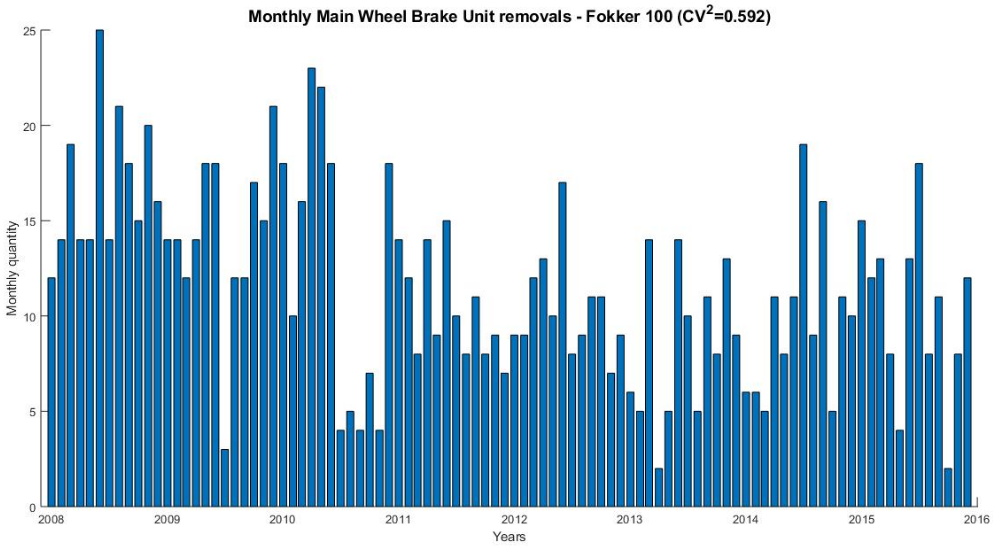

Figure A.4: Monthly component removals of 324-201 category for Fokker 100 between 2008 and 2016

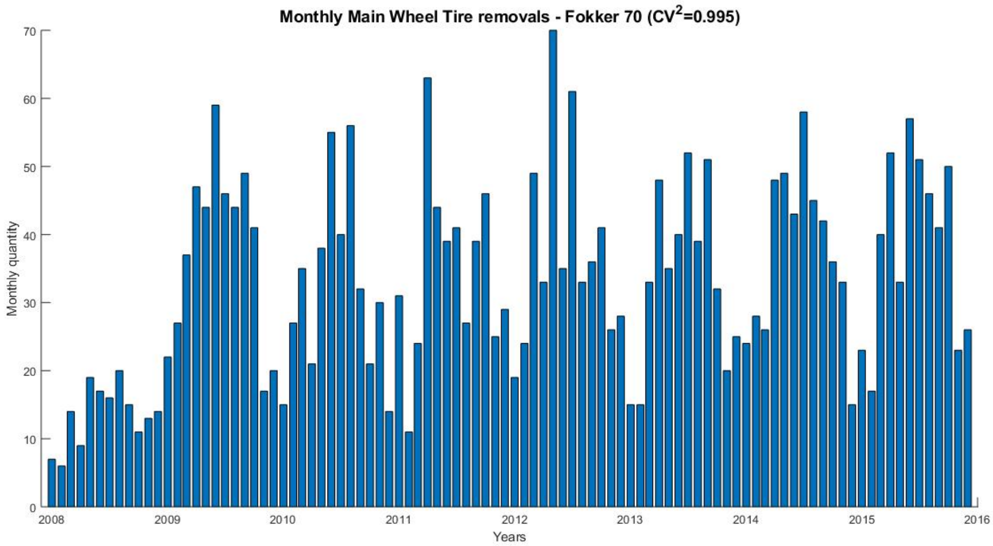

Figure A.5: Monthly component removals of 324-101 category for Fokker 70 between 2008 and 2016

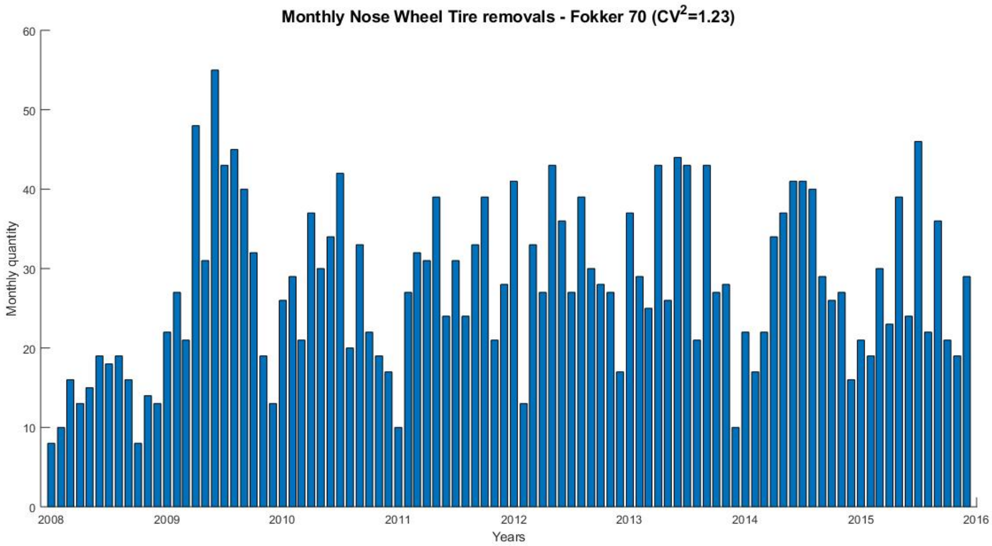

Figure A.6: Monthly component removals of 324-103 category for Fokker 70 between 2008 and 2016

