

2016

Aircraft Demand Forecasting

Kayla M. Monahan

University of Massachusetts Amherst

Follow this and additional works at: https://scholarworks.umass.edu/masters_theses_2



Part of the [Industrial Engineering Commons](#), and the [Operational Research Commons](#)

Recommended Citation

Monahan, Kayla M., "Aircraft Demand Forecasting" (2016). *Masters Theses*. 329.
https://scholarworks.umass.edu/masters_theses_2/329

This Open Access Thesis is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Masters Theses by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

AIRCRAFT DEMAND FORECASTING

A Thesis Presented

by

KAYLA MONAHAN

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

**MASTER OF SCIENCE IN INDUSTRIAL ENGINEERING AND OPERATIONS
RESEARCH**

February 2016

Mechanical and Industrial Engineering

© Copyright by Kayla Monahan 2016

All Rights Reserved

AIRCRAFT DEMAND FORECASTING

A Thesis Presented

by

KAYLA MONAHAN

Approved as to style and content by:

Ana Muriel, Chair

Hari Balasubramanian, Member

Erin Baker, Member

Sundar Krishnamurty, Department Head
Mechanical and Industrial Engineering

DEDICATION

To my parents, for their endless love, support and encouragement.

ACKNOWLEDGMENTS

Foremost, I'd like to thank my advisor, Ana Muriel, for persevering with me and providing constant support and guidance throughout this process. I'd also like to thank the members of my thesis committee, Hari Balasubramanian and Erin Baker, for their time and expertise to improve my work.

ABSTRACT

AIRCRAFT DEMAND FORECASTING

FEBRUARY 2016

KAYLA MONAHAN

B.S., RENSSELAER POLYTECHNIC INSTITUTE

M.S.I.E.O.R., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Ana Muriel

This thesis aims to forecast aircraft demand in the aerospace and defense industry, specifically aircraft orders and deliveries. Orders are often placed by airline companies with aircraft manufacturers, and then suddenly canceled due to changes in plans. Therefore, at some point during the three-year lead time, the number of orders placed and realized deliveries may be quite different. As a result, orders and deliveries are very difficult to predict and are influenced by many different factors. Among these factors are past trends, macroeconomic indicators as well as aircraft sales measures. These predictor variables were analyzed thoroughly, then used with time series and multiple regression forecasting methods to develop different forecasts for quarterly and annual orders and deliveries. The relative accuracies of forecasts were measured and compared through the use of Theil's U statistic. Finally, a linear program was used to aggregate multiple forecasts to develop an optimal combination of all forecasts. In conclusion, the methods employed in this thesis are quite effective and produce a wholesome aggregate forecast with an error that is generally quite low for a forecasting task as challenging as this one.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	v
ABSTRACT.....	vi
LIST OF TABLES	ix
LIST OF FIGURES	x
CHAPTER	
1. INTRODUCTION	1
1.1 Motivation.....	1
1.2 Objective.....	2
2. LITERATURE REVIEW	3
2.1 Background on Forecasting	3
2.1.1 Time Series Methods	4
2.1.2 Regression Models.....	11
2.1.3 Evaluating Forecast Accuracy	11
2.1.4 Seasonality	13
2.1.5 Combining Forecasts	14
2.2 Forecasting in the Airline Industry	15
3. METHODOLOGY	20
3.1 Description of Variables	20
3.2 Analysis of Input Variables	24
3.2.1 Seasonality	24
3.2.2 Volatility	26
3.2.3 Linear Regression	27
3.2.4 Correlation	31
4. ANALYSIS AND RESULTS.....	36
4.1 Forecasting Methods and Models	36
4.1.1 Naïve No Change Method	37

4.1.2 Holt’s Method	39
4.1.3 Holt-Winters Method	42
4.1.4 Forecasting using Aggregate Annual Data with Seasonal Factors	44
4.1.5 Multiple Regression with Lagged Values	46
4.1.6 ARIMA Forecasting.....	49
4.2 Evaluation of Forecasts with Theil’s U	55
4.3 Aggregation Linear Program	56
4.4 Summary of Model Performance.....	61
5. DISCUSSION AND CONCLUSION.....	63
5.1 Overall Performance of Selected Forecasting Models.....	63
5.2 Limitations	64
APPENDIX: SUPPORTING BASELINE GRAPHS	65
BIBLIOGRAPHY	72

LIST OF TABLES

Table	Page
1. R Squared Values for Input Variables for Orders.....	28
2. R Squared Values for Deliveries.....	29
3. Naive No Change Performance Statistics	39
4. Performance Statistics for Orders Using Holt's Method.....	40
5. Performance Statistics for Deliveries using Holt's Method	42
6. Performance Statistics for Orders and Deliveries using Holt-Winter's Method.....	43
7. Seasonal Factors for Orders and Deliveries.....	44
8. Performance Measures for Seasonal Factor Forecasts.....	45
9. Performance Statistics for Quarterly Deliveries	47
10. Performance Statistics for Quarterly Orders.....	49
11. AIC Values for ARIMA Models for Deliveries	51
12. ARIMA Performance Statistics for Deliveries	52
13. AIC Values for ARIMA Models for Orders	54
14. ARIMA Performance Statistics for Orders.....	54
15. Theil's U Statistics by Forecasting Model for Deliveries.....	56
16. Theil's U Statistics by Forecasting Model for Orders	56
17. Aggregate Linear Program Weights assigned to each Forecast.....	57
18. Performance Measures for Aggregate Linear Program	59
19. Performance Statistics for Simple Average Aggregate Forecasts	60
20. Summary of Performance Statistics for Orders	61
21. Summary of Performance Statistics for Deliveries.....	62

LIST OF FIGURES

Figure	Page
1. Flow Map of the Airline Industry (adapted from Lyneis (2000)).....	23
2. Seasonality of Retirements	25
3. Seasonality of Orders	25
4. Seasonality of Net Orders	26
5. Seasonality of Deliveries	26
6. Quarterly and Annual Volatility of Input Variables	27
7. R Squared Values for Orders	30
8. R Squared Values for Deliveries.....	30
9. Correlation of Input Variables	32
10. Lag Correlation for Orders.....	34
11. Lag Correlation for Deliveries	34
12. Naive No Change Forecast for Orders.....	38
13. Naive No Change forecast for Deliveries	38
14. Holt's Method for Annual Orders	40
15. Holt's Method for Quarterly Orders.....	40
16. Holt's Method for Annual Deliveries	41
17. Holt's Method for Quarterly Deliveries	41
18. Holt-Winter's Method for Quarterly Orders	43
19. Holt-Winter's Method for Quarterly Deliveries.....	43
20. Seasonal Factor Forecast for Deliveries	45
21. Seasonal Factor Forecast for Orders	45
22. Example of a 4 quarter lag	46

23. Lagged Multiple Regression Forecast for Quarterly Deliveries	47
24. Lagged Multiple Regression Forecast for Annual Deliveries.....	47
25. Lagged Multiple Regression for Quarterly Orders	48
26. Lagged Multiple Regression for Annual Orders.....	48
27. Initial ACF plot for Deliveries	50
28. ACF and PACF plots for Differenced Deliveries	51
29. ARIMA Forecast for Quarterly Deliveries	52
30. Initial ACF plot for Orders	53
31. ACF and PACF plots for Differenced Orders.....	53
32. ARIMA Forecast for Quarterly Orders.....	54
33. Linear Program Aggregate Forecast for Orders.....	58
34. Linear Program Aggregate Forecast for Deliveries	58
35. Simple Average Aggregate Forecast for Orders.....	60
36. Simple Average Aggregate Forecast for Deliveries	60

CHAPTER 1

INTRODUCTION

1.1 Motivation

Physicist Neils Bohr once said, "Prediction is very difficult, especially if it's about the future." Forecasting is indisputably an intrinsically difficult task. It has been referred to as "one of the 10 grand challenges of modern science" (Cheng et. al. 2015). However, if done correctly, it can have an immense impact. Forecasts are essential for making business decisions such as knowing how much to produce, the resources and capacity required, which products to develop, and the optimal time to develop them. A 10% improvement in forecasting accuracy can impact revenue gain by up to 4% (Yu 2012). For many large companies, even a 1% increase means an increase in millions of dollars of revenue, further stressing the importance of an accurate forecast.

Forecasting in the Aerospace and Defense industry is a unique and interesting problem to consider. The industry is extremely complex due to its competitive and dynamic nature. This is especially true after the passage of the 1978 Airline Deregulation Act, which removed government control over fares, routes and market entry, thus allowing the airline industry to succumb to a state of aggressive competition among carriers (Prousaloglou and Koppelman 1995). Additionally, there are many other factors that influence aircraft demand. To put it simply, economic conditions generate travel demands which creates growth in demand for aircraft and components from large companies such as Boeing, Airbus and Pratt and Whitney. However, during these long lead times, airline companies may suddenly cancel aircraft orders due to changes in plans. This makes it evident that many factors contribute to aircraft demand, and must be considered and evaluated when

forecasting. The dynamic nature of this problem makes it an exciting one to tackle, not to mention one whose solution is in high demand.

1.2 Objective

The goal of this thesis is to forecast aircraft demand in the aerospace and defense industry, specifically orders and deliveries. Orders are often placed by airline companies with aircraft suppliers and then suddenly canceled due to changes in plans. Therefore, at some point during the three year lead time, the number of orders and deliveries becomes much different. Thus, orders and deliveries are very difficult to predict, and much analysis must be done in order to forecast these variables adequately. First, it is essential to understand the background of research as well as the potential economic drivers of demand that have been previously indicated. Various time series forecasting techniques will be used, as well as multiple regression with different combinations of relevant and important economic indicators such as Gross Domestic Product (GDP), Revenue Passenger Mile (RPM), Load Factor, Fuel Price, Interest Rate, etc. The resulting forecasts will be analyzed and aggregated through the use of a linear program, producing an optimal combination of forecasts.

CHAPTER 2

LITERATURE REVIEW

2.1 Background on Forecasting

In general, forecasting techniques can be broken down into two different categories: qualitative or quantitative. Quantitative forecasting techniques consist of either time series analysis or causal models and rely heavily on historical data. Holt's Method, moving averages, and trend projection are just a few examples of time series techniques. Causal methods consist of many different regression models. To contrast, qualitative forecasting techniques are much less methodical and rely on judgement. Some examples are the Delphi Method and sales force composites.

There are a few well-known facts about forecasting that is important to always remember. First, forecasts, in general, are always wrong. No forecast is perfectly accurate; therefore the goal is to achieve a forecast with minimal error. Second, long term forecasts are usually less accurate than short term forecasts. Third, aggregate forecasts, where data is drawn from various sources are generally more accurate than disaggregate forecasts.

Forecasting models consist of two components: a systematic component and a random component. The systematic component is what we are trying to predict, and often exhibits trends, cycles or seasonality. Trends are any steady growth or decline in the forecast. Seasonality is defined as up and down swings exhibiting a pattern in a short or intermediate time, generally a year. Cycles are up or down swings over a long time. With any forecast, there is always a random component that cannot be explained, but the goal is to minimize this element as much as possible. The basic approach to forecasting is to

understand the objective, and then identify major factors that influence the variable in question. It is important to choose the appropriate forecasting technique and finally, evaluate performance and measure error.

In this section, an overview of different forecasting methods and models will be presented, as well as methods to measure forecast accuracy.

2.1.1 Time Series Methods

Methods for forecasting originated in the 1950s to 1960s and typically did not address the random component of a time series. The main idea was to develop methods for predicting the variable in question from its past data. Some of the most simple univariate forecasting methods are the naïve no-change method, naïve change and naïve seasonal change method. The naïve no-change method simply develops a forecast for the given period (\hat{Y}_{t+1}) that is the actual value from the previous period, (Y_t). The naïve change method develops a forecast for a given period (\hat{Y}_{t+1}) as the actual value from the previous period (Y_t) plus an extra component which is defined as the difference between that previous period (Y_t) and the period before the previous period, (Y_{t-1}). The naïve seasonal change method develops a forecast for the given period (\hat{Y}_{t+1}) as the actual value from the previous period (Y_t). plus an extra component defined as the difference between the value occurring one complete season before the forecast (Y_{t+1-s}) and one season before the previous period, (Y_{t-s}), where s is the seasonal component. All three formulas are presented below (Enders 2004).

$$\text{naïve no - change: } \hat{Y}_{t+1} = Y_t$$

$$\text{naïve change: } \hat{Y}_{t+1} = Y_t + (Y_t - Y_{t-1})$$

$$\text{naïve seasonal change: } \hat{Y}_{t+1} = Y_t + (Y_{t+1-s} - Y_{t-s})$$

Another method commonly used to forecast is the simple moving-average method. In contrast to the naïve method which typically is successful when the observations are relatively constant over time, the moving average method can be used to smooth data in order to see the trend. The forecast is calculated as follows:

$$\hat{Y}_{t+1} = \frac{Y_t + Y_{t-1} + Y_{t-2} + \dots + Y_{t-(k-1)}}{k}$$

where k is the number of values in the average. Typically for quarterly data, the k value would be 4, for monthly data the k value would be 12. Typically, the larger the value of k , the smoother is the series.

