TIME SERIES FORECASTING OF
TANKER TRAINING DEMAND

GRADUATE RESEARCH PAPER

Steven R. Hawkins, Major, USAF

AFIT-ENS-MS-19-J-033

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT-ENS-MS-19-J-033

TIME SERIES FORECASTING OF TANKER TRAINING DEMAND

GRADUATE RESEARCH PAPER

Presented to the Faculty

Department of Operational Sciences

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Operations Management

Steven R. Hawkins, BS, MBA

Major, USAF

June 2019

AFIT-ENS-MS-19-J-033

TIME SERIES FORECASTING OF TANKER TRAINING DEMAND

Steven R. Hawkins, BS, MBA

Major, USAF

Committee Membership:

William A. Cunningham III, Ph.D.
Chair (Research Advisor)

Lt Col John M. Dickens, Ph.D.
Member/Reader

AFIT-ENS-MS-19-J-033

Abstract

Since 1998 the USAF tanker fleet has decreased by 155 aircraft, with the additional anticipated retirement of the equivalent of 201 KC-135 tankers by 2029. Meanwhile, the new KC-46A is scheduled to replace only 179 boom-equipped tankers by 2029, resulting in a total decrease of 131 aircraft which corresponds to a reduction of 24.5% in total air refueling fuel capacity in just over three decades. This research examines historical tanker training requests, drawn from the 618th AOC's Air Refueling Scheduling Tool (ARST), and uses multiple forecasting techniques, including autoregressive integrated moving average (ARIMA) models, in order to create a model for predicting future air refueling training demand and communicating that demand in terms of aircraft flight hours. Air refueling remains a supply and demand problem in which there will always be more demand than the ability to supply with any realistically sized tanker fleet. The ability to understand, predict, and prepare for increased air refueling demand holds real value to planners, tanker units, and receiver units. This research is a first step in more clearly understanding unsupported air refueling training demand in terms of tanker aircraft flight hours.

To that receiver low on gas, may a tanker not be far away.

Acknowledgments

I would like to express my heartfelt gratitude to a number of vital supporters of this research. First of all, I would like to thank Lt Col Jason Evans for steering me to learn about challenges outside of my airlift community. Additionally, I would like to thank Mr. Ronald Boisjoli and his Air Refueling Liaison Office team at the 618th AOC for coordinating the sharing of data from the Air Refueling Scheduling Tool (ARST) with me for this research and answering countless questions throughout this process. I would like to also thank Maj. Gen. Ricky Rupp for his sponsorship of this project along with his staff's support in better understanding air refueling demand for future application. This paper could not have been completed without the contributions and assistance of the Expeditionary Center's research librarian, Ms. Pamela Bennett Bardot, and that of my AFIT research advisor, Dr. William Cunningham, I am grateful to you both. Finally, and most importantly, I would like to thank my wife and our three lovely children for their support and encouragement as I completed this project over the last year. Your steadfast love and support over many, very long days and missed family time truly made this endeavor possible.

Steven “Bandit” Hawkins

Table of Contents

List of Figures

List of Tables

List of Equations

TIME SERIES FORECASTING OF TANKER TRAINING DEMAND

I. Introduction

"There is nothing we do … without tankers,"
– Former United States Air Force Chief of Staff, General T. Michael Mosely

General Issue

The United States Air Force (USAF) is facing a dramatic decrease in its projected air refueling (AR) tanker fleet. In 1998 the USAF operated 610 boom-equipped air refueling tankers (HQ AMC Office of History, 2018); twenty-one years later that number is down to 455 (U.S. Air Force, 2014b, 2018), the smallest fleet since the early 1960s (Boeing, 2018). The USAF total force currently operates 396 KC-135R/T and 59 KC-10A boom aircraft as the preponderance of the land-based air refueling fleet for the Department of Defense (DoD) (U.S. Air Force, 2014b, 2018). This equates to 455 booms and 501 KC-135 equivalent tankers in terms of overall fuel carrying capacity at the end of the 2018 fiscal year. The USAF is currently in the process of acquiring the new KC-46A, which is a similar size to the older KC-135, as it plans to retire all of the much larger KC-10As and some older KC-135R/Ts (Gertler, 2013, p. 1). The exact future tanker force structure remains uncertain because of a range of factors, including delays in the acquisition of the KC-46A, undetermined divestment timing of the KC-10A, and an aging fleet of KC-135 R/Ts with unknown reliability in the coming decades. What is known is that the target air refueling tanker mix includes a purchase of 179 KC-46A with a 6% fuel capacity increase over the KC-135 R/T and the sustainment of about 300 of the KC-135 R/T aircraft by 2029 (U.S. Air Force, 2016). This understanding is important because it means that tanker availability or the actual number of booms in the air will remain roughly the same as the current status quo while decreasing in fuel capacity of up to 42 KC-135 equivalent aircraft in the next ten years, depicted in Figure 1.