APPENDIX B

Results of Baseline and Altered forecasting methods applied to W&B data sets

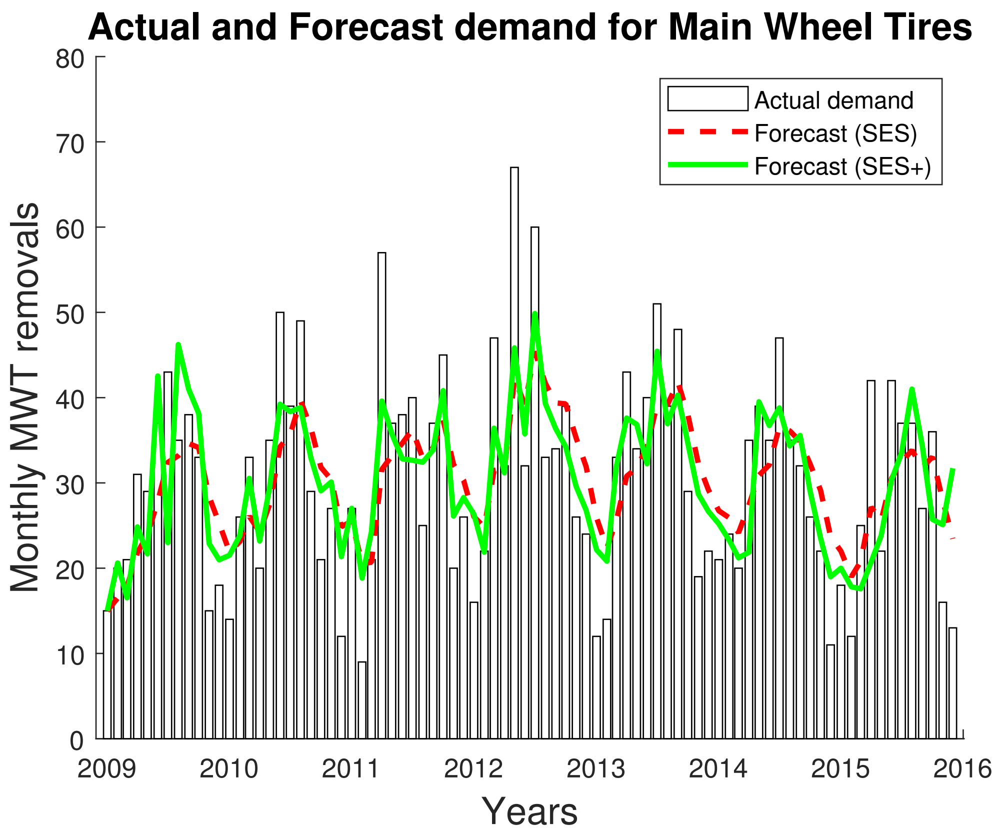

Figure B.1: Actual and SES Forecast demand for Main Wheel Tires

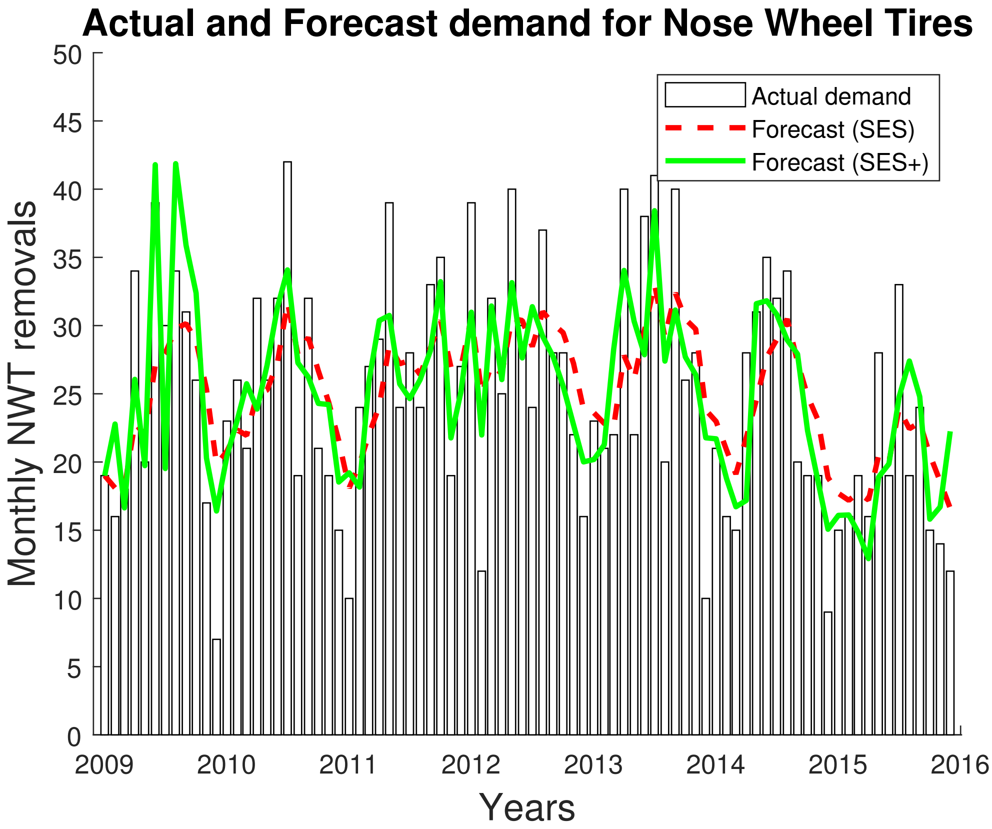

Figure B.2: Actual and SES Forecast demand for Nose Wheel Tires

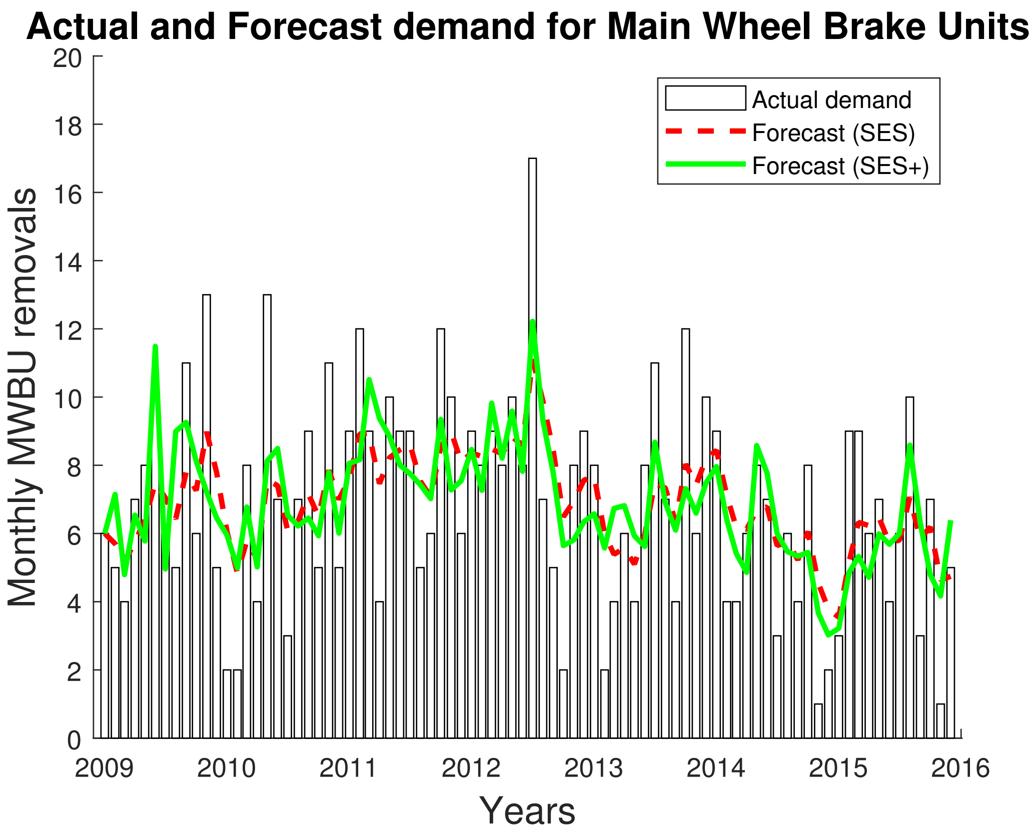

Figure B.3: Actual and SES Forecast demand for Main Wheel Brake Units

APPENDIX C

Results of Baseline and Altered forecasting methods applied to validation data sets