Brown, Holt and Winters pioneered many of the methods that use exponential smoothing (De Gooijer and Hyndman 2006). Exponential smoothing is a form of filtering or averaging. In contrast to the simple moving average method, exponential smoothing puts greater weight on more recent observations. Exponential smoothing takes into account three parameters of a data series: level, trend and seasonality. The single exponential smoothing method accommodates level only, whereas Holt's accommodates level and trend, and Holt-Winters' accommodates level, trend and seasonality. Choosing a method to use is somewhat ad hoc, one should simply infer based on the appearance of the appropriate parameters in the data series. The robustness of the exponential smoothing method has been commented on by many different researchers. It was generally shown that exponential smoothing is optimal for many data generating processes (De Gooijer and Hyndman 2006). Chatfeild et al. (2001) showed that simple exponential smoothing is highly applicable and optimal for understanding many different types of data generating processes. Additionally, Hyndman (2001) showed how simple exponential smoothing

outperformed first order ARIMA models. ARIMA Models, which will be discussed further later, stand for Auto Regressive Integrated Moving Average Models and attempt to describe autocorrelations in data.

In 1957, Charles C. Holt expanded on the previous work on simple exponential smoothing by Robert Goodell Brown. Holt's method involves a forecast equation and two smoothing equations, one for the level and one for the trend (Otexts 2015). The subsequent equations are presented below.

Holt's Method

$$\textbf{Forecast Equation: } \hat{y}_{t+h|t} = l_t + hb_t$$

$$\textbf{Level Equation: } l_t = \alpha y_t + (1 - \alpha)(l_{t-1} + b_{t-1})$$

$$\textbf{Trend Equation: } b_t = \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1}$$

where l_t is an estimate of the level of the series at time t

b_t denotes an estimate of the trend of the series at time t

h denotes the number of forecasted steps ahead

y_t denotes the actual value at time t

α is the smoothing parameter for the level, $0 \leq \alpha \leq 1$

β is the smoothing parameter for the trend, $0 \leq \beta \leq 1$

In 1960, Holt and Winters expanded on Holt's Method in order to incorporate seasonality, and thus, the Holt-Winters Seasonal Method was born. There are two different forms of this method, the additive and the multiplicative method. The difference between these two variations is in the seasonal component. Generally, the additive method is used

when the seasonal change is constant throughout the series. Conversely, the multiplicative method is used when the seasonal change is proportional to the level of the series. The additive and multiplicative method equations are presented below (Otexts 2015).

Holt – Winters Additive Method

Forecast Equation: $\hat{y}_{t+h|t} = l_t + hb_t + s_{t-m}$

Level Equation: $l_t = \alpha(y_t - s_{t-m}) + (1 - \alpha)(l_{t-1} + b_{t-1})$

Trend Equation: $b_t = \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1}$

Seasonality Equation: $s_t = \gamma(y_t - l_{t-1} - b_{t-1}) + (1 - \gamma) s_{t-m}$

where l_t is an estimate of the level of the series at time t

b_t denotes an estimate of the trend of the series at time t

s_t denotes an estimate of the seasonality of the series at time t

m denotes the period of the seasonality

h denotes the number of forecasted steps ahead

y_t denotes the actual value at time t

α is the smoothing parameter for the level, $0 \leq \alpha \leq 1$

β is the smoothing parameter for the trend, $0 \leq \beta \leq 1$

γ is the smoothing parameter for the trend, $0 \leq \gamma \leq 1$

Holt – Winters Multiplicative Method

Forecast Equation: $\hat{y}_{t+h|t} = (l_t + hb_t)s_{t-m}$

Level Equation: $l_t = \alpha\left(\frac{y_t}{s_{t-m}}\right) + (1 - \alpha)(l_{t-1} + b_{t-1})$

$$\textbf{Trend Equation: } b_t = \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1}$$

$$\textbf{Seasonality Equation: } s_t = \gamma \left(\frac{y_t}{(l_{t-1} + b_{t-1})} \right) + (1 - \gamma) s_{t-m}$$

where l_t is an estimate of the level of the series at time t

b_t denotes an estimate of the trend of the series at time t

s_t denotes an estimate of the seasonality of the series at time t

m denotes the period of the seasonality

α is the smoothing parameter for the level, $0 \leq \alpha \leq 1$

β is the smoothing parameter for the trend, $0 \leq \beta \leq 1$

γ is the smoothing parameter for the trend, $0 \leq \gamma \leq 1$

Researchers such as Slutsky, Walker, Whittle, Yaglom and Yule first identified the concepts of Autoregressive (AR) and Moving Average (MA). It wasn't until the publication of *Time Series Analysis: Forecasting and Control* by Box and Jenkins (1970), that a systematic process for time series identification, estimation and verification was formulated. This book tremendously impacted time series analysis and forecasting and popularized the use of Autoregressive Integrated Moving Average (ARIMA) models (De Gooijer and Hyndman 2006). Autoregressive (AR) models are used to model many different stationary processes. Among them are the ARMA and ARIMA models. The ARIMA model is generally represented as ARIMA (p,d,q) where p represents the order of the autoregressive model, d is the degree of the differencing needed for stationarity, and q is the order of the moving average model. A general rule of thumb is to keep the sum of p and q less than or equal to 2, to prevent overfitting of the model. Deciding on the value for

each of these parameters p and q depends on the autocorrelation and partial autocorrelation of the series. The differencing parameter d is defined by determining what order of difference is needed for the series to be stationary. Stationarity ultimately means that a series has a constant mean, variance and autocorrelation structure over time (De Gooijer and Hyndman 2006). The concept of differencing can be shown simply below, where $d=0$ represents no differencing, $d=1$ represents first differencing, and so forth.

$$\text{when } d = 0: y_t = y_t$$

$$\text{when } d = 1: y_t = y_t - y_{t-1}$$

$$\text{when } d = 2: y_t = (y_t - y_{t-1}) - (y_{t-1} - y_{t-2})$$

The general forecasting equation for an ARIMA model is presented below.

$$\hat{y}_t = \mu + \varphi_1 y_{t-1} + \dots + \varphi_p y_{t-p} - \theta_1 e_{t-1} - \dots - \theta_q e_{t-q}$$

where φ represents the autoregressive parameters, and θ represents the moving average parameters. When an ARIMA model simply contains the p parameter, it contains the AR component. Conversely, when an ARIMA model simply contains the q parameter, it contains the MA component only. Finally, when an ARIMA model just contains the d parameter, it is simply a random walk model.

The ARMA model lacks the “I” part since it consists of a linear relationship between lagged variables without the need for differencing. The ARIMA model focuses on reducing first-order non-stationarity through differencing (Cheng et. al. 2015). Overall, the ARMA model establishes a lagged relationship between the dependent variable Y_t and the independent variable Y_{t-1} . The estimation formula is presented below where β_0 and β_1 are estimated by the method of least squares and ε_t is a random disturbance with zero expected value and constant variance (Enders 2004).

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \varepsilon_t$$

Once the model is estimated where estimates for β_0 and β_1 are b_0 and b_1 , the forecast model is as follows:

$$\hat{Y}_{t+1} = b_0 + b_1 Y_t$$

There are many different variations of AR models, catering to different aspects of the data one is working with as well as the type of forecast that is needed. However, this means model selection errors are quite common, since selecting the specific type of AR model is mostly based on the researcher's interpretation of the look of the data, and subjective decisions are often made in this aspect. Therefore, simple exponential smoothing methods are often times preferred to AR methods for forecasting time series data (Hyndman 2001).

However, there is a way to help decide between competing ARIMA models through the Akaike Information Criterion (AIC). The AIC is ultimately a measure of the quality of fit for the given model on a data series, and will penalize a model for using a greater number of parameters. Thus, this rewards goodness of fit of a model, but discourages overfitting. The AIC is calculated as follows:

$$AIC = 2k - 2\ln(L)$$

where k is the number of parameters in the model and L is simply the Residual Sum of Squares divided by n , $\frac{RSS}{n}$, where n represents the number of observations. The winning model is the one with the lowest AIC value. This is an extremely useful metric considering initial model selection is somewhat ad hoc.

Additionally, there are many other approaches for forecasting such as qualitative methods and simulation methods. Qualitative methods are based on human judgement and

relies very little on historic data. An example of this method is the Delphi method, which is most commonly used in business environments. This method involves a team of experts that reach a consensus together, allowing for a complete evaluation of every member's argument. Additionally, simulation methods are often used for forecasting and involve imitating consumer choices to forecast the most likely scenario.

2.1.2 Regression Models

The first forecasting model that is of interest here is linear trend regression. This assumes a contemporaneous relationship between the dependent variable Y_t and the independent variable t . α_0 and α_1 are parameter estimates and are estimated by the method of least squares and ε_t is a random disturbance with zero expected value and constant variance. Parameter estimates for α_0 and α_1 are a_0 and a_1 . The estimation and forecast formula are presented below (Enders 2004).

$$Y_t = \alpha_0 + \alpha_1 t + \varepsilon_t$$

$$\hat{Y}_{t+1} = a_0 + a_1(t + 1)$$

There are many different regression models that capture elements such as causal relationships, trend, and indicator variables. To keep this review concise, all of the different regression based forecasting models will not be presented but it is important to note them.

2.1.3 Evaluating Forecast Accuracy

There are multiple ways to assess the accuracy of a forecast. Each technique involves comparing the forecasted value with the realized value of a variable of interest. The amount by which the forecast differs from the actual value Y_t is the forecast error e_t ,

where $e_t = Y_t - \hat{Y}_t$. Four simple and commonly used measures of forecast accuracy are presented below.

First, the Mean Absolute Error (MAE), also known as the Mean Absolute Deviation (MAD) is as follows where n is the total number of observations for the period.

$$MAE = \frac{\sum_{t=1}^n |e_t|}{n}$$

Another method used often is the Mean Absolute Percent Error (MAPE). The MAPE is a modification of MAD. MAPE looks at the size of the absolute value of the error relative to the actual value itself. MAPE is presented below.

$$MAPE = \frac{\sum_{t=1}^n \frac{|e_t|}{Y_t}}{n} \times 100$$

The third method for measuring error is Mean Square Error (MSE). This squares the individual errors as follows

$$MSE = \frac{\sum_{t=1}^n e_t^2}{n}$$

Finally, the Root Mean Square Error (RMSE) is simply the square root of the MSE. Each of these four measures can be used to determine forecast accuracy. MAPE is useful as it is unit free. When using MSE or RMSE, having one or two large errors may magnify the overall measure of error. Therefore, using MAD can avoid this. When all of the errors are of the same magnitude, RMSE and MSE are most useful (Newbold and Bos 1993).

Another, more unfamiliar measure of forecasting performance is Theil's U developed by Henri Theil. Theil's U measures the worth of a forecasting method that is deemed to be more advanced than the naïve no-change method. His "U statistic" is the ratio

of the RMSE of the more sophisticated method being analyzed to the RMSE of the naïve no-change method. It is represented below as

$$\textit{Theil's } U = \frac{\textit{RMSE of advanced approach}}{\textit{RMSE of naïve no – change}}$$

If $U > 1$, then the advanced approach has no value because it cannot perform as well as the naïve no change basic method. However, if $U < 1$, the advanced approach has more merit. The closer U gets to 0, the better the approach in question is. (Newbold and Bos 1993)

2.1.4 Seasonality

Seasonality is defined as “The estimated seasonal is that part of the series which, when extrapolated, repeats itself over any one-year time period and averages out to zero over such a time period” (Harvey 1990). The goal of any forecaster is to eliminate noise from the data. Therefore, by identifying and stripping a series of its seasonality, an analyst can identify period to period changes that are due to causes other than seasonality, such as trend. A deseasonalized series does just that; removes the seasonality to identify the general trend of the data series.

The most common method for deseasonalizing a series is the Simple Moving Average Method. This method originated in the 1920s and is the basis for the majority of methods to de-seasonalize and smooth out data (Otexts 2015). The general idea is to take a centered moving average of 4 points at a time for quarterly data, to remove the seasonality and smooth out the series. The centered data point is the one that assumes that value. In this case, the center of 4 data points is 2.5, so the average of the 2.5 and 3.5 values are used to find the value for time value 3. The formulas are presented below.

$$Y^*_{2.5} = (Y_1 + Y_2 + Y_3 + Y_4)/4$$

$$Y^*_{3.5} = (Y_2 + Y_3 + Y_4 + Y_5)/4$$

$$Y^{**}_3 = (Y^*_{2.5} + Y^*_{3.5})/2$$

From here, the seasonal indices can be found by first identifying the ratio between the original data point to the new deseasonalized data point. Then from there, the corresponding seasonal factor for each of the four quarters is its respective average over the entire time series.