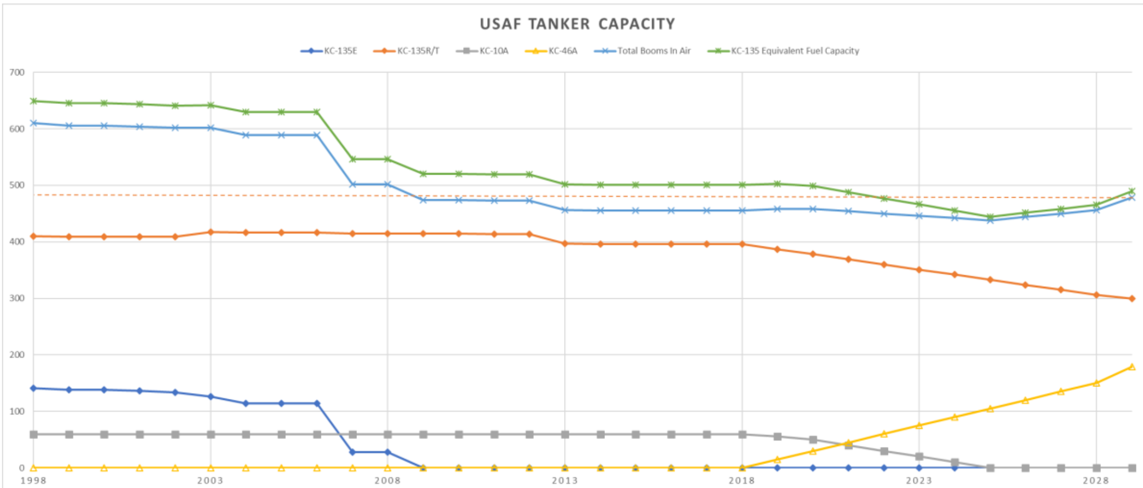

Figure 1: Historical and Projected USAF Tanker Capacity

The above estimates for projected tanker capacity by the researcher are strikingly different from USAF projections for tanker recapitalization released by the Air Force Association just ten years ago as shown in Figure 2 (Knight & Bolkcom, 2008). This shortfall was primarily a result of years of failed recapitalization plans for the tanker fleet. These failed plans therefore have reduced the overall air refueling capacity and that has global implications for the Joint Force's lethality, capability, and mobility as there will certainly be an increase in future air refueling requirements, not a decrease. Therefore, an analysis and time series forecast of the United States Transportation Command (USTRANSCOM) validated tanker training demand would inform decision making on future tanker force structure and mission readiness.

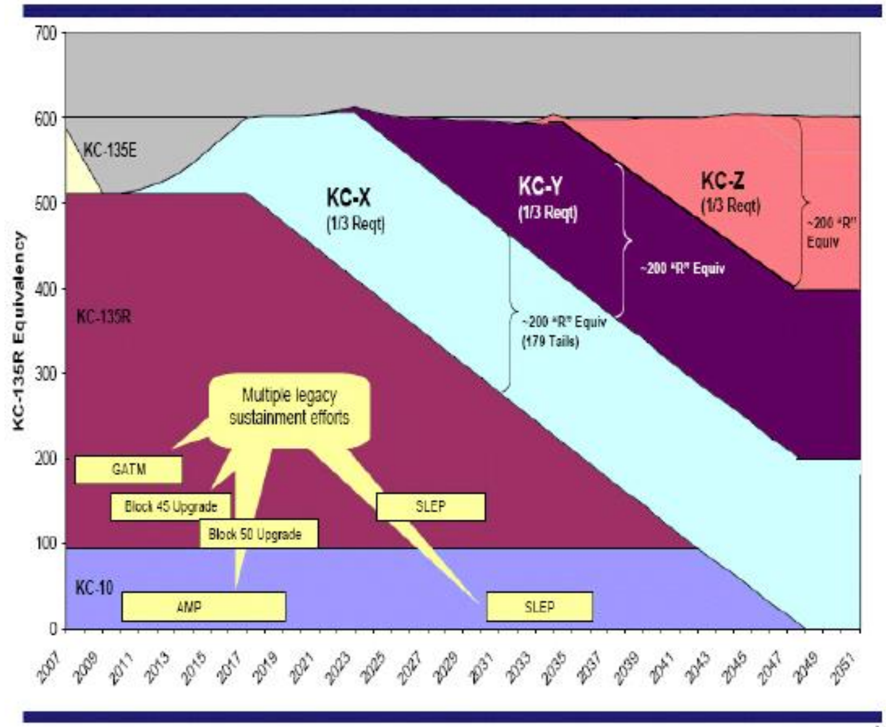

Figure 2: Air Force 2008 Tanker Recapitalization Plan

Source: (Knight & Bolkcom, 2008, p. 22)

This research focuses exclusively on tanker training data versus global operational missions because the training data was readily available over multiple years via the 618 Air Operations Center's (AOC) Air Refueling Scheduling Tool (ARST), "the USAF system of record for matching receiver training air refueling needs to tanker capacity" (HQ AMC/A3O, 2016). In particular, this computer-based system was designed to match lower priority air refueling requests to tanker availability. The database of air refueling requests provides an important demand signal as this lower level of training requests often represents the line between what can and cannot be supported due to several factors including aircraft and crew availability, higher priority taskings, and competing training opportunities. The most valuable aspect of the data set is that the requests are grouped by five different statuses, thus allowing some understanding for the researcher as to how air refueling requests were supported.

Using the currently available data the researcher utilized the definitions below (Table 1) to identify whether an air refueling request was supported by the tanker unit and then modeled all five request statuses in order to assist in future analysis and research.

Table 1: Air Refueling Scheduling Tool Request Statuses

Confirmed	All parties have been notified and have accepted the displayed event's terms. Supported AR by both tanker and receiver units
Bought	The tanker unit has agreed to support the request and is awaiting receiver unit to confirm AR details to confirm the transaction. Supported AR
Filled	The receiver unit may have pre-coordinated with the tanker unit and places some tanker information in the request. Potentially Supported AR
Published	Air Refueling requests that have not been supported and have been entered into the system without tanker details.
Cancelled	Requests that have been cancelled or not supported by tanker units.