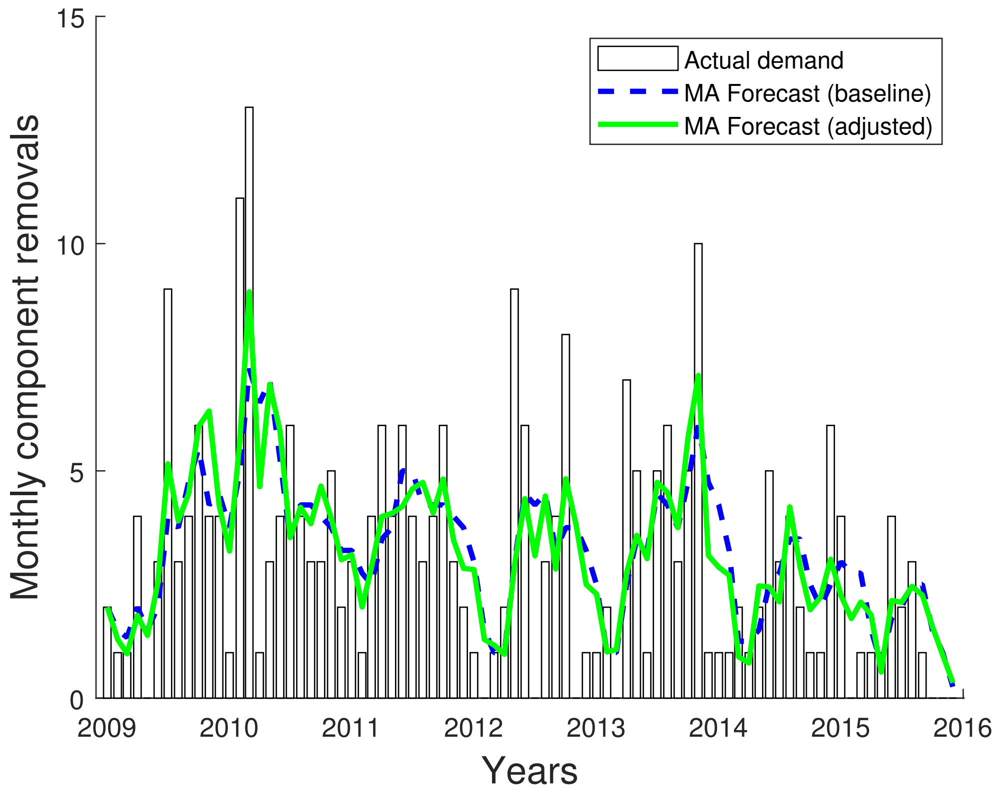

Figure C.1: Actual demand, MA forecast and adjusted forecast for component type 2

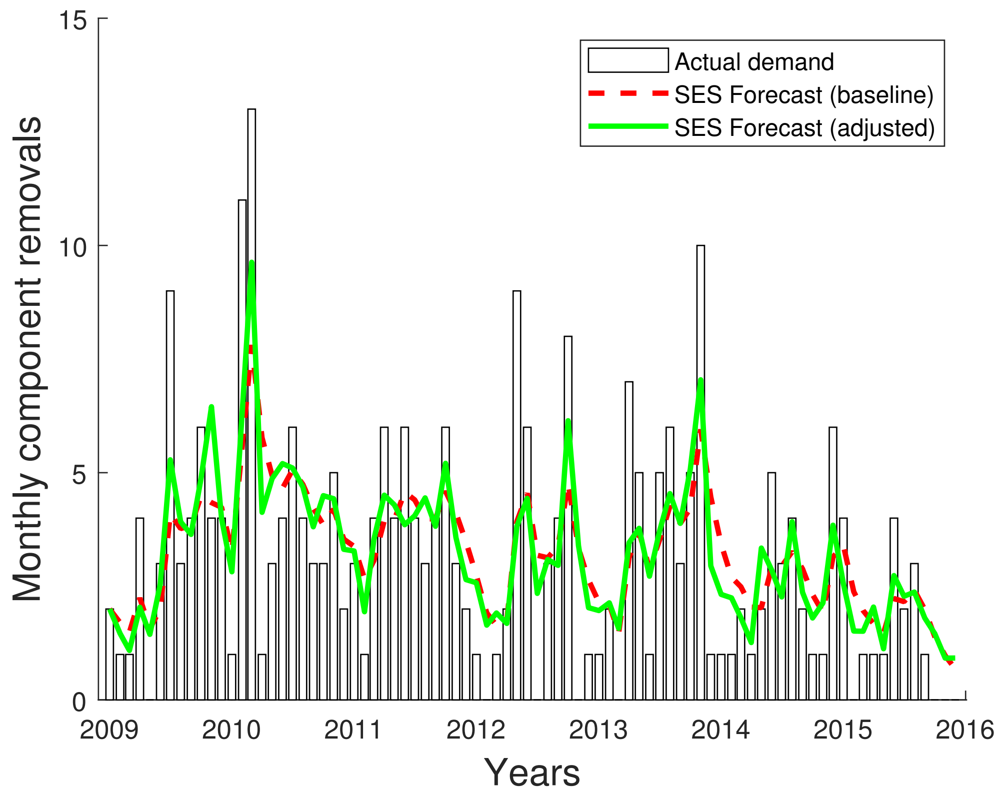

Figure C.2: Actual demand, SES forecast and adjusted forecast for component type 2

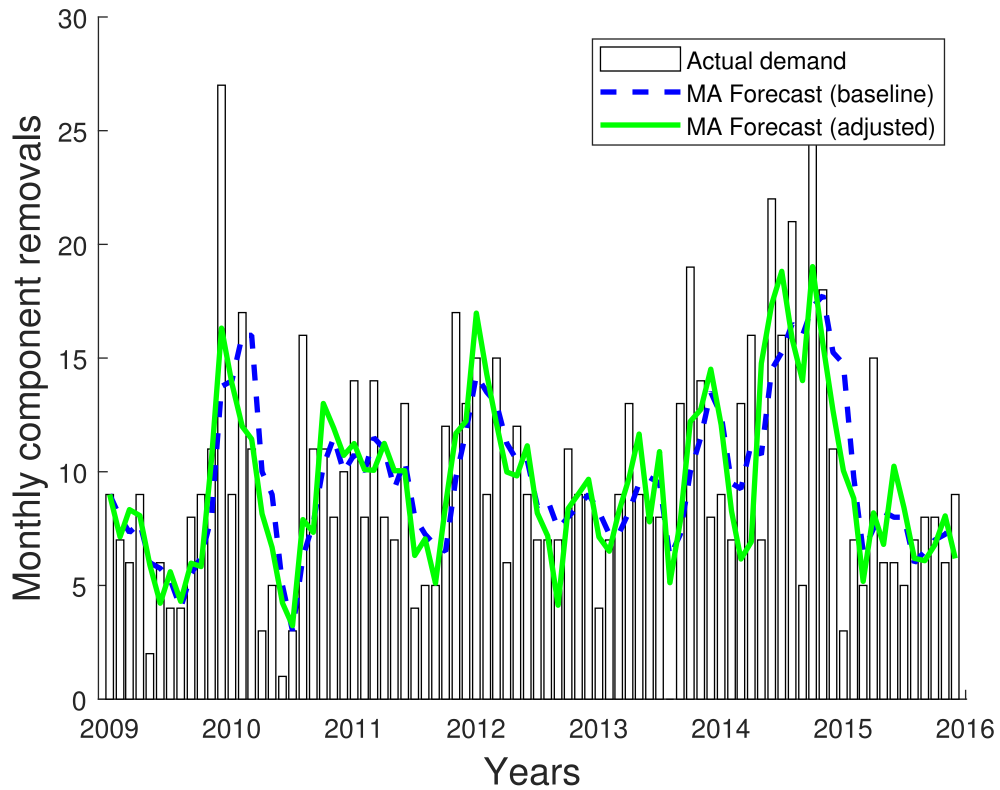

Figure C.3: Actual demand, MA forecast and adjusted forecast for component type 3