This is an important step in analyzing a data series as it allows reliable comparison of observations at different points in time. Additionally, it allows for other behavior patterns and trends in the data to be seen clearly.

There are some models that incorporate a seasonal parameter in the forecasting process itself. Among them are regression with seasonality, Holt-Winter's Method, and an ARIMA model. It was found that the latter two models are quite robust when seasonality is apparent (De Gooijer and Hyndman 2006).

2.1.5 Combining Forecasts

Combining different forecasts obtained from different but valid forecasting techniques is a common practice for many forecasters. Early researchers such as Bates, Granger, Newbold, Winkler and Makridakis presented significant evidence toward the effectiveness of combining forecasts (De Gooijer and Hyndman 2006).

There are many different methods that have been used to aggregate different forecasts. First, the simple average is the most simple and widely used method, however it has been criticized for lacking the ability to utilize information on the accuracy of the

forecasts. This method essentially assigns even weights to the forecasts regardless of their given accuracy (Clemen 1989). Another method commonly used simply assigns weights to the forecast using Ordinary Least Squares. These weights are determined based on the overall accuracy of the forecasts (Granger and Ramanathan 1984). More complicated methods for combining forecasts exist and among them are Bayesian shrinkage techniques, and methods which update and change the weights, rather than using fixed weights. However, Miller et al. (1992) suggest that simpler methods for aggregating forecasts are generally more successful than the more complex methods since nonstationarity can occur. Ultimately though, by combining different competing forecasts one can obtain a vastly superior forecast (Fang 2003). It is also perceivably less risky in practice to use a combined forecast rather than selecting a single forecast (Hibon and Evgeniou 2009).

2.2 Forecasting in the Airline Industry

For lack of literature specifically in methods for forecasting aircraft demand (as top competitors such as Airbus and Boeing keep that very private) we will focus on relevant previous research done on forecasting commercial airline demand.

Beginning in the 1950s, the gravity model was widely used for forecasting demand between pairs of cities or airports. The gravity model assumes that trips produced at an origin and attracted to a destination are directly proportional to the total trip productions at the origin and the total attractions at the destination. Essentially, that the sum total of trip production from areas is equivalent to the sum total of trip attractions for those same areas. (Harvey 1951). There is a friction factor involved that represents the reluctance of people to make certain trips. For example, as travel times increase, travelers are less likely to make

the trip. The model incorporates socioeconomic factors for individual trips. This literature review primarily focuses on time series and causal models, and this model reinforces the importance of using economic factors for forecasting in a causal manner (Verleger 1962).

A new perspective for forecasting air transportation demand emerged in the 1960s. The goal was to predict Revenue Passenger Miles (RPM). RPM is a product of the number of revenue-paying passengers and the distance traveled. In Bartlett (1965) 26 regression equations were used with different explanatory variable combinations. The resulting best fit regression equation was one that accounted for a measure of leisure time activity. However, even accounting for the monetary value of time, this model only produced an R squared value of 0.58. Another forecasting model to predict RPM was developed by Vitek and Taneja (1975). This model used time series data along with price of flight, income and inflation as explanatory variables. Linear regression was performed and forecasts were provided up to 1990. Results showed that price is the most stable and significant determinant of demand. Income and rate of inflation were significant as well but much more variable. An R squared value of 0.96 was achieved. Brown and Watkins (1968) utilize time series techniques as well as multiple regression with predictors such as income, fare, travel time and number of stops to predict air traffic demand, or RPM.

A linear regression based model developed by Jacobson (1970) predicted trips generated at an airport based on average income and airfare in the US. The model used 18 years of data from airports in Virginia. The resulting model had an R squared value of 0.82. Another model that focused on one airport was that of Haney (1975), which forecasted total annual traffic for the St. Louis airport. Socioeconomic variables were chosen to represent the area around the airport. The resulting regression model had an R squared

value of 0.99. A third paper focusing on predicting the demand at a single airport was that of Thomet and Sultan (1979). The focus here was to forecast the number of passengers originating and terminating at an airport in Saudi Arabia. The input variables were related to crude oil and petroleum, as most travel to and from Riyadh International Airport was oil-related. The resulting forecast achieved an R squared value of 0.99. A fourth model based around a single airport was developed by (Mellman et.al. 1980). The focus here was Boston Logan International Airport. The authors employed multiple regression techniques and identified key predictors such as change in economic conditions of Boston, load factor and fuel prices. Certain forecasting scenarios were explored and the results concluded that air passenger volumes are likely to increase at about five percent per year, contingent upon changes in regional income and fares.

Karlaftis et. al. (1996) developed a model to forecast demand at two major airports: Miami and Frankfurt. The authors used simple time series demand forecasting models to achieve R squared values of 0.72 and 0.94. Additionally, the authors cautioned against using too many predictor variables, as it has not been proven to improve the quality of forecast. In turn, it may be detrimental to the forecast, so predictor variables must be chosen wisely. Lyneis (2000) brings up similar insights but contributes the idea of lagging certain input variables, as that can be a beneficial contributor to explaining demand.

In a paper by Littlewood (1972) in collaboration with Scandinavian Airlines, some simple forecasting models for flight bookings were proposed. The models used here addressed topics such as removing outliers corresponding to non-recurrent events in time series data; in this case the Gulf War. Duncanson (1974), while working at Scandinavian Airlines, also developed time series models to forecast passenger traffic. Simplicity was

emphasized in these models, which incorporated seasonal analysis and exponential smoothing methods to forecast passenger traffic.

In his book titled *Airline Traffic Forecasting*, Taneja (1978) argues that regression models are the most popular and most successful methods for forecasting airline traffic. He primarily presents models that forecast total airline traffic regionally, nationally and internationally.

Oberhausen and Koppelman (1982) produce short term forecasts for inter-city air travel demand. ARIMA methods are used on time series data as well as a bivariate time series model incorporating air fare as an explanatory variable. In this case, the bivariate model did not produce a significantly better forecast.

In a paper by Carson et. al. (2011) forecasting aggregate demand for US commercial air travel is explored using exogenous macroeconomic indicators as the independent variables and a ratio of the number of passengers originating at an airport and the population served by that airport as the dependent variable. The economic indicators used were population, income and energy prices.

The nature of the aircraft industry is extremely cyclical and is in part due to how heavily it is influenced by business cycles. It is thought that this cyclic behavior began after the airline market was deregulated in 1978. In Liehr et. al. (2001) an attempt is made to understand the underlying business cycles that drive the airline industry. The period of the cycle is about seven to ten years and seems to be quite sensitive to fluctuations in GDP. The points brought up in this paper highlight key variables that drive the aircraft industry. Among those variables are GDP, load Factor, revenue, and RPM.

An important perspective to consider is the market outlook forecasts for Boeing and Airbus. Since these companies are the most concerned with predicting aircraft demand, their forecasts provide a benchmark to consider for comparison. It is also of interest to contrast the variables they consider with the variables used in this thesis. The forecasting methods used by Boeing and Airbus are not made public. However, we can infer what variables they may have used in their analysis from what they identify as drivers in their market outlook presentations. Airbus (2015) highlights Worldwide GDP, RPM, and Oil Price as important indicators for demand growth in the airline industry. Congruently, Boeing (2015) identifies RPM, Load Factor, and GDP as strong drivers of aircraft demand. Both papers strongly stress the impact macroeconomic indicators have on driving the demand upwards in the coming years.

To conclude this review of the existing literature, it is important to note that many models keep their analysis simple by using either time series or regression methods. Congruently, the main methods employed in this thesis will be time series and multiple regression techniques using time series data and various economic indicators to forecast aircraft demand. Previous research suggests that air traffic demand is driven by the same macroeconomic factors as aircraft demand, and is sufficient to support the analysis in this paper.

CHAPTER 3

METHODOLOGY

In this chapter, the relative input variables will be discussed in greater detail, as well as their anticipated impact on the multiple regression forecast. Additionally, this chapter will explore greater analysis of the input variables will be explored in terms of seasonality, volatility, correlation with one another, and ultimately correlation with orders and deliveries. This analysis will aid in understanding the underlying relationship between the variables to create a more precise forecast. Again, the overall goal of this thesis is to forecast aircraft orders and deliveries, and those subsequent forecasting methods and models will be presented further in Chapter 4.

3.1 Description of Variables

The following list of variables were identified by top competitors in the aerospace industry, such as Boeing (2014) and Airbus (2015) as potential drivers of demand. The variables can be grouped into two categories - global macroeconomic indicators and aircraft sales figures - are listed and defined below. All data is for the time period 1995 to 2013, and is in quarterly and annual increments.

1. Global Macroeconomic Indicators:

GDP-Worldwide: Gross Domestic Product - The monetary value of all the finished goods and services (In 2005 billion).

GDP Growth: Year over year Percent change

Rate of Inflation Worldwide: Percentage; The rate at which the general level of prices for goods and services is rising

Long Term Interest Rate-Worldwide: Average of daily rates, measured as a percentage

Long Term Interest Rate-US: Average of daily rates, measured as a percentage

Jet Fuel Prices: Price per gallon

Crude Oil Prices: West Texas Intermediate Price per barrel

2. Aircraft Sales Figures:

Aircraft Orders: Number of aircraft ordered

Aircraft Deliveries: Number of aircraft actually delivered

Aircraft Order Cancellations: Number of aircraft cancelled

Aircraft Net Orders: Orders minus Cancellations

Installed Base-Active: Number of aircraft in active use

Retirements: Number of aircraft retired

Revenue Passenger Mile (RPM): In Billions, measures of traffic for airline flights; product of the number of revenue-paying passengers aboard the vehicle and the distance traveled

RPM Growth: Year over year Percent change

Available Seat Mile (ASM): In Billions, measure of a flight's revenue-generating abilities based upon traffic; product of number of seats available and the number of miles flown

ASM Growth: Year over year Percent change

Load Factor: Percentage (RPM/ASM)

Operating Revenue: In millions, revenue worldwide

Operating Profits: In millions, profits worldwide

Net Profits: In millions, net profits worldwide

It is expected that as GDP and GDP growth increases, the number of orders and deliveries will increase in kind, as the national wealth increased. Consequently, it is expected that as the fuel price, oil price, rate of inflation and interest rate (in both the U.S. and worldwide) increases, the number of orders and deliveries will decrease due to the increased financial burden.

Aircraft orders and deliveries are linked through aircraft order cancellations. The nature of the aircraft sales industry is such that aircraft are ordered and possibly cancelled during the approximately three year lead time before delivery. The lead time is not concrete, and may take more or less time for an order. Therefore, we cannot simply subtract the number of cancellations from orders three years ago to obtain the number of deliveries in that year. This complicates the problem further, however we can hypothesize that an increase in the number of cancellations will cause a decrease in the number of deliveries.

Intuitively, we can estimate that an increase in aircraft net orders will cause an increase in aircraft deliveries, as more orders are expected to cause more deliveries, with a time lag. It also is expected that a decrease in installed base will cause an increase in orders and deliveries, as there will be less aircraft in the total fleet. Similarly, it is expected that an increase in the number of aircraft retirements will cause an increase in orders and deliveries, as newer aircraft may be desired. Next, it is anticipated that as RPM, ASM and load factor increase, the number of orders and deliveries will increase because this indicates that the company is increasing in revenue and may need to acquire more aircraft to meet booming demands. Similarly, as the measures for profits, revenue and net profits increase, we can expect aircraft orders and deliveries to increase as well.

We can anticipate how each of these variables will in individually affect orders and deliveries, but how do these variables interact with each other? In Lyneis (2000) the airline industry was represented using an interaction map. As it applies here, the flow map presented in Lyneis (2000) has been amended to represent the anticipated interaction between variables in this analysis. Here, demand by passengers for travel is driven by economic conditions. That demand then determines an airline’s revenue, load factor, and fleet utilization. These factors are also influenced by fuel price and oil price. All of these conditions then determine the fares and number of flights an aircraft will have, which in turn affects passenger demand. Also, the success of an airline then determines the number of orders aircraft manufacturers will receive. From here, lead times and cancellations affect the deliveries of new aircraft to the airlines. As you can see in the amended flow map presented in the figure below, there is a feedback loop between all of these variables, as they all affect each other either directly or indirectly.




Figure 1: Flow Map of the Airline Industry (adapted from Lyneis (2000))

It should now be clear that there are many moving parts in determining the number of aircraft orders and deliveries in a given quarter or year. Before beginning to forecast, we must understand the input variables further. In the next section, additional analysis of the input variables is performed.

3.2 Analysis of Input Variables

3.2.1 Seasonality

An important step in the forecasting process is understanding the underlying workings of each input variable. To do so, each input variable was graphed separately to identify any type of trend, seasonality or cycle. The corresponding graphs are presented in the Appendix.