The way that the ARST system works is that an air refueling receiver unit enters details of their air refueling request in the computer database. If the user includes specific details of their requested tanker unit, the status is recorded as Filled. If no tanker details are entered the request is entered in the Published status. Once a tanker unit agrees to support the request, the former clicks a "buy" button which enters their tanker unit details and changes the status to Bought. The last step is for the receiver unit to confirm all the entered or updated air refueling details to include air refueling track location, rendezvous time, altitude, etc. by clicking the "confirm" button, thus changing the request status to Confirmed. At any time either the tanker or receiver unit can cancel the request, for any reason, by clicking "cancel," changing the status to Cancelled.

The Confirmed and Bought statuses most likely indicate supported requests that the tanker unit was capable of supporting. In contrast, the Filled and Cancelled statuses currently have no way of determining if they were or were not supported by the tanker unit in the ARST system. This is because a request can be cancelled by either the receiver unit or the tanker unit for a multitude of reasons or never even supported. Currently, the Published category is the clearest representation of unsupported air refueling requests that were never supported by a tanker unit in the ARST. The time series forecasting presented here focused on all five request statuses in order to build representative models in the hopes that future updates to the ARST will provide more fidelity on unsupported requests and reasoning for more detailed future analysis. For example, if either the tanker or receiver unit could annotate why a request was being unsupported or why it was cancelled, it would assist in future capacity and bottleneck analysis. Sorties could be unsupported for common reasons like crew or aircraft availability, but also because the request has insufficient training value, a low designated priority, or any number of other reasons. Such a modification to the ARST would not alter this research but would assist future models and provide relevant and timely feedback to receiver units as to why their request was not supported while providing the receiver units the details to better communicate the impact of a lack of support back to the tanker units, 618th AOC, and Air Mobility Command. An aggregated analysis of supported and unsupported unit types, aircraft types, and request priority was also analyzed and is further discussed in Chapter 5.

Forecasting of air refueling demand and, in particular, training demand allows tanker units and planners to use assets more efficiently and to prepare for periods of increased/decreased tanker capacity. This represents a new way of thinking about tanker scheduling and is proactive rather than reactive; it is a new way to think about the data the 618th AOC is already capturing as part of the ARST program. Moreover, a reliable forecast has second-order effects on unit maintenance and support activities by allowing for more effective operations and an increase in overall readiness. In the end, the whole point of this system is to provide the most training available to requested receivers, thus enabling combat readiness. Air refueling remains a supply and demand problem in which there will always be more demand than the ability to supply with any realistically sized tanker fleet. Therefore, being able to understand better, predict, and prepare for increased air refueling demand holds real value to planners, tanker units, and receiver units.

Focused research on this problem is warranted now because the tanker fleet is not expected to increase in the near future while increased joint and international air refueling requirements are being added, thus increasing both the mission and training demand with the same or fewer air refueling assets for the near future. A specific example includes the next generation fighters that use two to two and a half times as much fuel for much shorter ranges than older legacy aircraft (U.S. Air Force, 2014a, 2015b, 2015a). Therefore, it is critical that planners gain a better understanding of what the training demand data is telling us with regard to request frequency in order to better understand the overall demand market for air refueling.

As previously stated, the current air refueling demand vastly outpaces the USAF's supply. One of the most promising solutions to reduce a segment of the demand remains the inclusion of commercial air refueling tankers which could reduce a segment of the training demand. This would likely include routine receiver training, some test and evaluation support, and non-combat related flight extension. The concept has been studied multiple times starting in 1998 by the US Transportation Command in its "Concept Development Report on Contracted Aerial Refueling" (USTRANSCOM TCJ5, 1999). This report led to research including the Defense Science Board Task Force Report on "Aerial Refueling Requirements" released in 2004 and a RAND Project Air Force report in 2006 entitled "Analysis of Alternatives (AoA) for KC-135 Recapitalization (Defense Science Board, 2004; Kennedy et al., 2006). The barriers to market entry range from policy issues such as no FAA certification for commercial air refueling to a significant financial burden for any organization without a known or guaranteed market size. Therefore, due to numerous technical and bureaucratic challenges, no commercial enterprise currently offers boom air refueling service. Omega Air Refueling Services has operated a small fleet of probe and drogue only aircraft (utilizing a basket at the end of a hose that the receiver connects to versus the USAF's mainstay of a flying boom) since 2001, primarily servicing the US Navy along with some international partners (Omega Air, 2018). The root cause of the problems highlighted above all focus on not truly understanding the DoD air refueling market demand. If air refueling demand could be forecasted, units could plan more efficiently longer into the future and, consequently, more training could be accomplished. Furthermore, if commercial augmentation was, at some future time, supported by the DoD as a viable option to moderate some training demand by supporting less desirable training opportunities, planners could make informed decisions based on a forecast from historical employment. It is this question that this research seeks to answer. Utilizing historical tanker training requests, this research will focus on multiple forecasting techniques in order to create a model for predicting future air refueling training demand.

Problem Statement

Currently, no known forecast method exists for predicting tanker training demand.

Purpose Statement

This research will consider multiple forecasting techniques utilizing historical tanker training requests in order to model future training demand. The future demand will then be estimated using the metric of air refueling flight hours.

Research Questions/Objective

RQ1: Which forecasting technique best predicts the various air refueling training requests for each designated status in the air refueling scheduling tool?

RQ2: Can the number of requested air refueling flight hours be predicted, with up to 90% accuracy annually, for each designated status in the air refueling scheduling tool?