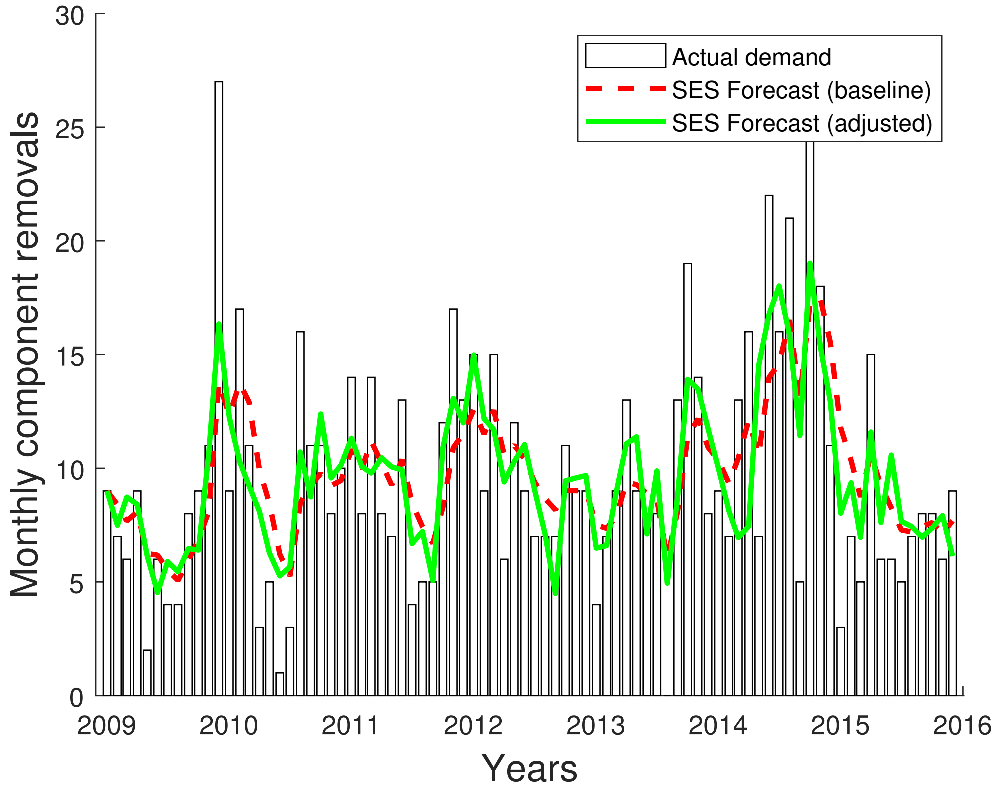

Figure C.4: Actual demand, SES forecast and adjusted forecast for component type 3

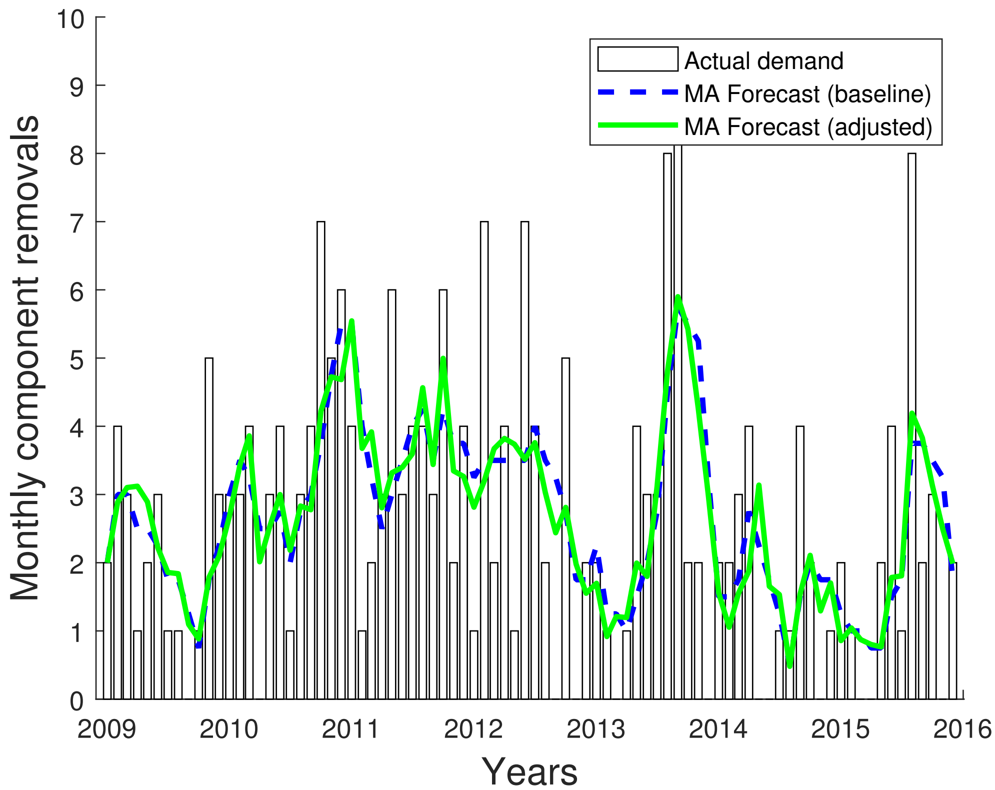

Figure C.5: Actual demand, MA forecast and adjusted forecast for component type 4

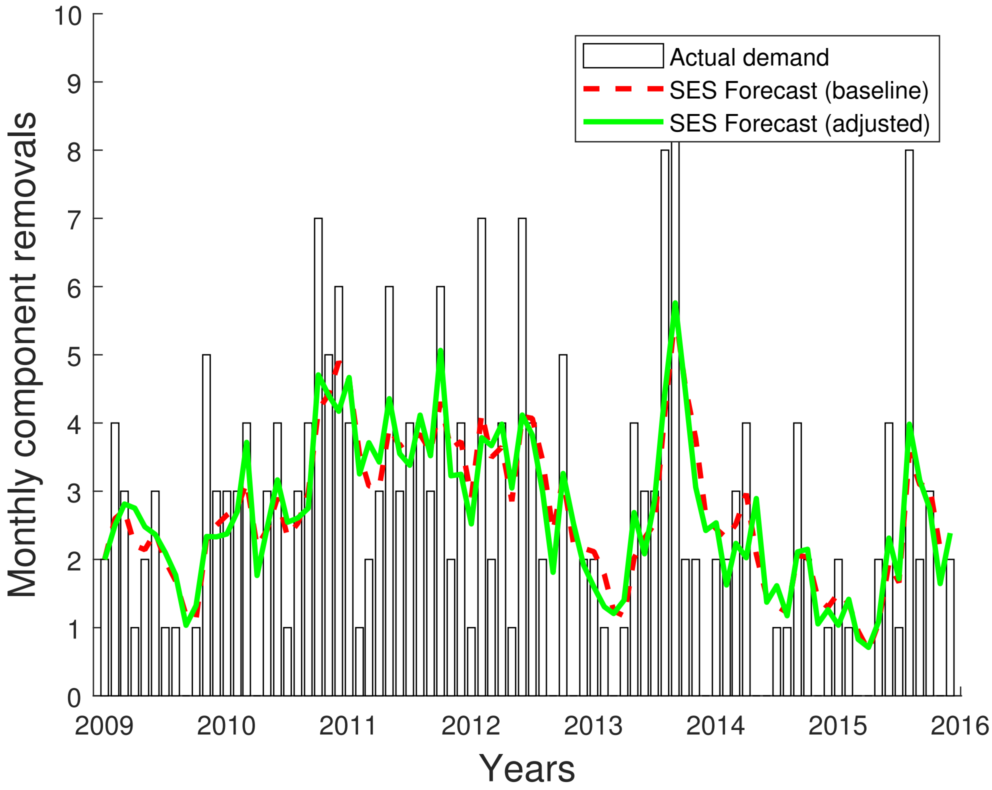

Figure C.6: Actual demand, SES forecast and adjusted forecast for component type 4