The moving average method was used to de-seasonalize the data. It is a simple but robust tool for de-seasonalizing data and is therefore sufficient for this analysis. For the quarterly data, a centered moving average of 4 periods at a time was used to eliminate any seasonality, as the data exhibits upswings every 4th quarter of each year. The idea behind a moving average is to smooth out the seasonal variation by taking a rolling average of the data. Then, the seasonal factors are computed by dividing the original data by the averaged data values. Next, an average is taken for each quarter's seasonal factors to establish one seasonal factor for each of the four quarters. Finally, the original data is divided by the corresponding seasonal factor to generate a de-seasonalized data set.

Major seasonality was identified in the quarterly data for retirements. Minor seasonality was identified in net orders, orders and deliveries. The resulting graphs are presented below where the seasonality was removed for use in future analysis.

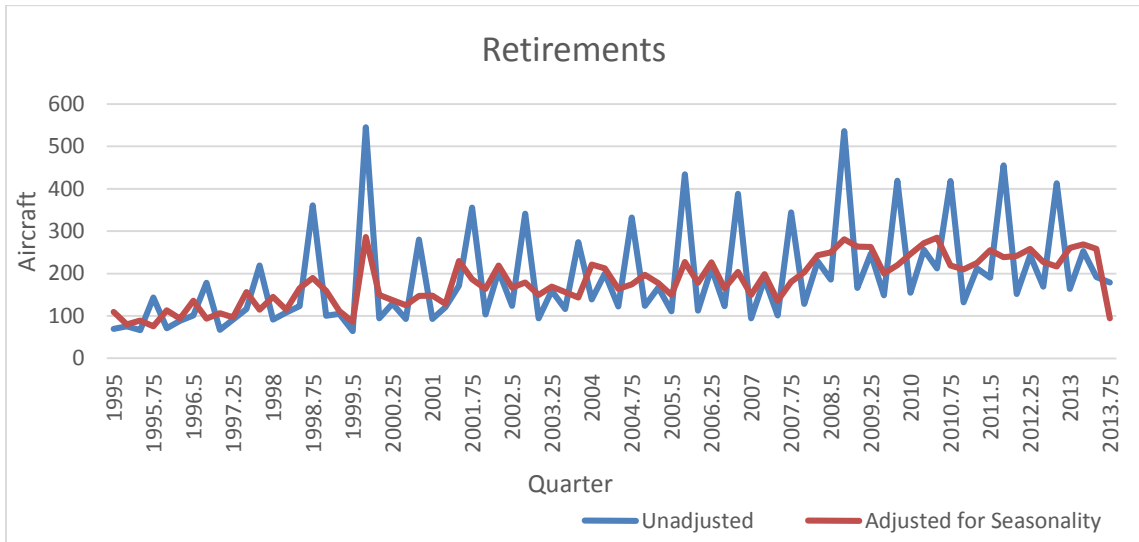


Figure 2: Seasonality of Retirements




Figure 3: Seasonality of Orders

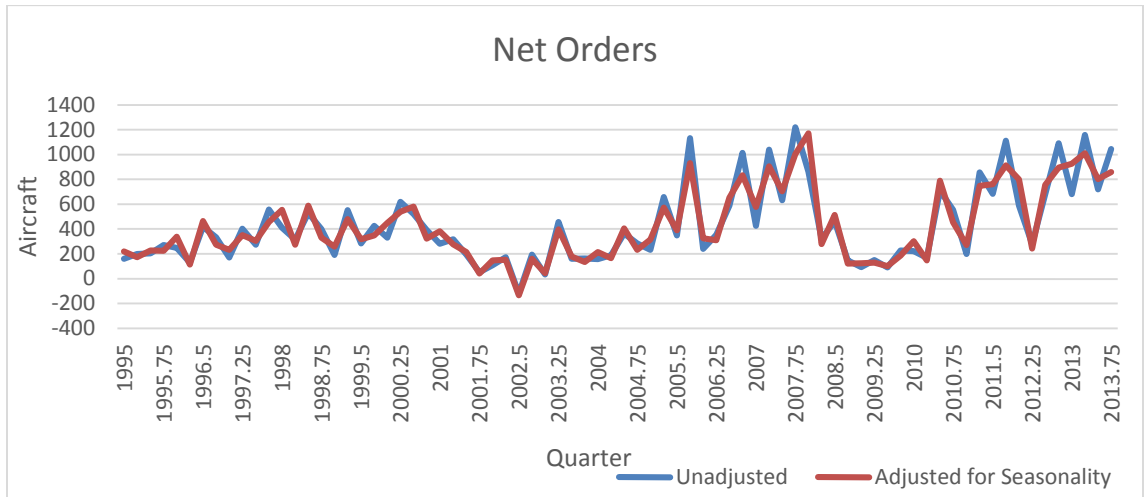


Figure 4: Seasonality of Net Orders




Figure 5: Seasonality of Deliveries

3.2.2 Volatility

Next, the volatility of each variable was measured. Here, volatility is represented as a sliding measure of the coefficient of variation (standard deviation divided by the mean). A time frame of eight quarters was used for the quarterly data, and four years for the annual data. The chart below represents the coefficient of variation for each variable under four different scenarios. The average volatility is measured for both quarterly and

annual data. In addition, the volatility of the most recent data (2009-2013) is measured for both quarterly and annual data.

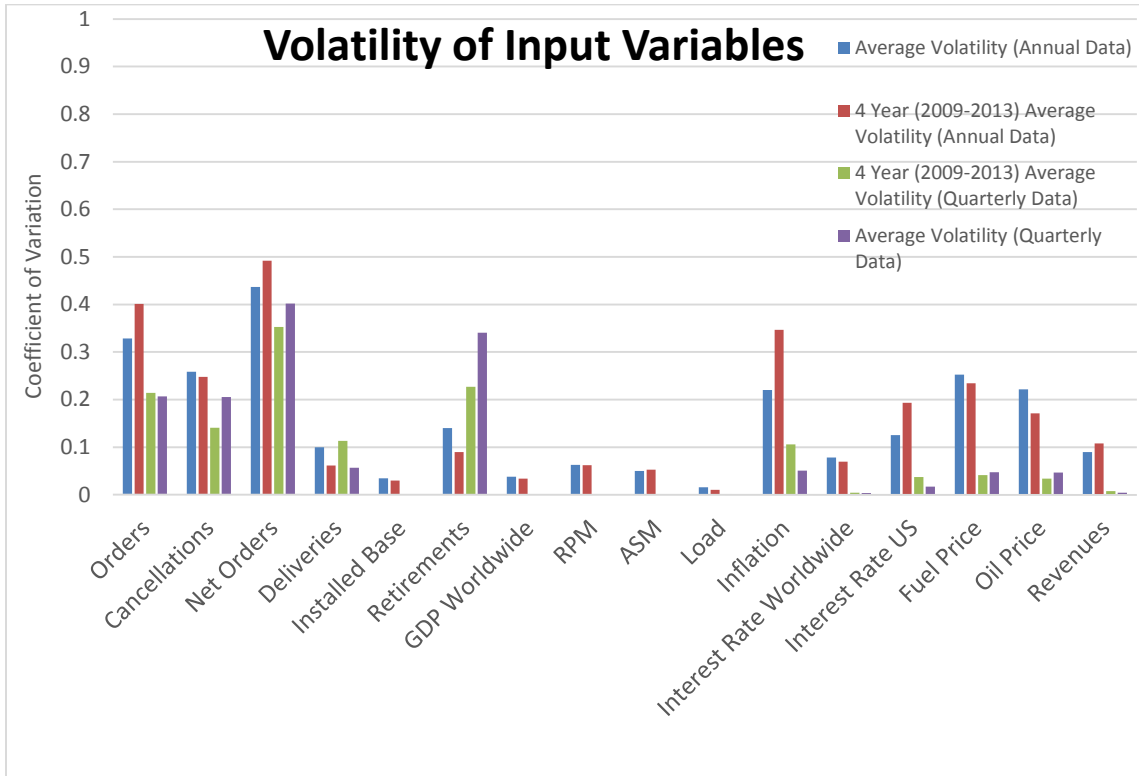


Figure 6: Quarterly and Annual Volatility of Input Variables

3.2.3 Linear Regression

Next, linear regression was performed on each input variable against both Orders and Deliveries separately. This was done to evaluate the predicting capacity of each variable as well as to evaluate the predicting ability of both quarterly and annual data. The tables below summarize the results in order from highest to lowest.

Table 1: R Squared Values for Input Variables for Orders

Linear Regression for Orders		
	Annual R Squared	Quarterly R Squared
Net Orders	0.9704	0.9684
Operating Revenue	0.6318	0.3533
ASM	0.6165	0.3438
Fuel Price	0.6174	0.3644
RPM	0.6033	0.3364
Oil Price	0.5781	0.3623
GDP	0.5342	0.2472
Installed Base	0.531	0.2394
Load	0.4958	0.2756
Cancellations	0.3631	0.2762
Interest Rate US	0.295	0.1778
Interest Rate Worldwide	0.2561	0.1362
Retirements	0.152	0.0825
GDP Growth	0.1505	0.0472
RPM Growth	0.0731	0.0428
ASM Growth	0.0345	0.0203
Inflation	0.0157	0.000

Table 2: R Squared Values for Deliveries

Linear Regression for Deliveries		
	Annual R Squared	Quarterly R Squared
Installed Base	0.6531	0.4625
ASM	0.6309	0.6186
RPM	0.6381	0.6142
Load	0.621	0.4871
GDP	0.6168	0.4432
Operating Revenue	0.605	0.6541
Interest Rate US	0.5721	0.4924
Retirements	0.4852	0.1817
Interest Rate Worldwide	0.4725	0.2877
Fuel Price	0.4744	0.5662
Oil Price	0.4781	0.5792
Cancellations	0.4214	0.3081
Net Orders	0.3103	0.5192
Inflation	0.2459	0.007
RPM Growth	0.0832	0.0021
ASM Growth	0.073	0.0016
GDP Growth	0.0322	0.0081

Additionally, the figures presented below show an additional view of the R squared values in decreasing order for orders and deliveries.



Figure 7: R Squared Values for Orders

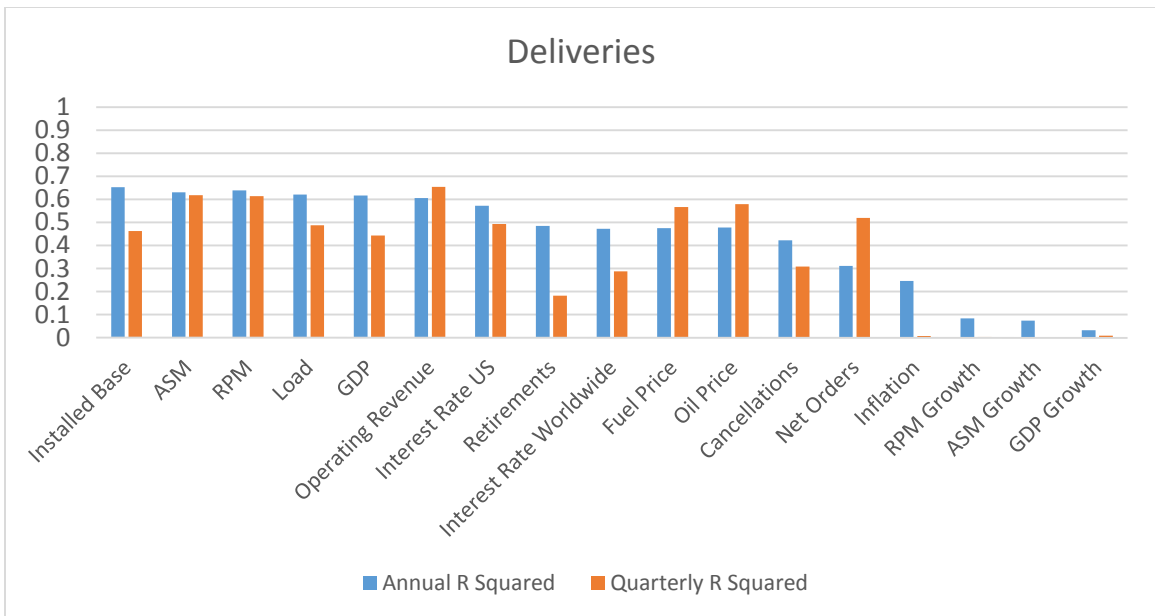


Figure 8: R Squared Values for Deliveries

Based on this analysis, deliveries seem to be overall easier to predict as the regression coefficients are higher than orders for most of the input variables. The variables

RPM Growth, GDP Growth, ASM Growth were consequently eliminated from further analysis due to the lack of accurate quarterly data and insignificant correlation to the dependent variables.

3.2.4 Correlation

An important consideration when forecasting is the relationship between each of the input variables. If two input variables are highly correlated with each other, using both in a regression model can cause error in forecasts. Essentially, we want the input variables to explain different portions of the variance for the dependent variable, and ideally when all the variables are put together in the model, all of the variance is explained. The table below presents the correlations between each of the input variables as well as the two dependent variables, orders and deliveries.