Methodology

This research utilizes time series forecasting to analyze the available ARST data from Jun 2010 – March 2018. The initial data collection only included Published and Confirmed statuses and, because there is no way to identify the other three statuses from that earlier data, the analysis started with the April 2012 data, when all five statuses started to be recorded. The data included several air refueling details about each event which were further aggregated in order to learn more about patterns within the data set. The focus of this time series analysis was on the ARST system's five designated statuses of Confirmed, Bought, Filled, Published, and Cancelled over a regular time period. Annual analysis was rejected due to a limited number of data points; daily time periods were problematic as some statuses had zero requests on a given day resulting in errors in the computations for variance. Following a careful review, the researcher settled on weekly time periods for analysis, the reasoning for which will be discussed in Chapter 4. After running multiple time series analysis techniques for each of the five statuses, the researcher used validation techniques against one year of reserved data to verify the best model for each request status. Once the best forecast model was determined, the researcher applied multiple computational techniques to estimate total annual flight hours by each designated status.

Assumptions/Limitations

As discussed above, the ARST data used for forecasting was restricted to April 2012 through March 2018 due to the limited statuses recorded Between June 2010 and April 2012. The database management of the ARST was changed in March of 2018; which resulted in the collected data from that time forward being unavailable to this researcher. The data that was inputted into the ARST system also limited this research; there is a high probability that not all air refueling training requests were inputted into the system. Additionally, air refueling requests were most likely both supported and unsupported that were never entered into the ARST. However, the system shows stability over time, with a similar number of requests (an average of 21350 per year from 16 Mar 2014 – 10 Mar 2018 with the sums each year remaining within 2-4% of the mean) which leads to the conclusion that it represents a consistent proportion of the air refueling demand over time. Furthermore, future data can be entered into the proposed models thereby improving the model accuracy over time. Lastly, commercial air refueling costing structures were not attainable by the researcher because a boom capable commercial air refueling aircraft does not exist and, therefore, this research does not include a potential cost comparison for future commercial tankers. For this reason, all cost comparisons will use DoD cost structures for current USAF tanker aircraft. When commercial tankers do come to the market, the cost will likely be on a per flight hour basis. Therefore, this research methodology is designed to predict air refueling flight hours to easily apply to future cost comparisons between different tanker aircraft types for future research endeavors and decision making.

Implications

Air refueling demand will always outpace the available supply. This research is designed to find the best model for forecasting future air refueling requests as defined by the ARST status. The analysis outlined here can then use the forecast data to estimate the flight hour duration requested annually in terms of the same ARST statuses. These research questions are a first step to a better understanding of unsupported air refueling training demand in terms of aircraft flight hours. While several metrics are important, this research focused on the use of flight hours in the final analysis because they have a defined value to the DoD. Furthermore, research based on using flight hours can be used in the future in order to either justify increased expenditures to execute unsupported air refueling sorties or to compare USAF flight hour costs to future commercial tanker flight hour costs. When commercial air refueling service is available to the public market, the research and methods outlined here can provide an actual cost comparison given the quantitative forecasting background of the models. Such a comparison is likely to be far more accurate when the commercial tanker costs are available and this research, with its outlined mathematical rigor, is applied to known costs and more detailed needs rather than on educated assumptions about the unsupported air refueling training market using current ARST data.

II. Literature Review

"Aerial refueling will be the biggest shortfall in our Mobility Air Forces"
- Secretary of the United States Air Force, Heather Wilson

Chapter Overview

This research began with a focus on understanding the air refueling challenges and, specifically, on exploring how the USAF fleet capacity has changed and is expected to change in the future. The air refueling fleet clearly has a supply and demand problem, and this researcher started thinking about the problem with respect to how future commercial air refueling tankers could be used to increase the overall supply faster than the USAF could through additional KC-46A procurement. However, as more information was gathered from previous reports, articles, and academic papers, it became clear that the demand side of the equation required analysis before any supply proposals could be explored.

Increased supply will certainly help in the short-term, and commercial air refueling has the potential to solve only certain portions of the overall demand, namely peacetime training purposes for a host of reasons. It is unlikely that any commercial enterprise will invest the time, partner with industry for aircraft acquisition, breakthrough long-standing barriers to market-entry, or even spend what would probably amount to a significant research and development cost without the DoD’s support and a better understanding of the potential market. Therefore, gaining a better understanding of the unsupported air refueling requests (market demand) should be a first step in actually making headway in this process, which began over 20 years ago.

In order to fully understand the market demand of the USAF tanker enterprise, the researcher first focused on literature detailing two decades of recapitalization efforts along with multiple commercial air refueling studies as background to the problem. With that background in mind, the second part of this literature review will detail why demand forecasting is important, in particular, to commercial industries with a focus on the aviation industry. Lastly, the literature review will focus on the background and theory behind the primary forecasting technique used in the following analysis which is the Box-Jenkins methodology for Autoregressive, Integrated, Moving Average (ARIMA) forecast models.

Air Refueling Tanker Recapitalization

Unfortunately, the multiple plans to replace the over 400 “Eisenhower-era KC-135’s” have faced significant challenges, controversies, corruption, and ultimately several failed acquisition programs dating back to the early 2000s (Grismer, 2011, p. 63). These recapitalization plans range from Boeing’s 767 lease proposal that failed in 2001 to Northrop Grumman/European Aeronautic Defense and Space Company’s (EADS) 2008 contract award that was later cancelled under protest from Boeing (Grismer, 2011, p. 63). Then, in February 2011, the Boeing company was finally awarded the contract to build the KC-46A, a contract valued at approximately $35 billion (Gertler, 2013, p. 1). Again, after more delays, the USAF only recently took delivery of the first KC-46A on 10 Jan 2019, while still working to reconcile “major technical problems” (Insinna, 2019). This brief background of the KC-46A’s acquisition is important because, in the over 19 years that it has taken for the USAF to take the delivery of a new tanker, the overall fleet size has decreased by 151 boom tanker aircraft. Moreover, as the USAF acquires the projected 179 KC-46A aircraft, 59 KC-10A and approximately 96 KC-135R/T aircraft are projected to be retired, further reducing the overall tanker fleet (Pawlyk, 2018). Therefore, although the USAF will increase the number of booms available by 24, from 455 to 479 with the current plan, the capacity of fuel available actually decreases by over 2.2M pounds of fuel, from 100,204,000 pounds at the end of fiscal year 2018 to a projected 98,001521 pounds of fuel at the end of fiscal year 2029 (Gertler, 2013, p. 6). As stated above, work on this problem started decades ago with the full knowledge of how important it was to replace this high-demand, low-density resource quickly. It is truly unfortunate that years' worth of studies and recapitalization efforts have languished for so long, resulting in the retirement of more aircraft and only recently resulting in a new tanker acquisition.