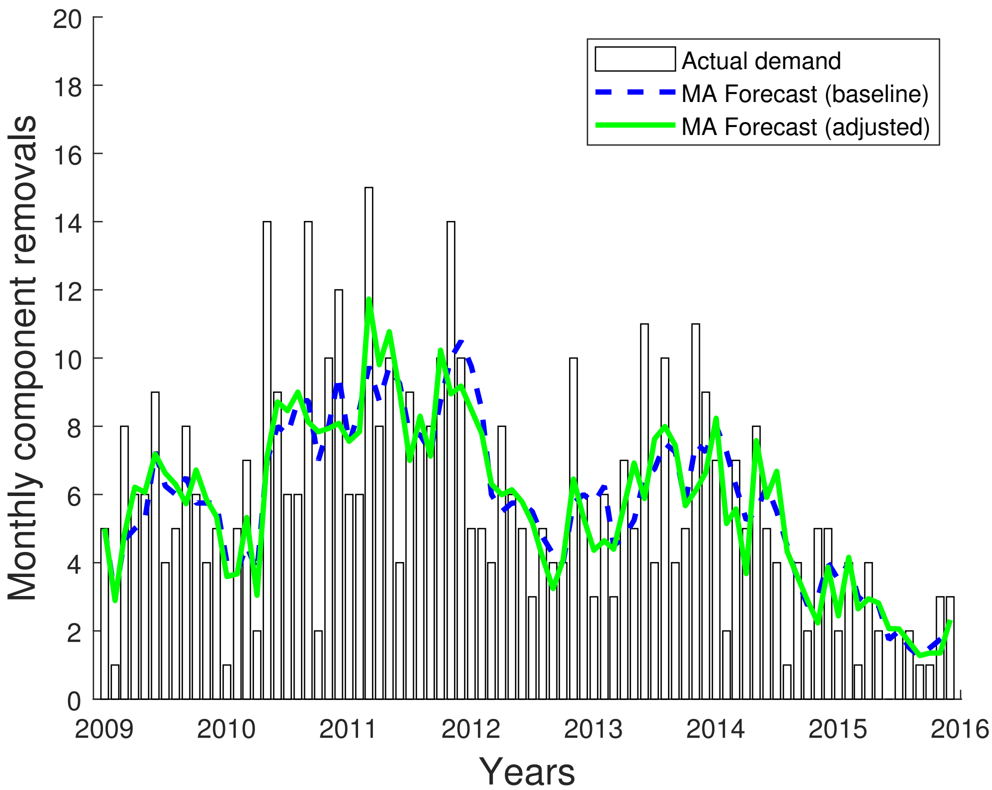

Figure C.7: Actual demand, MA forecast and adjusted forecast for component type 5

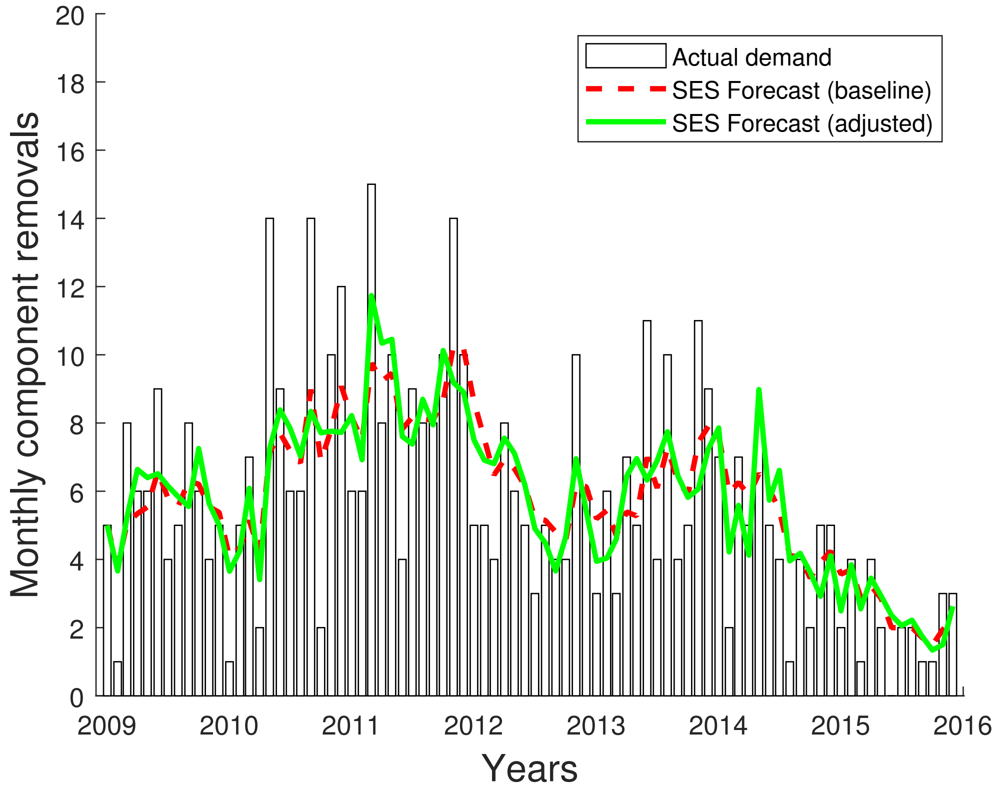

Figure C.8: Actual demand, SES forecast and adjusted forecast for component type 5

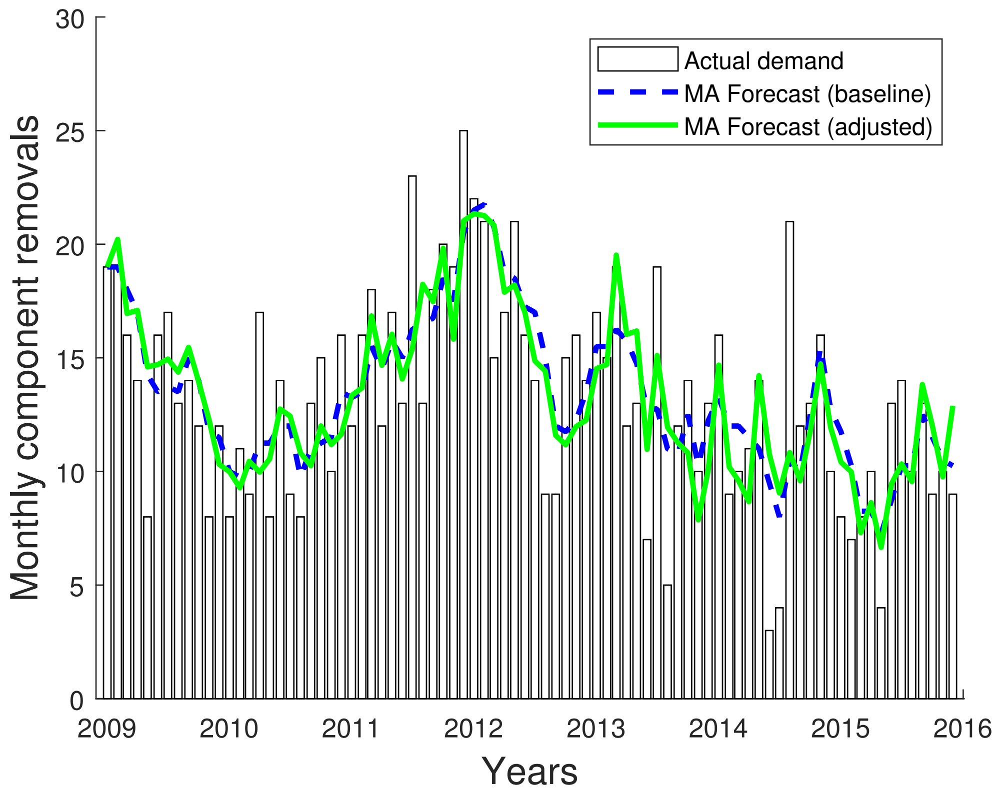

Figure C.9: Actual demand, MA forecast and adjusted forecast for component type 6