	Cancellations	Installed Base	Retirements	GDP	RPM	ASM	Load	Interest Rate Worldwide	Interest Rate US	Fuel Price	Oil Price	Revenue	Orders	Deliveries
Cancellations	1	0.58	0.47	0.58	0.55	0.55	0.56	-0.2	-0.56	0.55	0.52	0.53	0.53	0.56
Installed Base	0.58	1	0.75	0.98	0.95	0.95	0.98	-0.37	-0.91	0.88	0.88	0.91	0.49	0.68
Retirements	0.47	0.75	1	0.72	0.7	0.7	0.74	-0.33	-0.74	0.66	0.61	0.71	0.29	0.43
GDP	0.58	0.98	0.72	1	0.98	0.98	0.99	-0.31	-0.9	0.92	0.93	0.95	0.57	0.75
RPM	0.55	0.95	0.7	0.98	1	0.999	0.97	-0.25	-0.88	0.93	0.94	0.98	0.58	0.78
ASM	0.55	0.95	0.7	0.98	0.999	1	0.97	-0.24	-0.87	0.93	0.94	0.98	0.59	0.79
Load	0.56	0.98	0.74	0.99	0.97	0.97	1	-0.35	-0.89	0.89	0.9	0.93	0.52	0.7
Interest Rate World	-0.43	-0.87	-0.71	-0.84	-0.77	-0.76	-0.85	1	0.89	-0.6	-0.58	-0.66	-0.37	-0.54
Interest Rate US	-0.56	-0.91	-0.74	-0.9	-0.88	-0.87	-0.89	0.35	1	-0.82	-0.79	-0.84	-0.42	-0.7
Fuel Price	0.55	0.88	0.66	0.92	0.93	0.93	0.89	0.02	-0.82	1	0.92	0.95	0.6	0.75
Oil Price	0.52	0.88	0.61	0.93	0.94	0.94	0.9	-0.03	-0.79	0.92	1	0.95	0.6	0.76
Revenue	0.53	0.91	0.71	0.95	0.98	0.98	0.93	-0.09	-0.84	0.95	0.95	1	0.6	0.8

Figure 9: Correlation of Input Variables

It is clear here that RPM and ASM, Fuel Price and Oil Price, GDP and Load Factor, as well as GDP and Installed Base are highly correlated with each other. Many of the other variables have high correlations as well. This is important to note for further analysis, as highly correlated variables may hinder a forecast, and therefore only one of those variables may need to be selected for the model.

Further delving into the correlations between the input variables and the dependent variables, lag correlations were investigated. Lags of 0 to 8 quarters behind were investigated to determine if there was a delayed relationship between the input variables and the dependent variables. A lag of one, for example, uses a predictor variable such as jet fuel price a quarter ago, to predict orders in the current quarter. The same idea is then followed for the remaining lags to determine the correlations for each variable. The results are presented in the figures below. Only variables with significant correlations at the 95% confidence level are shown. This means that we can be 95% sure this is representative of the correlation, and the probability of observing a value outside this one is .05%.

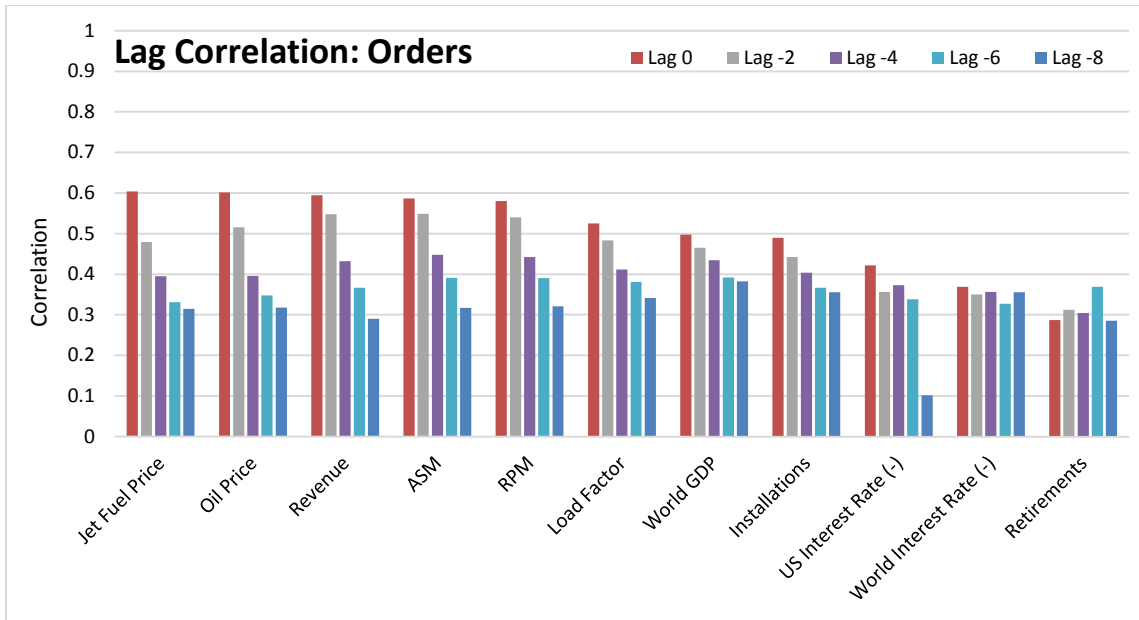


Figure 10: Lag Correlation for Orders

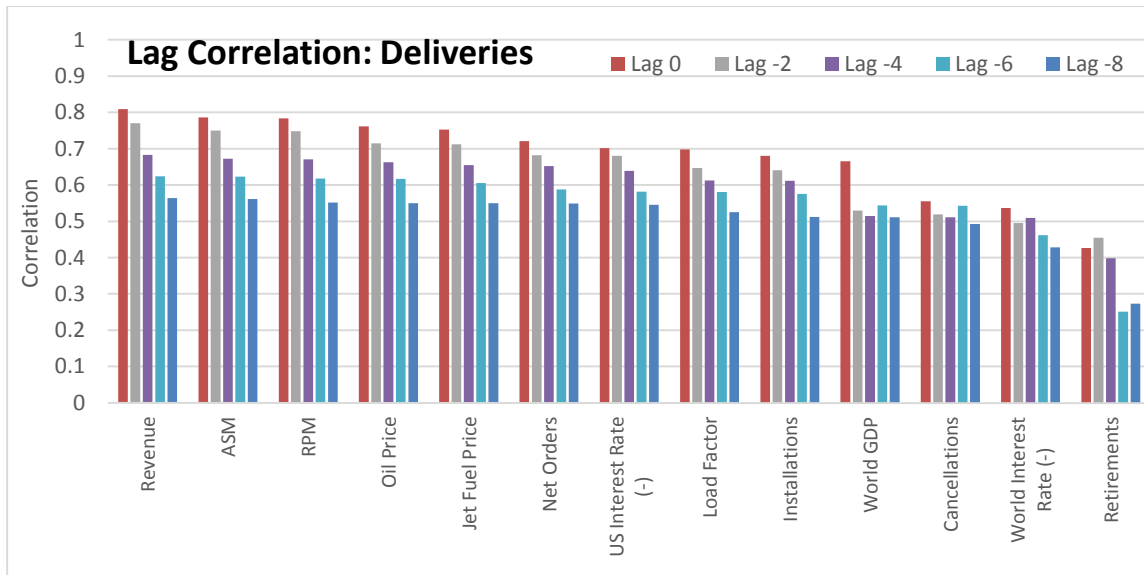


Figure 11: Lag Correlation for Deliveries

It is clear that deliveries as a whole are more correlated with the input variables than orders. Generally, a lag of zero provides the strongest correlation, with a few exceptions.

To conclude this analysis, it is evident that the input variables provide significant explanatory value. Additionally, adjusting for seasonality is critical to understanding the data better. The volatility of each input variable varies significantly from variable to variable. It is important to keep this in mind, as well as the inter-correlation of input variables, when forecasting.

CHAPTER 4

ANALYSIS AND RESULTS

4.1 Forecasting Methods and Models

Keeping the forecast simple for easy transfer to industry use was a primary consideration. It was found that Microsoft Excel worked well for performing most time series and regression methods with this data, and was the preferred platform for our industry partners. SAS software was used for ARIMA forecasting. It was important to keep this analysis relatively user friendly. Excel is not only user friendly, but is relatively inexpensive. Green and Armstrong (2015) focus on similar objectives, keeping the method simple with respect to the forecasting method and the number of input variables. Regression analysis was recommended as a sound forecasting technique. In addition, it was recommended to use a weighted combination of different forecasts. The following section presents a few different methods for forecasting orders and deliveries. Among them are Holt's Method, Holt-Winter's Method, Seasonal Factor Forecasting, Lagged Multiple Regression, and ARIMA forecasting.

With each method, the data for forecasting orders and deliveries will be broken up into two categories: a training set and a test set. The training set is the within sample data; all values that are being used to create a forecast model. The period being forecasted is the test set, and the actual realized values of orders and deliveries over the test set will be compared with the forecasted values to determine the forecast accuracy. In this thesis, our within sample period is 1995-2011, and the post sample period is 2012-2013. The forecast model applied to the training set is also compared with the actual values during that period

to provide a metric for model fit, that is, how well the model fits the data over the training set. The two main metrics of forecast fit and accuracy that are reported in this section for each forecasting method are the Mean Absolute Percent Error (MAPE) and the Root Mean Square Error (RMSE). The RMSE is then used to calculate Theil's U, which provides a metric for comparing different forecasting techniques and establishing their predictive validity.

4.1.1 Naïve No Change Method

To provide a baseline for evaluating more advanced methods, the naïve no change method was used to forecast orders and deliveries. The naïve no-change method simply develops a forecast for the given period (\hat{Y}_{t+1}) that is the actual value from the previous period, (Y_t). As rudimentary as this seems, this provides a baseline for which more sophisticated methods and models should have greater accuracy.

The corresponding naïve no change forecasts for orders and deliveries are presented below. As you can see, the forecast is simply the actual values shifted ahead one period.




Figure 12: Naive No Change Forecast for Orders

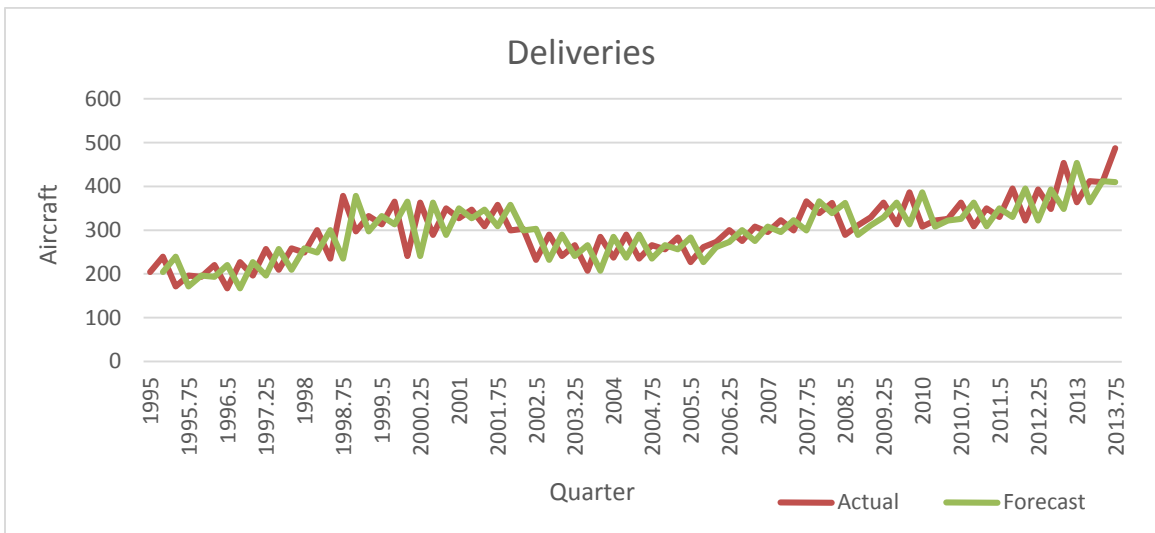


Figure 13: Naive No Change forecast for Deliveries

The performance statistics for orders and deliveries are presented in the table below. Again, the Fit values are for the period 1995-2011 and the Forecast values are for the period of 2011-2013.

Table 3: Naive No Change Performance Statistics

	Orders	Deliveries
MAPE Fit	51.75%	16.04%
RMSE Fit	327.383	52.68
MAPE Forecast	57.84%	16.21%
RMSE Forecast	499.964	70.68

The results from the naïve no change method are quite primitive, however they provide a great baseline for both orders and deliveries. The set of performance statistics presented above will be used further with Theil’s U statistic when assessing the validity of more sophisticated forecasting methods.

4.1.2 Holt’s Method

First, Holt’s Method is used on the annual and quarterly data to forecast orders and deliveries. Data from 1995 to 2011 was used to forecast for 2012 and 2013. The actual values for 2012 and 2013 were then compared to the forecasted values. The corresponding graphs for orders are presented below.