The USAF, DoD, and other research agencies have conducted multiple studies on increasing the tanker enterprise over the years. In an effort to highlight some of the most applicable reports presented in chronological order, the researcher started in 1996. That year the GAO conducted a study entitled, "U.S. Combat Air Power: Aging Refueling Aircraft are Costly to Maintain and Operate." The basic conclusions of this report to Congress included an observation that, "although the services' air refueling tanker aircraft meet current needs, satisfying future requirements may be difficult" as the aircraft age and require increasingly more money to operate (Meredith, Stone, Dey, Newell, & Ragsdale, 1996, p. 34). The report also recommended that Congress consider dual use airlift and tanker aircraft in future acquisitions programs (Meredith et al., 1996, p. 35). Reports to Congress such as these highlighted the potential for a deficit of tankers in the future and opened up studies of alternate methods to meet future tanker requirements, including contracted or commercial tankers.

The United States Transportation Command's "Concept Development Report on Contracted Aerial Refueling" was one of the first DoD studies and was completed on 1 Oct 1997, revised on 1 March 1998, and revised again on 21 June 1999. At the time of the original report, the USAF still maintained approximately 610 boom air refueling tankers and had only started talks with industry about a replacement for the KC-135, then roughly 40 years old. Moreover, the 1997 report represents the first time that a commercial organization, Omega Air, Inc., was included in USTRANSCOM working groups. This USTRASCOM report provided in-depth operational, policy, legal, contractual, and cost considerations for what a contract air refueling provider might provide to USTRASCOM, along with the feasibility of such a contract in a "CRAF-like, indefinite delivery, indefinite quantity (IDIQ) contract" (USTRANSCOM TCJ5, 1999). This report focused on using contract tankers to complete only probe and drogue air refueling which represent a small portion of the annual requirement and the only type of air refueling that Omega Air was at the time (and remains) capable of performing. Due to the focus remaining on a small subset of the air refueling requirement (probe and drogue), legal concerns about how commercial tankers would be integrated into combat operations, and concerns over reduced training opportunities for USAF crews the idea was rejected by USTRANSCOM, AMC, and the Joint Staff. In a November 1997 "Report to Congress on Private Sourcing of Airlift of Military Personnel and Cargo," required by the 1996 National Defense Authorization Act, the Office of the Secretary of Defense (OSD) made their position on commercial air refueling clear. OSD highlighted that the combat integration, training, and full range of mission sets make “private-sector sources” for air refueling not suitable (USTRANSCOM TCJ5, 1999, p. 23).

Then, in 2004, the Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics received their requested report from the “Defense Science Board Task Force on Aerial Refueling Requirements.” The report focused on evaluating the current state and future requirements “of the USAF tanker fleet, the corrosion and maintenance issues associated with it, the studies pertaining to the KC-135, and several near-term options that the DoD has with regard to recapitalizing the fleet” (Defense Science Board, 2004, p. 4). Although the report found that the corrosion which resulted from the age and exposure to naturally-occurring elements was being controlled, the cost of that process along with the increasing costs of maintenance, lead to the task force to recommend multiple options for replacing the KC-135 fleet within a “reasonable timeframe” (Defense Science Board, 2004, p. 42). The proposed options included purchasing and converting used aircraft for air refueling, such as available DC-10-30’s, contracting commercial organizations for specific mission sets, and working with large airframe manufactures for procurement of the next generation of air refueling tanker in the near-term (Defense Science Board, 2004, p. 36). Significantly, this report continued to highlight the need for multiple courses of action while considering alternate methods of recapitalization from a traditional acquisition approach. Furthermore, the 2004 Defense Science Board Report directly lead to a follow-on RAND study for an analysis of alternatives for the future of the KC-135.

In 2006 RAND Project Air Force published their report “Analysis of Alternatives (AOA) for KC-135 Recapitalization.” The Under Secretary of Defense for Acquisition, Technology, and Logistics at the time had requested the AOA based on the findings of the 2004 DSB study. Furthermore, the AOA was to include considerations of the tanker requirements from the forthcoming 2005 Mobility Capabilities Requirements Study. The RAND AOA study looked specifically at two research questions: "What is the most cost-effective alternative for recapitalizing the KC-135 fleet?" and "When should the recapitalization assets be acquired?" (Kennedy et al., 2006, p. 16) After reviewing multiple options from new and used commercial-derivatives to new designs to unmanned and even commercial sources, the study recommended new commercial-derivatives in the 300,000 – 1,000,000 pound max gross takeoff weight categories as the most cost-effective (Kennedy et al., 2006, p. 21). Smaller tankers, unmanned, and stealthy tankers were all considered to be not cost-effective in the report. The idea of commercial tankers was rejected primarily on the assumption that all tankers "must be capable of carrying out wartime missions" and, therefore, commercial tankers were considered cost prohibitive due to required defensive equipment (Kennedy et al., 2006, p. 22). Furthermore, the report found that the timing of the recapitalization of the tanker fleet is more dependent on factors such as the risk of catastrophic technical problems or critical maintenance issues rather than overall life-cycle costs of operating the legacy fleet (Kennedy et al., 2006, p. 25). However, the AOA conceded that the purchasing of a new tanker fleet quickly after waiting too long for recapitalization would be much more costly than purchasing it before the legacy tanker fleet wears out (Kennedy et al., 2006, p. 25). The AOA represents the most recent government report highlighting the need to recapitalize the KC-135 tanker fleet and, while it considered alternatives to a normal acquisition process, it failed to look for broad and innovative solutions to a supply and demand problem that was only getting worse by the year.