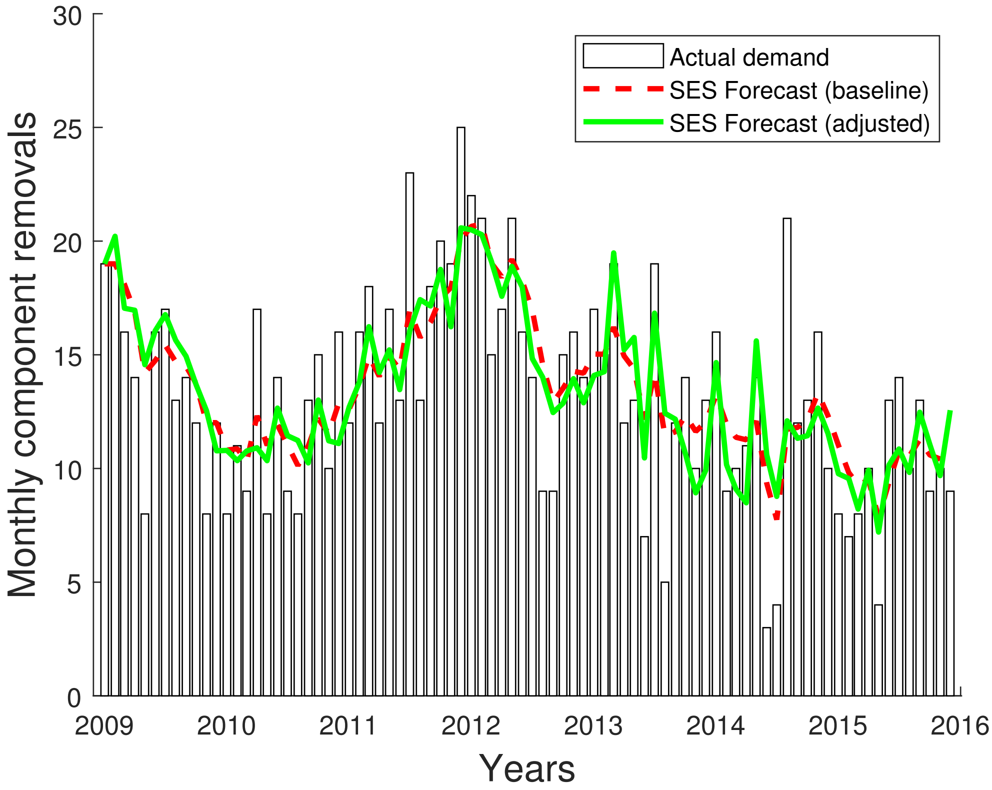

Figure C.10: Actual demand, SES forecast and adjusted forecast for component type 6

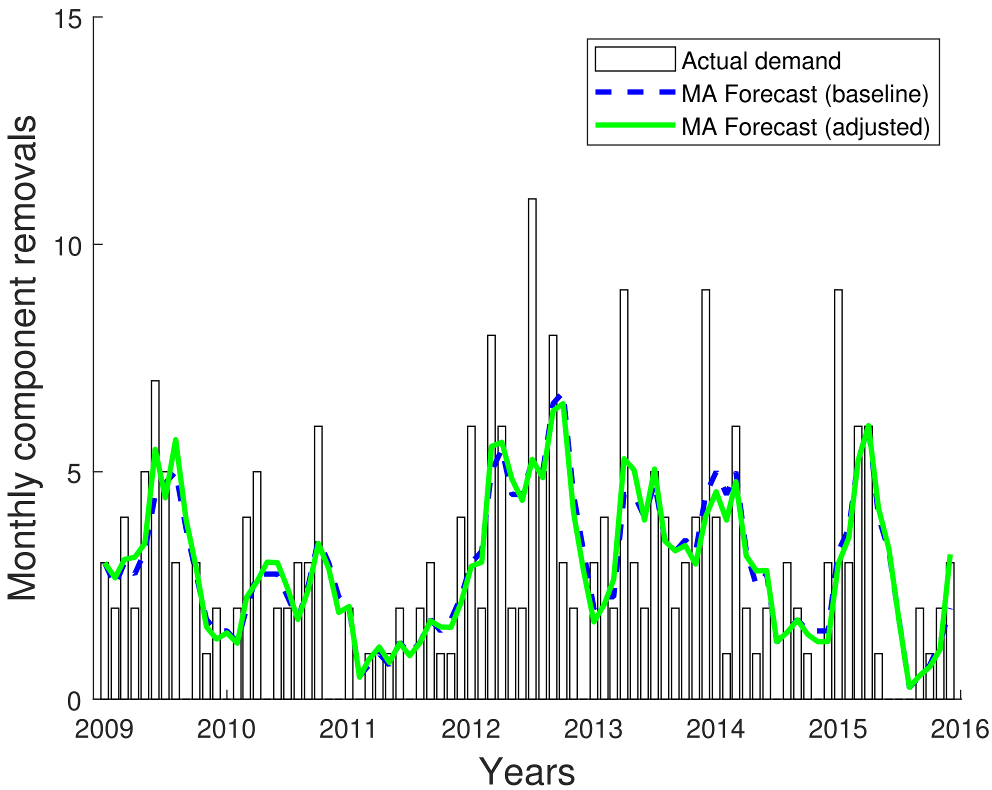

Figure C.11: Actual demand, MA forecast and adjusted forecast for component type 8

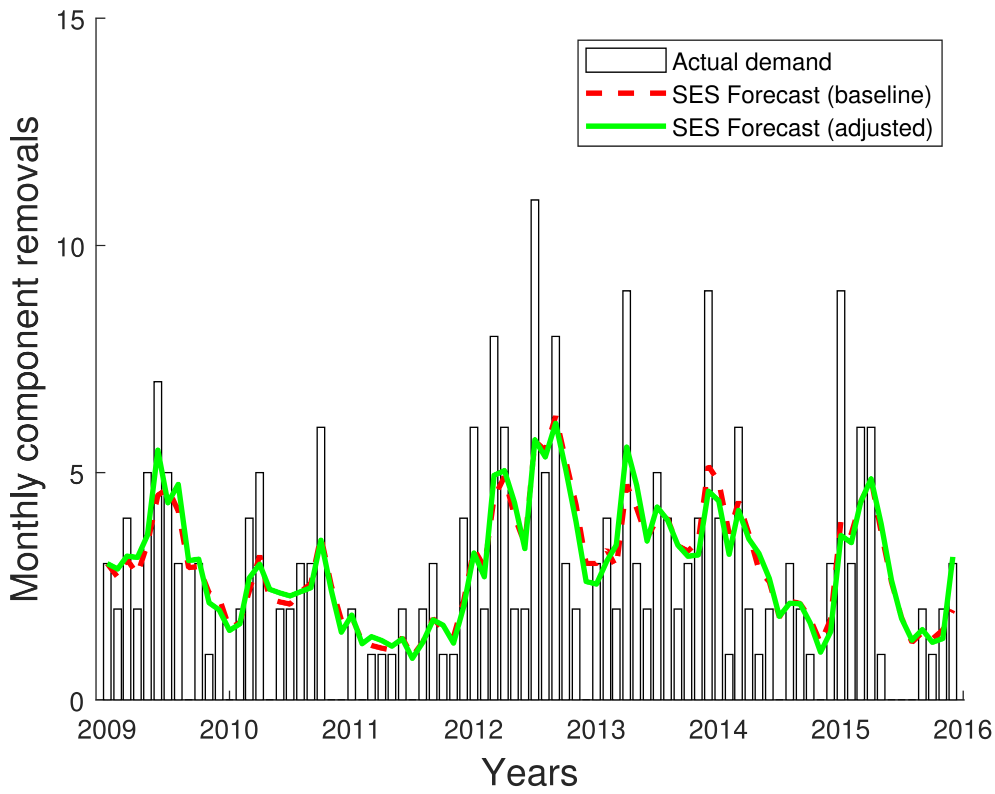

Figure C.12: Actual demand, SES forecast and adjusted forecast for component type 8