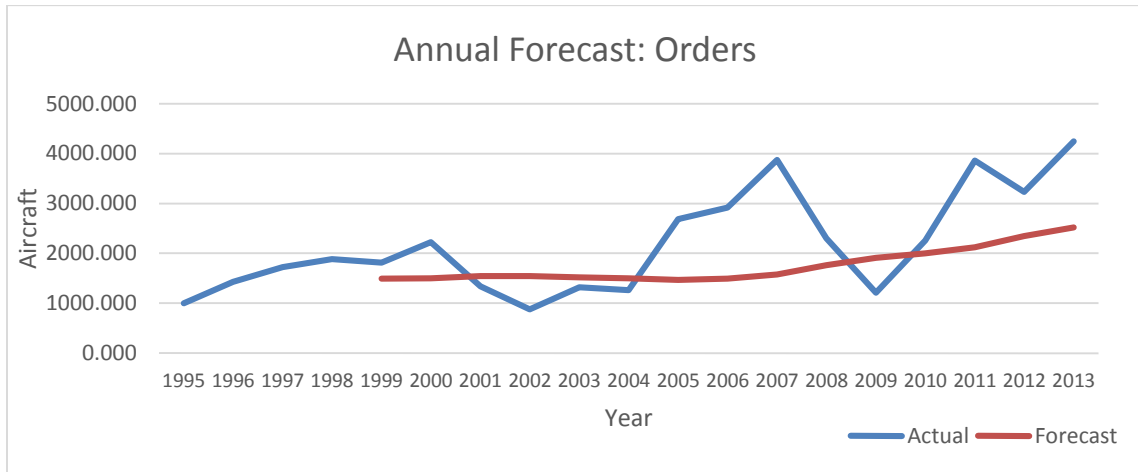


Figure 14: Holt's Method for Annual Orders




Figure 15: Holt's Method for Quarterly Orders

The performance statistics are presented in the table below for both the quarterly and annual forecasts for orders.

Table 4: Performance Statistics for Orders Using Holt's Method

	Quarterly	Annual
MAPE Fit	58.58%	35.93%
RMSE Fit	321.94	1033.27
MAPE Forecast	26.02%	34.10%
RMSE Forecast	315.60	1373.23

The corresponding graphs for quarterly and annual forecasts for deliveries are presented below.




Figure 16: Holt's Method for Annual Deliverieies

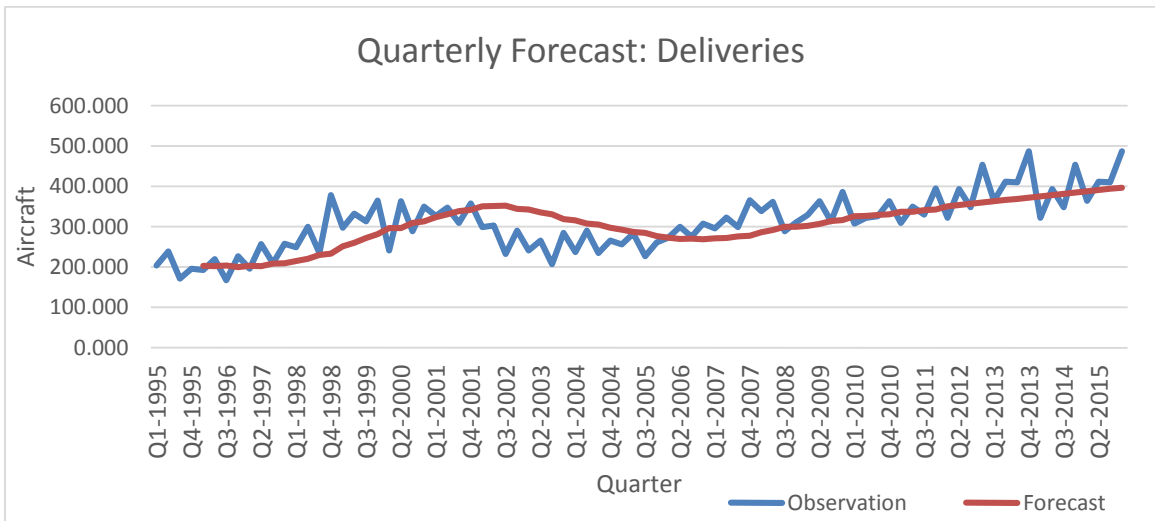


Figure 17: Holt's Method for Quarterly Deliveries

The performance statistics are presented in the table below for both the quarterly and annual forecasts for deliveries.

Table 5: Performance Statistics for Deliveries using Holt's Method

	Quarterly	Annual
MAPE Fit	14.17%	9.72%
RMSE Fit	51.57	166.86
MAPE Forecast	10.89%	14.15%
RMSE Forecast	47.95	270.43

Overall, Holt's method seems to provide a much more reliable forecast for deliveries than orders. Additionally, it seems as though the MAPE fit denotes that the annual model is more accurate, whereas the MAPE forecast indicates a more accurate quarterly model. In both cases the RMSE Fit and Forecast reveals a more accurate quarterly model.

4.1.3 Holt-Winters Method

Next, Holt-Winters Method was used to accommodate a potential additional factor of seasonality. Holt-Winters method was used explicitly on the quarterly data for orders and deliveries, as minor seasonality was found in both variables during the analysis of input variables. The resulting forecasts for orders and deliveries are presented below.




Figure 18: Holt-Winter's Method for Quarterly Orders

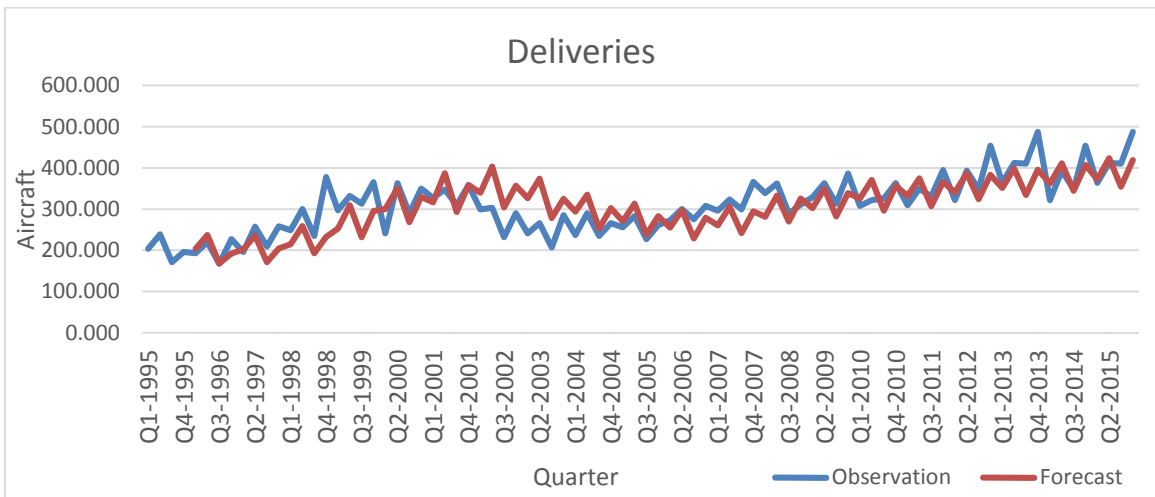


Figure 19: Holt-Winter's Method for Quarterly Deliveries

Next, the performance statistics for each forecast are presented in the table below.

Table 6: Performance Statistics for Orders and Deliveries using Holt-Winter's Method

	Orders	Deliveries
MAPE Fit	63.31%	12.93%
RMSE Fit	311.46	46.17
MAPE Forecast	29.59%	9.23%

RMSE Forecast	283.82	50.62
----------------------	--------	-------

Based on these results, it seems that Holt’s method provides a more accurate forecast for orders, whereas Holt-Winter’s method provides a slightly more accurate forecast for deliveries. Deliveries seem to be slightly more seasonal, so intuitively it makes sense for Holt-Winters Method to be more applicable. In general though, both methods seem to produce similar results, and are quite robust for forecasting deliveries.

4.1.4 Forecasting using Aggregate Annual Data with Seasonal Factors

Here, forecasting using aggregate annual data and seasonal factors was explored to determine if a more reliable quarterly forecast could be generated. By dividing each original annual data forecast (by four and multiplying those values by the corresponding quarterly seasonal factor, we obtain a quarterly forecast. The subsequent forecasts and seasonal factors are presented below.

Table 7: Seasonal Factors for Orders and Deliveries

	Orders	Deliveries
Quarter 1	0.77	0.90
Quarter 2	1.04	1.07
Quarter 3	0.97	0.88
Quarter 4	1.23	1.15




Figure 20: Seasonal Factor Forecast for Deliveries




Figure 21: Seasonal Factor Forecast for Orders

Finally, the performance measures for both orders and deliveries are presented in the table below.

Table 8: Performance Measures for Seasonal Factor Forecasts

	Orders	Deliveries
MAPE Fit	30.63%	5.57%
MAPE Forecast	24.76%	8.50%
RMSE Fit	172.624	18.988
RMSE Forecast	222.313	40.986

The results from this forecasting method indicate again that deliveries produce a much more reliable forecast than orders. It seems the forecast for deliveries is quite accurate, with a minimal MAPE and RMSE. Overall, this method provides an alternative technique for forecasting seasonal data.

4.1.5 Multiple Regression with Lagged Values

In this section, lagged multiple regression is used to forecast eight quarters ahead for quarterly data and three years ahead for annual data. Beginning with one period lagged and increasing up to 8 periods lagged, multiple regression is performed on all of the input variables to create a forecast that is the average of all of the lags. The figure below displays a visualization of the concept of a lag of 4 quarters. Here, the value from 4 quarters back is used to predict the current value.

Time	Predictor Variable (GDP, \$Billions)	Response Variable (Aircraft Deliveries, #Aircraft)
1	50,345	299
2	50,802	318
3	51,170	232
4	51,418	290
5	51,617	241
6	52,023	312
7	52,750	245
8	53,396	287
9	54,051	237
10	54,538	307

Figure 22: Example of a 4 quarter lag

This is particularly useful to capture any lagged relationships between the variables. The resulting forecasts for 2012 and 2013 for deliveries compared against actual values are presented below.




Figure 23: Lagged Multiple Regression Forecast for Quarterly Deliveries

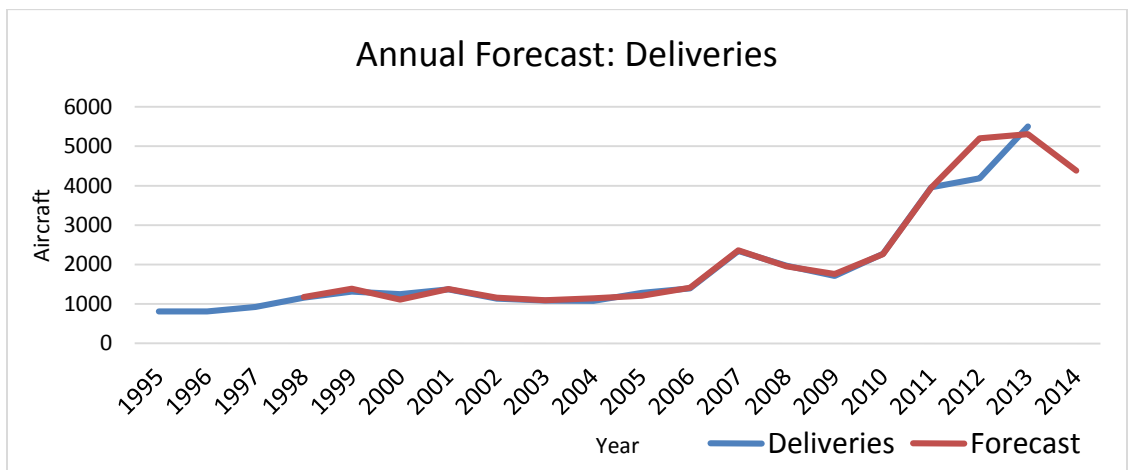


Figure 24: Lagged Multiple Regression Forecast for Annual Deliveries

The performance statistics for the quarterly forecast for Deliveries, including MAPE and RMSE for both the fit and forecast are presented in the table below.

Table 9: Performance Statistics for Quarterly Deliveries

MAPE Fit	9.85%
MAPE Forecast	6.12%
RMSE Fit	34.045
RMSE Forecast	32.304

For the annual forecast for deliveries, a MAPE Forecast of 24% and a RMSE Forecast of 1100.32 was achieved.

This analysis was then repeated for Orders, lagging values from one period to eight periods lagged, and then finding the average forecast 2012 and 2013 for all lags. The quarterly and annual forecasts are presented below.