Following the 2006 AOA report, the tanker recapitalization focus was on the 2008 contract that was initially awarded to EADS, challenged by Boeing, and then rescinded by the USAF. Again, timing is important because the leadup to the contract award coincided with the USAF's decision to retire a majority of the KC-135E model aircraft due to corrosion issues with the engine bolts. Although a small number were upgraded to KC-135R aircraft, approximately 114 were grounded and ultimately sent to long-term storage in Arizona, further reducing the overall fleet (Knight & Bolkcom, 2008, p. 31). The next eleven years were marked by continued calls for recapitalization, followed by yet another contract bid for a new tanker replacement and countless delays that only now have resulted in the USAF taking the possession of the first few KC-46As. This history of multiple working groups, analysis, reports, and recapitalization efforts at multiple levels within multiple organizations is important for the reader to understand that just because an innovative or non-standard solution such as outsourcing was deemed not cost effective or rejected in the past, that conclusion does not necessarily apply in today's environment which is radically different than ten or twenty years ago.

Aviation Demand Forecasting

With a clearer understanding of the environment of dwindling tanker fleet supply and increased demand that the Joint Force faces today, the focus must be on using the tanker fleet as efficiently as possible, and this includes forecasting. The air transportation industry is heavily reliant on forecasting in order to use its assets as efficiently as possible. "Every day, at all levels of management, within all segments of the air transportation industry, decisions are made about what is likely to happen in the future" (Wensveen, 2015, p. 268). Forecasting expected demand can look at multiple different time periods and include an evaluation using several different methods. "The choice of forecasting methods should be based on several factors, including availability of data, accuracy of available data, management sophistication, intended forecast use, and availability of electronic processing" (Wensveen, 2015, p. 270).

For the reasons above, the decision as to the specific forecasting method is almost as important as the data available. There are countless studies of forecasting methods for the aviation industry to include Richard Wickham's 1995 MIT thesis, "Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation." Wickham utilized eleven different models – 3 time series models, 2 regression models, and 6 different pickup models - in his analysis of an eighteen week, short-term air transportation demand forecast for up to an eight week future time period (Wickham, 1995, p. 55). Wickham's research was focused on determining the best forecasting method while varying the forecasting period and the size of the historical data set to see how these changes affected the accuracy of the different methods; he limited his methods to a simple mean and various exponential smoothing techniques. His research findings showed that each model's accuracy decreased as the forecast horizon was increased, yet for short periods of less than four weeks, the results were very similar (Wickham, 1995, p. 110). This is to be expected in that longer forecasts are more challenging, in particular when the data is highly variable as was the case in Wickham's study. However, Wickham's third finding was much more interesting. Some of the models that decreased in accuracy did so as the sample size was increased. He states that "if the focus is shifted to the local booking activity, *micro-trends* can be observed where the data immediately before the point of observation gives some indications of the preceding booking behavior" (Wickham, 1995, p. 111). The point is that, with his short-term demand data and its volatile behavior, the more recent observations were more important to the short-term forecast than the older observations. This is also why his exponential smoothing models resulted in less forecast error; those models weight the more recent observations greater than the older observations. It should be noted that Wickham chose not to include corrections for seasonal fluctuations in his analysis based on the narrowed focus of only eighteen weeks of data, although a larger data set might require such corrections and have different results as "seasonal variation occurs quite naturally in the demand for air travel" (Wickham, 1995, p. 23). Wickham's analysis provides a comparison of multiple forecasting methods and their application to the air transportation industry, while the next article speaks to the applicability of one method of analysis over another to the aviation industry.

In "Predicting Air-Transport Demand," Pitfield focused on a comparison of ARIMA models and regression models for the purpose of forecasting air-transport passengers by the route. Pitfield pointed out numerous cases where regression analysis is used for demand forecasting and identifying explanatory variables in the air transportation industry such as UK domestic passengers who use business services and passenger traffic between airports (Pitfield, 1993, p. 459). The main challenge Pitfield identified with extremely complicated systems like air transportation was that defining all relevant variables accurately for regression models is difficult if not impossible. His paper examines domestic air travel time series data from the UK by applying both ARIMA modeling techniques and regression modeling to the same data set for comparison (Pitfield, 1993, p. 460). The final comparison of the error statistics showed significantly lower error in every category for the ARIMA model versus the regression-based models, leading to Pitfield to conclude that the ARIMA model is a superior forecasting tool over the regression model that struggled with defining the explanatory variables outside of the airlines' control (Pitfield, 1993, pp. 465–466). This article highlights the difficulty in defining and quantifying the multiple variables associated with forecasting in the airline industry. It also highlights the potential accuracy and utilization of ARIMA forecasting models, despite multiple ill-defined variables in such an industry.