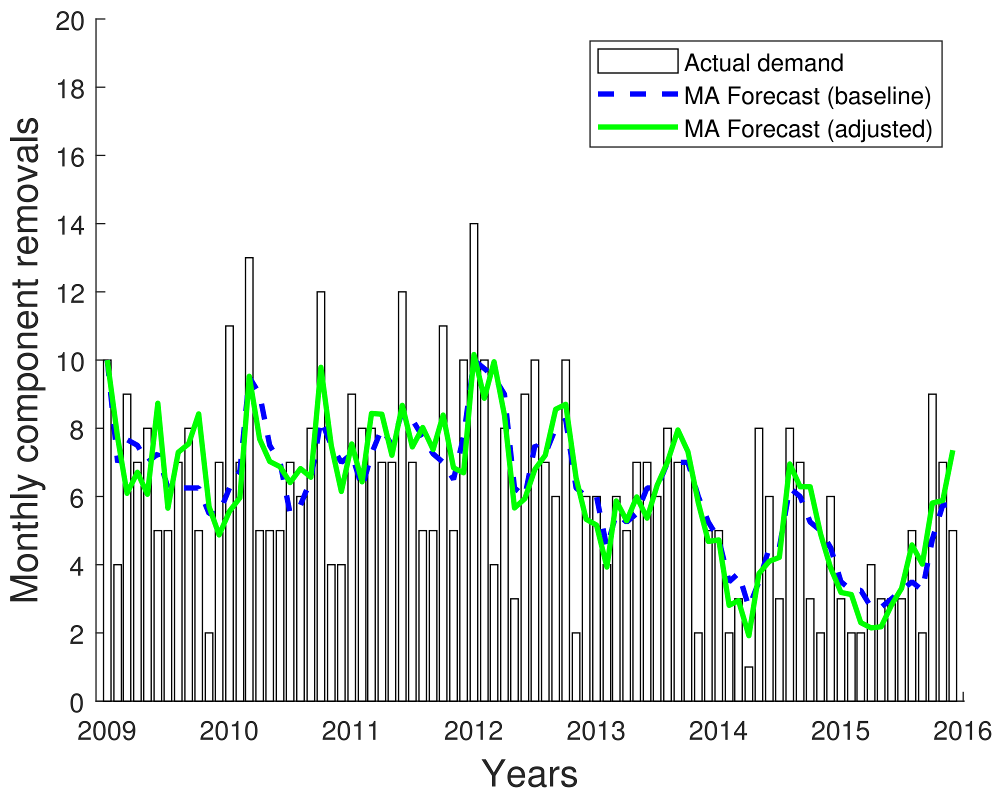

Figure C.13: Actual demand, MA forecast and adjusted forecast for component type 9

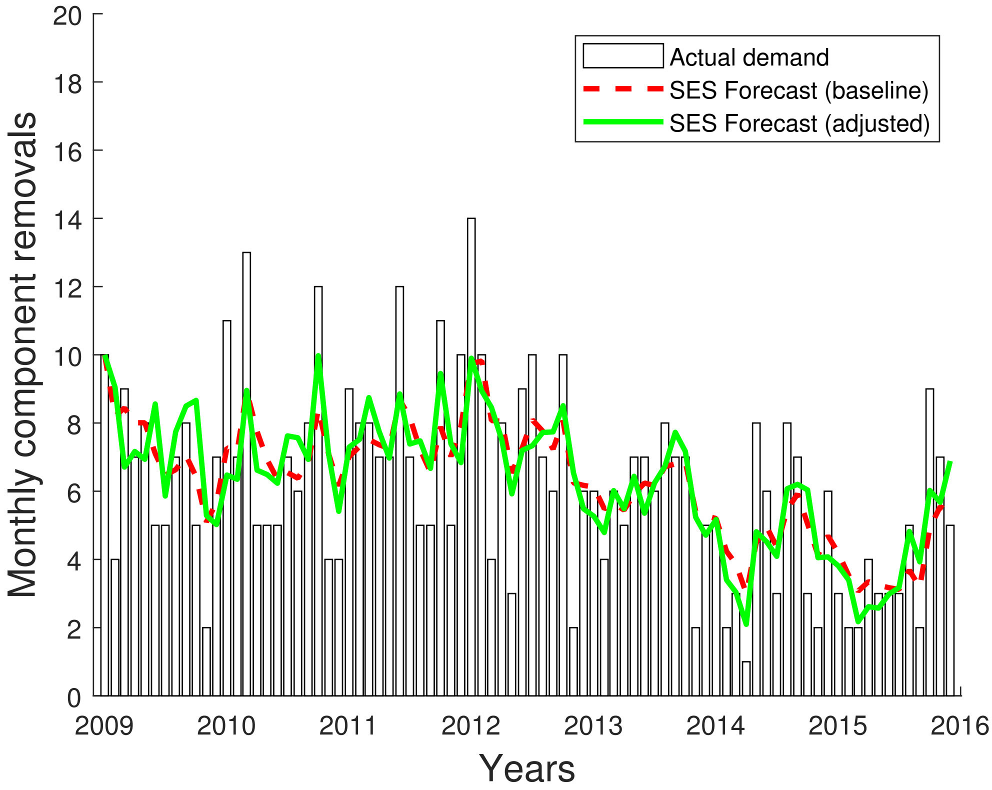

Figure C.14: Actual demand, SES forecast and adjusted forecast for component type 9

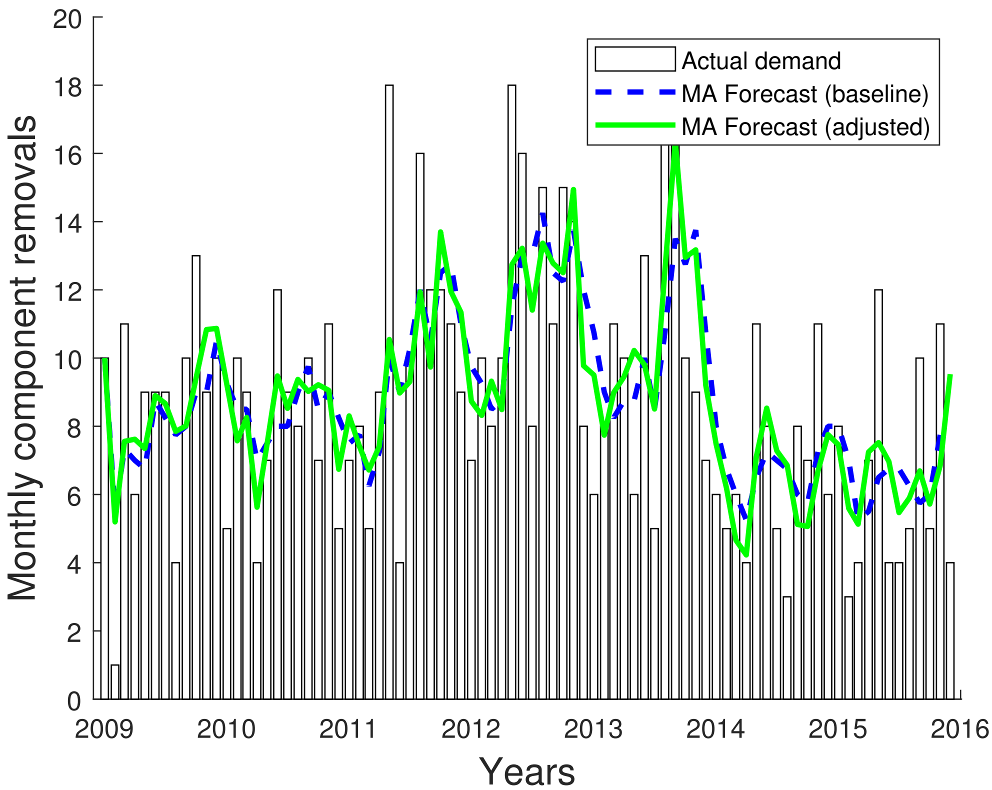

Figure C.15: Actual demand, MA forecast and adjusted forecast for component type 10