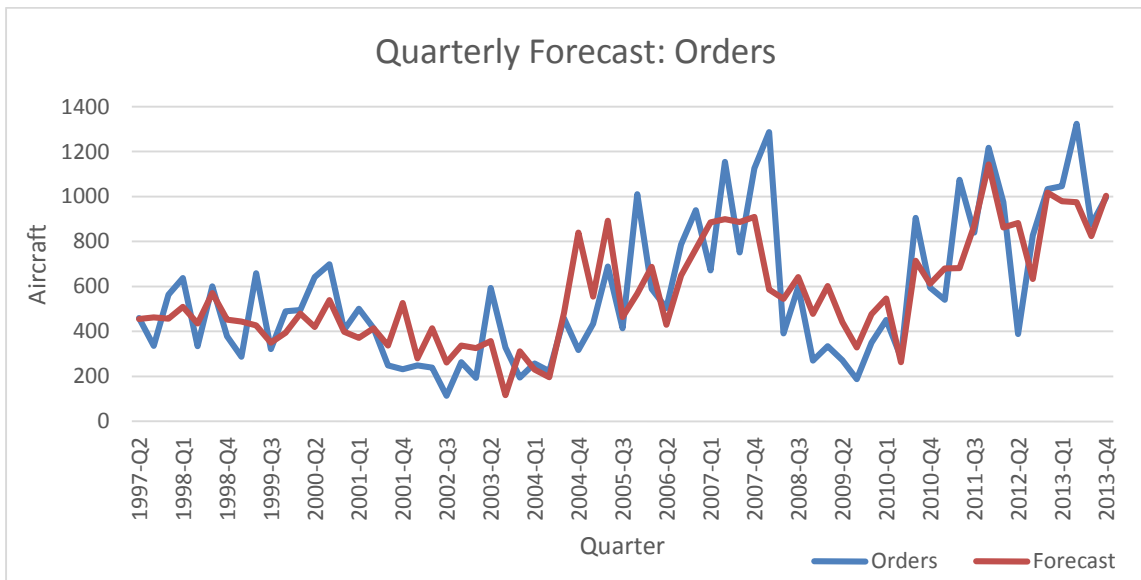


Figure 25: Lagged Multiple Regression for Quarterly Orders

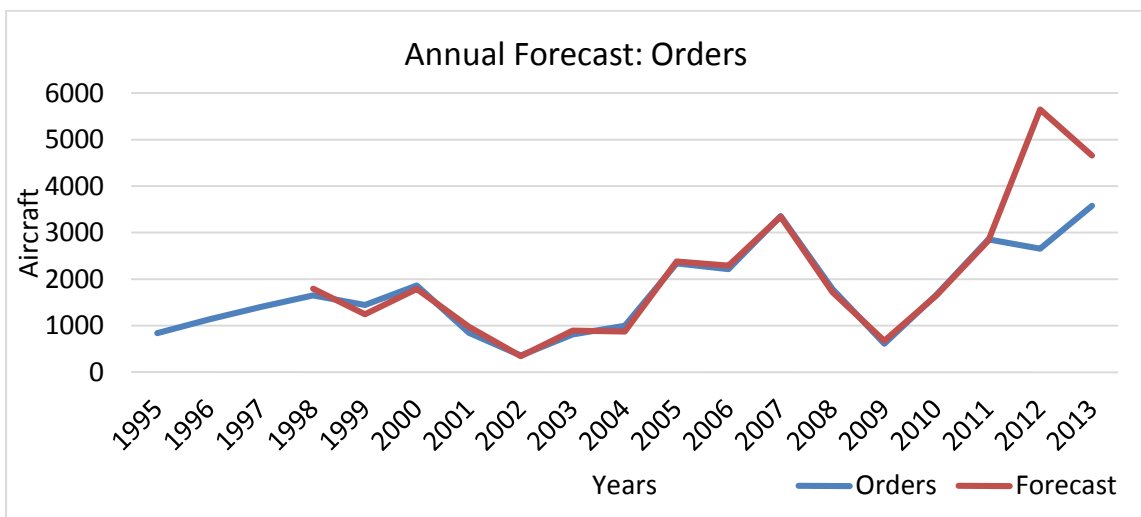


Figure 26: Lagged Multiple Regression for Annual Orders

The performance statistics for the quarterly forecast for Orders, including MAPE and RMSE for both the fit and forecast are presented in the table below.

Table 10: Performance Statistics for Quarterly Orders

MAPE Fit	38.35%
MAPE Forecast	25.32%
RMSE Fit	103.523
RMSE Forecast	229.773

Additionally, the MAPE Forecast was 53% and the RMSE Forecast was 2086.09 for the annual forecast for orders. This value is surprisingly high compared the performance measures for the quarterly forecast. Overall, it seems the quarterly forecast using multiple regression is more accurate than the annual forecast for both orders and deliveries.

Unmistakably it is again clear that the forecasts for orders are significantly less accurate than those for deliveries. Intuitively, this makes sense as it is commonplace in the aerospace and defense industry for orders to be placed and then cancelled. Those that are placing the orders themselves are basing their order on an expectation of the future, and is therefore subject to change.

4.1.6 ARIMA Forecasting

In this section, SAS software was used to analyze the series and ultimately generate forecasts for orders and deliveries using the Autoregressive Integrated Moving Average (ARIMA) model.

The first step of this analysis is to identify the correct ARIMA model to use for each variable. SAS was used to run a sequence plot of the respective variable. This aids in determining if the series is stationary. Stationarity needs to be achieved before an ARIMA

model can be used. In the SAS output presented in the figure below, it is clear that the series for deliveries is non-stationary since its autocorrelation function (ACF) plot decays very slowly.

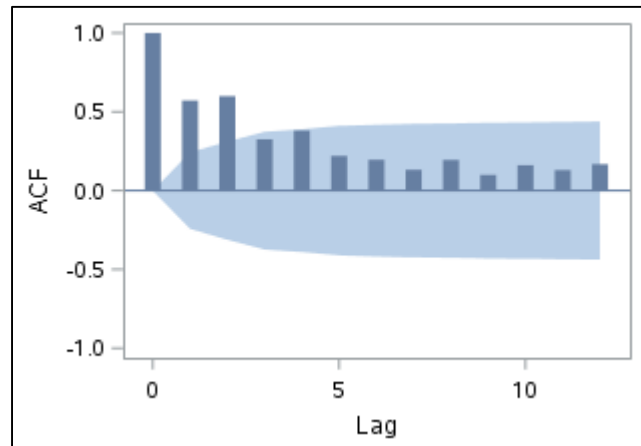


Figure 27: Initial ACF plot for Deliveries

Since the data is non-stationary, it was first differenced in SAS by taking the logarithm of the data. The figure below displays the autocorrelation function (ACF) plot as well as the partial autocorrelation plot (PACF) for the differenced series. The autocorrelation plot of the differenced series suggests that the series is now stationary. The ACF plot cuts off after the 3rd lag (above the 95% confidence level), therefore this implies that an ARIMA (0,1,2) model could be used. Essentially, when a plot “cuts off,” it means the lags suddenly cut off after a certain number of lags, and dip lower than the 95% confidence band. However, looking at the PACF, it seems that an ARIMA (1,1,0) model may be sufficient since the lags are not significant past the first one.




Figure 28: ACF and PACF plots for Differenced Deliveries

Next, the AIC criterion will be used to decide between the two possible models.

The AIC values are presented in the table below for both models.

Table 11: AIC Values for ARIMA Models for Deliveries

Model	AIC Value
ARIMA(1,1,0)	838.698
ARIMA(0,1,2)	842.576

Since the AIC Value is less for the ARIMA (1,1,0) model, it will be selected for all further analysis. The prediction equation is presented below.

$$\hat{Y}_t = \mu + Y_{t-1} - \varphi_1(Y_{t-1} - Y_{t-2})$$

From here, the resulting forecast from this model is presented below.




Figure 29: ARIMA Forecast for Quarterly Deliveries

The performance statistics are presented in the table below for the within and post sample for deliveries.

Table 12: ARIMA Performance Statistics for Deliveries

MAPE Fit	14.95%
RMSE Fit	51.806
MAPE Forecast	15.02%
RMSE Forecast	60.271

Next running the same SAS analysis for orders, the autocorrelation plot is presented below.




Figure 30: Initial ACF plot for Orders

Again, it is clear that the data is not stationary since it tails off slowly, therefore first differencing must be done. The resulting ACF and PACF graphs are presented below for orders.




Figure 31: ACF and PACF plots for Differenced Orders

Based on these graphs, an $ARIMA(2,1,0)$ could be used, since the PACF cuts off after lag 2. Conversely, an $ARIMA(0,1,2)$ model could be used, since the ACF clearly cuts off after a lag of 2. Therefore, we will run both models, and use the AIC to decide between them. The AIC presented below indicates that the $ARIMA(0,1,2)$ wins out slightly.

Table 13: AIC Values for ARIMA Models for Orders

Model	AIC Value
ARIMA(2,1,0)	918.858
ARIMA(0,1,2)	917.401

The prediction equation for the ARIMA (0,1,2) model is presented below.

$$\hat{Y}_t = \mu + Y_{t-1} - \theta_1 e_{t-1} - \theta_2 e_{t-2}$$

Thus, proceeding further with this model, the resulting forecast for quarterly orders is presented below.




Figure 32: ARIMA Forecast for Quarterly Orders

The subsequent performance statistic are presented in the table below for the within and post sample for orders.

Table 14: ARIMA Performance Statistics for Orders

MAPE Fit	42.25%
RMSE Fit	248.761
MAPE Forecast	26.03%
RMSE Forecast	260.437

It is clear that ARIMA forecasting is quite robust for forecasting quarterly deliveries, however the fit is not ideal for quarterly orders.

4.2 Evaluation of Forecasts with Theil's U

In this section, each of the forecasts presented in the previous section will be compared and evaluated using Theil's U Statistic. It is important to again state that Theil's U measures the worth of a forecasting method that is deemed to be more advanced than the naïve no-change method. The "U statistic" is the ratio of the RMSE of the more sophisticated method being analyzed to the RMSE of the naïve no-change method. The formula is presented again below:

$$\textit{Theil's U} = \frac{\textit{RMSE of advanced approach}}{\textit{RMSE of naïve no – change}}$$

If $U > 1$, then the advanced approach has no value because it cannot perform as well as the naïve no change basic method. However, if $U < 1$, the advanced approach has more merit. The closer U gets to 0, the better the approach in question is. (Newbold and Bos 1993)

The table below indicates the Theil's U statistic for each of the quarterly forecasting methods used in this thesis. For the purpose of comparing and aggregating, quarterly forecasts will be focused on from this point further. The within and post sample values for Theil's U were calculated to evaluate the accuracy of the model fit (1995 to 2011) and the model forecast (2012-2013), respectively.

Table 15: Theil's U Statistics by Forecasting Model for Deliveries

Model	Within Theil's U	Post Theil's U
Holt's Method	0.9789	0.6784
Holt-Winter's Method	0.8764	0.7161
Seasonal Factor Forecast	0.3604	0.5799
Multiple Regression	0.6463	0.4570
ARIMA	0.9834	0.85271

Table 16: Theil's U Statistics by Forecasting Model for Orders

Model	Within Theil's U	Post Theil's U
Holt's Method	0.9833	0.6312
Holt-Winter's Method	0.9514	0.5677
Seasonal Factor Forecast	0.5273	0.4447
Multiple Regression	0.3162	0.4596
ARIMA	0.7598	0.5209

From here, it can be concluded that all forecasts can be used for the aggregation linear program to create a final forecast for 2012-2013 orders and deliveries, since the Theil's U statistics are less than 1.

4.3 Aggregation Linear Program

The following linear program was used to aggregate the best forecasts to produce an optimal combination of multiple forecasts. The linear program minimizes the Mean

Absolute Percent Error by assigning specific weights, x_i , to each forecast. This creates an aggregate forecast that is ultimately the optimal combination of forecasts.

Let:

T: Total Forecasted Periods

K: Total number of Forecast Models

d_i : Actual value at time i

f_{ij} : Forecast for time period i in forecast model j

y_i : Error

j : Weight assigned to forecast model j

$$\text{Minimize MAPE: } \frac{1}{T} \sum_{i=1}^T \left[\frac{y_i}{d_i} * 100 \right]$$

Subject to:

$$\text{c1: } y_i \geq d_i - (\sum_{j=1}^K x_j f_{ij}) \quad \forall i$$

$$\text{c2: } y_i \geq -(d_i - (\sum_{j=1}^K x_j f_{ij})) \quad \forall i$$

$$\text{c3: } x_j \geq 0$$

This linear program was implemented in AMPL and used to aggregate forecasts for both orders and deliveries. The resulting aggregate forecasts for 2012-2013 are presented below compared the actual values for those years. The linear program only assigned weights to the Holt-Winter's, Multiple Regression and ARIMA forecasts for both Orders and Deliveries, deeming the forecasts from Holt's Method and the Seasonal Factor forecast not necessary for aggregation. The table presented below shows the weights assigned to each forecasting method.