Historical Perspective of Time Series Analysis

The concept of time series forecasting is not new. In fact, there are documented cases of merchants utilizing rudimentary forecasting and quantitative reasoning in order to determine expected values for their profits and losses as the markets changed going back to at least mid-17th century and probably earlier (Klein, 1997, p. 1). Though many statisticians have contributed to the early work in time series analysis, George Udny Yule is credited as one of the first statisticians to utilize applied correlation and regression for what he called the "time-correlation problem" (Klein, 1997, p. 222). Yule's work in the 1920s laid the foundation for applied statistics and time series in multiple fields of study by providing an initial understanding of non-stationary data. "Udny Yule's specification of an autoregressive stochastic process was a by-product of his attempts to explain why statisticians, and in particular those who worked with economic and social data, often got strange correlations from time series data" (Klein, 1997, pp. 264–265). In doing so, "Yule's specification of the autoregressive process…was to become one of the most common tools of forecasting with univariate time series models" (Klein, 1997, p. 270). This work enabled others, including Herman Wold who, in 1938, defined the discrete stationary process, which is a stochastic or random process. "Wold spelled out the autoregressive, moving-average, modeling approach that would be used in the coming decades to investigate stochastic processes, his work spurred further development in spectral analysis, he gave justification for the use of sum of squares in the analysis of stationary times series, and he put correlograms front and center as the key means of model specification" (Klein, 1997, p. 289). These statisticians are highlighted for their foundational work that was continued by the statisticians George Box and Gwilym Jenkins, who popularized the use of the ARIMA process for economic and business process in their seminal work "Time Series Analysis: Forecasting and Control" (1970).

AutoRegressive, Integrated, Moving Average (ARIMA)

"ARIMA processes are mathematical models used for forecasting" (Hyndman, 2001, p. 1). Due to George Box and Gwilym Jenkins' extensive study of these models and their applicability to time series forecasting, the ARIMA processes are often called Box-Jenkins Methods. "The ARIMA approach to forecasting is based on the following ideas: 1) The forecasts are based on linear functions of the sample observations; 2) The aim is to find the simplest modes that provide an adequate description of the observed data" (Hyndman, 2001, p. 1). This concept of the fewest parameters possible is also known as parsimony. "Parsimony may often be achieved by representation of the linear process in terms of a small number of autoregressive and moving average terms" (Box & Jenkins, 1970, p. 46). The intent is to reduce the model's complexity and, therefore, to avoid the potential of over-fitting the model with excessive parameters. The ARIMA process is based in three parameters: autoregressive (AR), integrated (I) , and moving average (MA); each are explained in detail with their applicable equations in Chapter 3 (Commandeur & Koopman, 2007, p. 122).

By putting all three parameters together, the ARIMA process forms a powerful forecasting tool with a defined mathematical structure allowing for seasonal and multivariate time series forecasting (Hyndman, 2001, p. 2). The research presented in this paper did not focus on multivariate or forecasting with multiple variables; however, it does focus on multiple seasonal time series independently for each of the five ARST defined statuses. Seasonal analysis uses the same basic structure as ARIMA but adds an extra set of AR, I, and MA parameters to model the seasonal elements of the time series, each characterized by a capital letter. Therefore, the shorthand for a seasonal ARIMA is (p,d,q)(P,D,Q). In order to determine the value of those parameters, Box and Jenkins developed a methodology to approach the process.

The Box-Jenkins Methodology

As discussed above, the Box-Jenkins modeling methodology was developed by two statisticians, George Box and Gwilym Jenkins, as a way to apply an organized and iterative approach for determining the correct ARIMA values given a forecasting problem. "Box-Jenkins modeling involves identifying an appropriate ARIMA process, fitting it to the data, and then using the fitted model for forecasting" (Hyndman, 2000, p.

1). Box and Jenkins applied the process through an iterative three-step process that they called model selection, parameter estimation, and model checking; in recent years the first step of data preparation and last stage of the model application or forecasting has been added (Hyndman, 2000, p. 1). The full process, as described by Makridakis, Wheelwright, and Hyndman, is presented in Table 2 below.