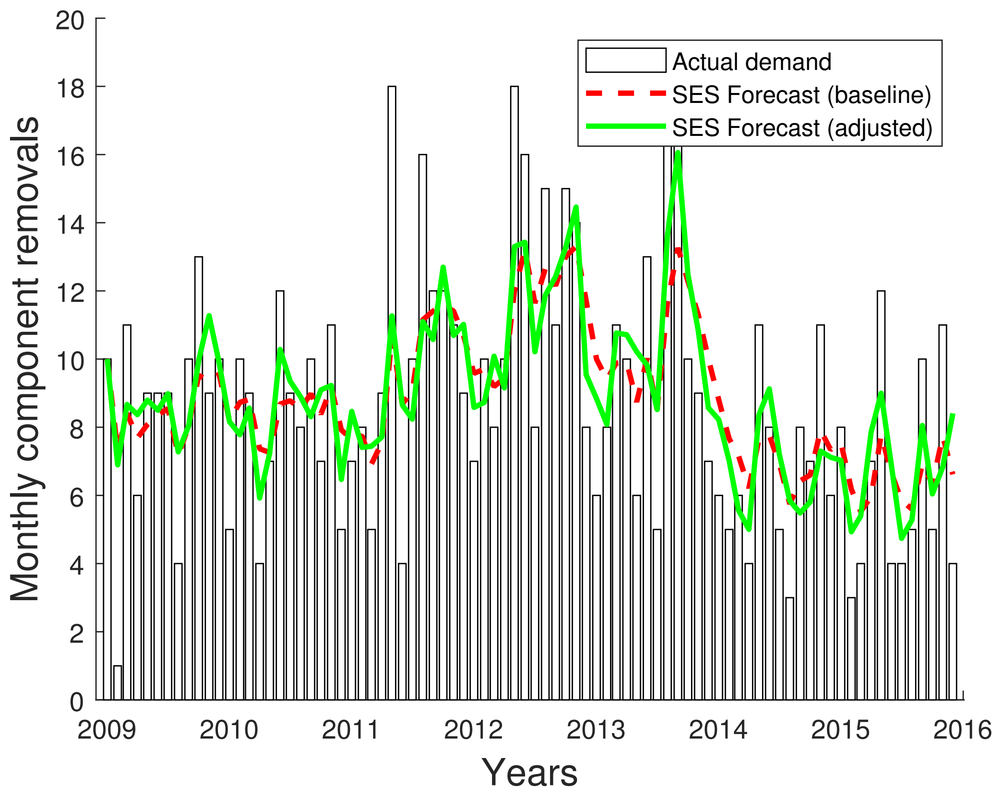

Figure C.16: Actual demand, SES forecast and adjusted forecast for component type 10

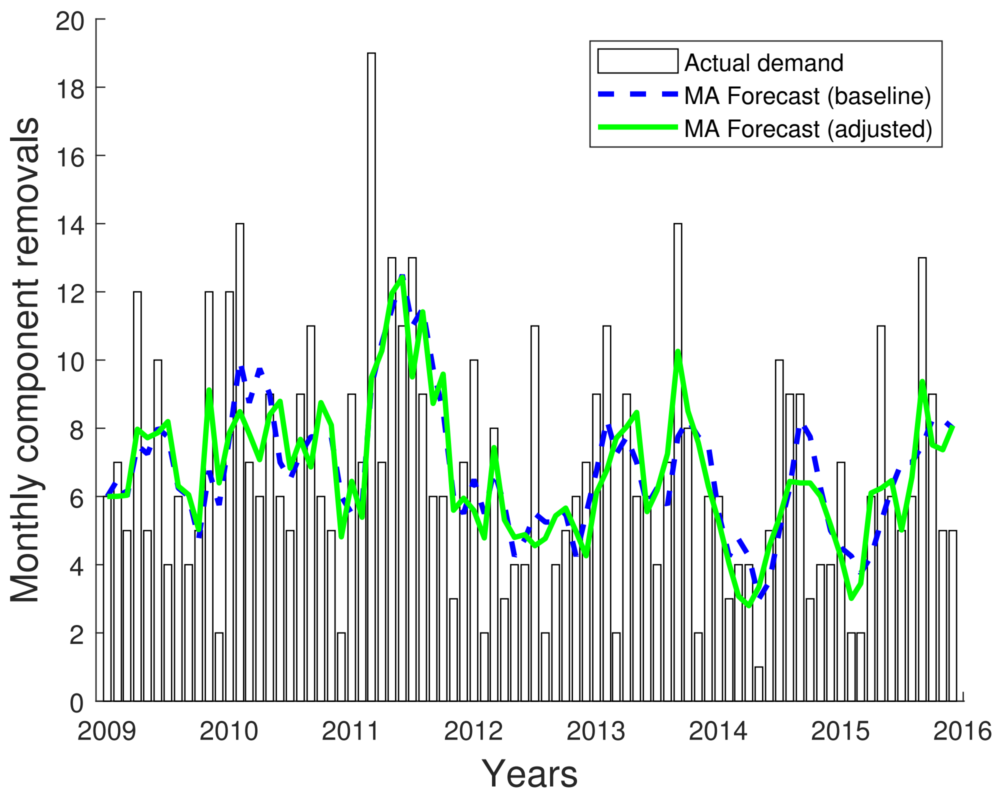

Figure C.17: Actual demand, MA forecast and adjusted forecast for component type 11

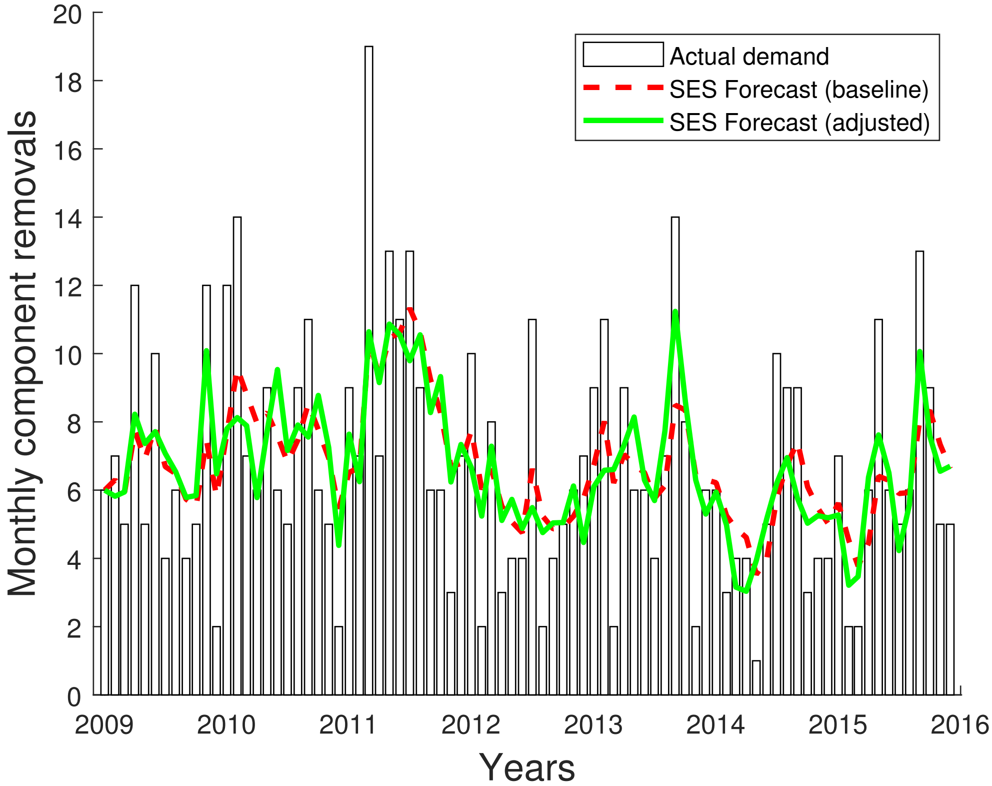

Figure C.18: Actual demand, SES forecast and adjusted forecast for component type 11