Table 17: Aggregate Linear Program Weights assigned to each Forecast

	Orders	Deliveries
Seasonal Factor	0	0
Holt	0	0
Holt-Winters	.21	.32
Multiple Regression	.42	.35
ARIMA	.37	.33

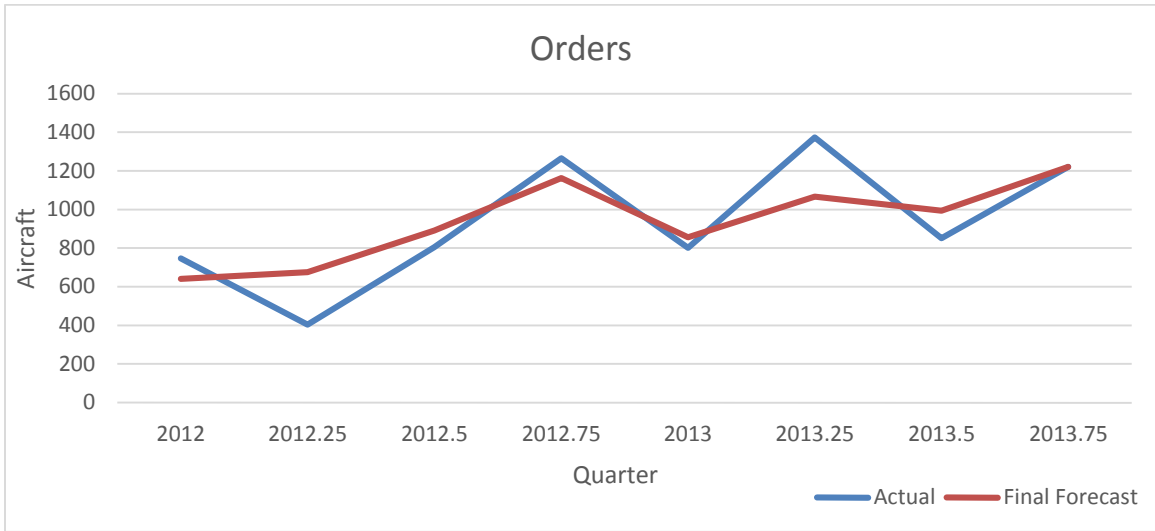


Figure 33: Linear Program Aggregate Forecast for Orders




Figure 34: Linear Program Aggregate Forecast for Deliveries

Additionally, the performance measures for each of the forecasts are presented below.

Table 18: Performance Measures for Aggregate Linear Program

	Orders	Deliveries
MAPE	18.35%	7.22%
RMSE	166.452	32.033
Theil's U	.250	.453

From here, it can be concluded that the aggregate forecast is highly applicable to both orders and deliveries, and it is an effective method for producing a more balanced forecast from multiple methods.

This linear program is a unique way to aggregate multiple forecasts. The more commonplace method is to simply average all values, essentially giving an equal weight to each forecast. Therefore, to contrast, the results from the simple average method are presented below.




Figure 35: Simple Average Aggregate Forecast for Orders




Figure 36: Simple Average Aggregate Forecast for Deliveries

The performance statistics for this simple average aggregate forecast are presented in the table below.

Table 19: Performance Statistics for Simple Average Aggregate Forecasts

	Orders	Deliveries
MAPE	20.45%	7.57%
RMSE	194.726	33.430
Theil's U	.389	.473

According to the MAPE, the average aggregation method performs slightly worse than the linear program aggregation method for orders and deliveries, but still provides a sufficient collective solution. The aggregation linear program that minimizes the MAPE is a unique and substantial method to combine multiple forecasts and produce an optimal final forecast. To conclude, both methods are beneficial to improving a final forecast through aggregation.

4.4 Summary of Model Performance

As a refresher and for comparison purposes, this section will briefly summarize the performance statistics for each of the seven forecasting models developed in this thesis. The tables below are for orders and deliveries, respectively.

Table 20: Summary of Performance Statistics for Orders

Performance Statistics for Orders				
	MAPE Fit	RMSE Fit	MAPE Forecast	RMSE Forecast
Naïve	51.75%	327.3831	57.84%	499.964
Holt	58.58%	321.94	26.02%	315.6
Holt-Winters	63.31%	311.46	29.59%	283.82
Seasonal Factor	30.63%	172.624	24.76%	222.313
Multiple Regression	38.35%	103.523	25.32%	229.773
ARIMA	42.25%	248.761	26.03%	260.437
Aggregate	16.67%	112.491	18.35%	166.452

Table 21: Summary of Performance Statistics for Deliveries

Performance Statistics for Deliveries				
	MAPE Fit	RMSE Fit	MAPE Forecast	RMSE Forecast
Naïve	16.04%	52.68	16.21%	70.68
Holt	14.17%	51.57	10.89%	47.95
Holt-Winters	12.93%	46.17	9.23%	50.62
Seasonal Factor	5.57%	18.988	8.50%	40.986
Multiple Regression	9.85%	34.045	6.12%	32.304
ARIMA	14.95%	51.806	15.02%	60.271
Aggregate	6.14%	28.461	7.22%	32.033

From here, it is clear that the most accurate forecasting model for orders according to the MAPE is the aggregate model, with a MAPE Forecast of a little over 18%. The results are not as clear for deliveries, where the seasonal factor model has a superior MAPE Fit of 5.57%, but the aggregate model has a superior MAPE Forecast of 7.22%. Therefore, we can conclude that both methods are sufficient for forecasting deliveries.

CHAPTER 5

DISCUSSION AND CONCLUSION

5.1 Overall Performance of Selected Forecasting Models

This thesis implemented different methods and models for forecasting aircraft orders and deliveries. Based on the results presented in the previous section, it is first important to note that all forecasting techniques were deemed more accurate than the Naïve No-Change forecast, according to Theil's U. This indicates that each forecast is more sophisticated than the most rudimentary method, and was sufficient for further analysis. After aggregation with the Linear Program, it became apparent that the Multiple Regression, Holt-Winters, and ARIMA quarterly forecasts were superior to the Holt and Seasonal Factor forecasts for both Orders and Deliveries, over the forecasting horizon of 8 quarters.

The Multiple Regression model captured the past behavior of the economic indicators for forecasting. It was extremely important to first analyze the input variables for the regression model prior to forecasting, as correlations between predictor variables needed to be identified. Highly correlated input variables can hinder a forecast, therefore it was important to eliminate highly correlated input variables for the regression analysis. The time series models used in this thesis effectively captured the pure data generating process of orders and deliveries to create a forecast. Both methods are arguably necessary to produce a wholesome aggregate forecast for both the variables of interest. Additionally, based on the performance measures from the aggregate forecasts, it can be concluded that the error is generally quite low for a forecasting task as challenging as this one. It is clear

that forecasting for orders is much more difficult than deliveries, which is probably due to the lack of congruency in the industry between orders and deliveries themselves. Deliveries, generally, are more stable as well. Additionally, from a practical supply chain perspective, it is inherently more valuable to have a prediction for deliveries, rather than orders, since orders can change during the three year lead time.


Of course, the forecasting models used in this thesis have their own shortcomings as they could only capture the behavior of the intended variables to a certain degree of accuracy. However, in terms of applicability to industry, the models developed in this thesis are simple, inexpensive, user friendly and sufficiently accurate.




5.2 Limitations




This thesis focused on a forecasting horizon of two years, or eight quarters which maintained the relative accuracy of the forecasts given the respective models. However, proceeding further out to a longer forecasting horizon would undoubtedly negatively impact the forecasting accuracy as a whole. Therefore, the aggregate forecasting methods employed in this thesis are limited. More robust machine learning methods could be considered to forecast a longer horizon and are expected to improve accuracy.




APPENDIX



SUPPORTING BASELINE GRAPHS




















BIBLIOGRAPHY

- Airbus (2015) *Global Market Forecast 2015-2034*.
- Bartlett, H. (1965). *The demand for air transportation, 1947–1962*. Unpublished Ph. D. dissertation. University of Michigan.
- Boeing (2015) *Current Market Outlook 2015-2033*.
- Box, G.E.P., & Jenkins, G.M. (1970) *Time series analysis: Forecasting and Control*. San Francisco: Holden Day.
- Brown S.M. & Watkins W.S. (1968) *The Demand for Air Travel: A regression study of time-series and cross sectional data in the US domestic market*. Highway Research Record
- Carson, R. T., Cenesizoglu, T., & Parker, R. (2011). Forecasting (aggregate) demand for US commercial air travel. *International Journal of Forecasting*, 27(3), 923-941.
- Chatfield, C., Koehler, A. B., Ord, J. K., & Snyder, R. D. (2001). A new look at models for exponential smoothing. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 50(2), 147-159.
- Cheng, C., Sa-Ngasoongsong, A., Beyca, O., Le, T., Yang, H., Kong, Z., & Bukkapatnam, S. T. (2015). Time series forecasting for nonlinear and non-stationary processes: a review and comparative study. *IIE Transactions*, (ahead-of-print), 1-19.
- Clemen, R. T. (1989). Combining forecasts: A review and annotated bibliography. *International journal of forecasting*, 5(4), 559-583.
- De Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International journal of forecasting*, 22(3), 443-473.
- Duncanson A.V (1974) *Short-Term Traffic Forecasting*. Proceedings of the 14th Annual AGIFORS Symposium.
- Enders, Walter. (2004). *Applied Econometric Time Series*. Hoboken, NJ: J. Wiley, Print.
- Fang, Y. (2003). Forecasting combination and encompassing tests. *International Journal of Forecasting*, 19(1), 87-94.
- Granger, C. W., & Ramanathan, R. (1984). Improved methods of combining forecasts. *Journal of Forecasting*, 3(2), 197-204.
- Green, K.C., Armstrong, J.C. (2015) *Simple Forecasting: Avoid Tears before Bedtime (Working Paper)*. Journal of Business Research.

- Haney, D. (1975). *Review of aviation forecasting methodology*. Rep. dot-40176, 6.
- Harvey, A. C. (1990). *Forecasting, structural time series models and the Kalman filter*. Cambridge university press.
- Harvey, D.A. (1951) *Airline passenger traffic patterns within the United States*. Journal of Air Law and Commerce.
- Hibon, M., & Evgeniou, T. (2005). To combine or not to combine: selecting among forecasts and their combinations. *International Journal of Forecasting*, 21(1), 15-24.
- Ho, K. K. (2008). *Demand forecasting for aircraft engine aftermarket* (Doctoral dissertation, Massachusetts Institute of Technology).
- Hyndman, R. J. (2001). It's time to move from what to why'. *International Journal of Forecasting*, 17(1), 567-570.
- Karlaftis, M. G., Zografos, K. G., Papastavrou, J. D., & Charnes, J. M. (1996). Methodological framework for air-travel demand forecasting. *Journal of Transportation Engineering*, 122(2), 96-104.
- Liehr, M., Größler, A., Klein, M., & Milling, P. M. (2001). Cycles in the sky: understanding and managing business cycles in the airline market. *System Dynamics Review*, 17(4), 311-332.
- Littlewood, K. (1972) *Forecasting and Control of Passenger Bookings*. Proceedings of the 12th Annual AGIFORS Symposium.
- Lyneis, J. M. (2000). System dynamics for market forecasting and structural analysis. *System Dynamics Review*, 16(1), 3-25.
- Mellman, R., Nelson, M., & Piro, J. (1980). *FORECASTS OF PASSENGER AND AIR-CARGO ACTIVITY AT LOGAN INTERNATIONAL AIRPORT* (No. 768).
- Miller, C. M., Clemen, R. T., & Winkler, R. L. (1992). The effect of nonstationarity on combined forecasts. *International Journal of Forecasting*, 7(4), 515-529.
- Newbold, P., Bos, T. (1993) *Introductory Business and economic Forecasting*. Second edition, South-Western Publishing, Cincinnati, Ohio.
- Oberhausen P.J. and Koppelman F.S. (1982) *Time series Analysis of Intercity Air Travel Volume*. Transportation and Research Record 1221.
- Otexts. (2015). *Holt's Linear Trend Method*. Retrieved from <https://www.otexts.org/fpp/7/2>

Proussaloglou, K., & Koppelman, F. (1995). Air carrier demand. *Transportation*, 22(4), 371-388.

Taneja, N. K. (1979). AIRLINE TRAFFIC FORECASTING; A REGRESSION ANALYSIS APPROACH.

Thomet, M. A., & Sultan, S. M. (1979). Traffic Forecasting for the New Riyadh International Airport. *ICAO manual on air traffic forecasting*.

Verleger P.K. (1972) *Models of the Demand for Air Transportation*. The Bell Journal of Economics and Management Science.

Vitek, R. L., & Taneja, N. K. (1975). *The impact of high inflation rates on the demand for air passenger transportation*. Cambridge, Mass.: Flight Transportation Laboratory, Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, [1975].

Yu, G. (Ed.). (2012). *Operations research in the airline industry* (Vol. 9). Springer Science & Business Media.