Table 2: The Box Jenkins Methodology for ARIMA models

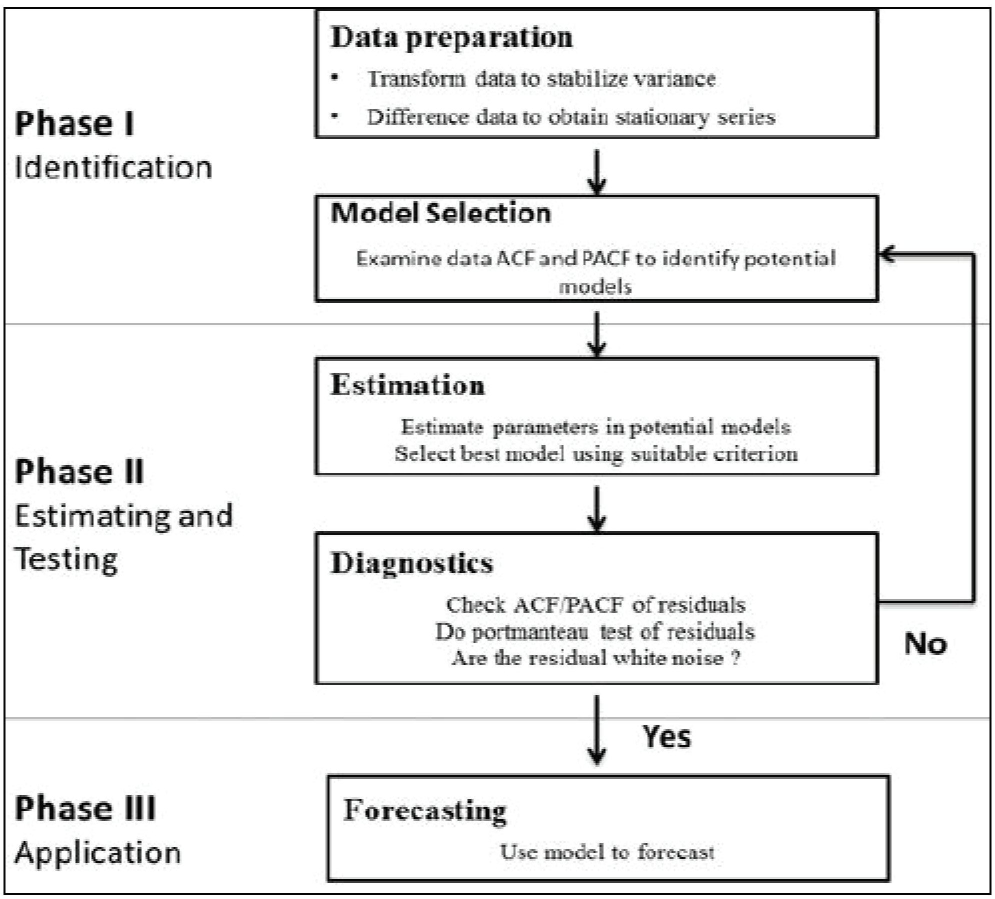

Source: (Makridakis, Wheelwright, & Hyndman, 1998, p. 314)

Data Preparation

Data transformations such as taking the logarithm, square root, or other mathematical transformations of the data in order to stabilize variance are often required for business or economic data as the first step in data preparation (Hyndman, 2000, p. 1). Data may then be differenced which involves subtracting successive observations in order to remove patterns such as trend or seasonality thus increasing ease of modeling the data without unpredictable trends or fluctuations (Hyndman, 2000, p. 1). The data transformation and differencing techniques that are included in the data preparation steps are not required with every data set; however, a thorough analysis of the data set is required before a model is selected to determine what type of preparation, if any, is required.

Model Selection

After the data is prepared, ARIMA models are determined based on the use of autocorrelation functions (ACF) and partial autocorrelation functions (PACF) of differenced data series to determine the best data fit (Hyndman, 2000, p. 1). "If the future values can be described only in terms of a probability distribution, the time series is said to be non-deterministic or simply a statistical time series" (Box & Jenkins, 1970, p. 24).

Parameter Estimation

This involves determining the ARIMA coefficients of (p,d,q) or (p,d,q)(P,D,Q) for seasonal models which best fit the series data (Hyndman, 2000, p. 1). There are multiple methods for determining the goodness-of-fit of a model, including the Akaike's Information Criterion and Schwarz Bayesian Information Criterion, both of which are utilized in the research below.

Model Checking

Checking the model is an iterative process of identifying where or how the model is not as good of a fit to the data series and re-accomplishing the model selection, parameter estimation, and model checking steps until the best fit is obtained (Hyndman, 2000, p. 1).

Forecasting

Utilizing the best model to accomplish a forecast is the end result of the process and is often accomplished via a computer program for large data sets (Hyndman, 2000, p. 1). Once the data is forecast, it can be compared to reserved data observations or future observations in order to validate the model via a multitude of error statistics. Validation is important because all the goodness-of-fit measures used in the iterative model checking process are only designed to determine how well the model fits the data used to forecast. Whereas, the validation process is used to determine how well the end result accurately predicts future events, the end goal of forecasting.

Summary

This literature review presented a background on over twenty years of USAF tanker recapitalization efforts from failed contracts to numerous reports analyzing tanker recapitalization strategies. All of these efforts were focused on producing a new tanker or tanker alternative before the tanker fleet was reduced by aging aircraft and the increasing cost of maintenance. Unfortunately, multiple delays have reduced the DoD's options and strained the air refueling system, only recently producing a new tanker in the KC-46A. Efficient use of the remaining assets to include demand forecasting strategies that have been used in the air transportation industry for decades may alleviate some of the strain. Several articles on different forecasting strategies were presented along with a background on time series analysis and ARIMA models. This background helps to explain the power and utilization of ARIMA models specifically for the air transportation industry. Finally, a discussion of the Box-Jenkin's methodology was presented in preparation for the mathematical methodology in Chapter 3.

III. Methodology

"Prediction is very difficult, especially if it is about the future" - *Niels Bohr*

Chapter Overview

This research employed time series forecasting techniques, including autoregressive integrated moving average (ARIMA), models developed by statisticians George Box and Gwilym Jenkins for business and economic data application in the 1970s (Nau, 2014b). Time series forecasting develops a model based on past observations of the same variable over equally spaced time periods in order to describe an underlying relationship (Zhang, 2003). Such a model can then be used to predict future outcomes of that variable over future time periods of the data series. This chapter will explain how the researcher analyzed the available tanker training ARST data which is comprised primarily of priority 3 and 4 air refueling training missions versus operational missions and is focused on the air refueling request status portion of the data set. Next, this chapter will include a brief discussion on data transformations to provide stationary time series including differencing and logarithmic transformations for application in multiple forecasting methods. Lastly, the chapter will focus on the description and understanding of the various time series forecasting methods that were utilized in increasing complexity to analyze the available data.

Data Scope

As previously stated, this researcher's data source focused on tanker training data received from the 618 Air Operations Center's (AOC) Air Refueling Scheduling Tool. 136,466 rows of data from 30 Jun 2010 to 10 Mar 2018 were initially provided. The data provided also contained information about the air refueling requests including receiver organization, aircraft type, fuel required, priority, duration, etc. for a total of 22 database columns. As stated in Chapter 1, the request status is most effective in determining whether the air refueling request was supported. Due to differences in how the request status was originally recorded prior to April 2012, the data set was scoped to start on 08 Apr 2012 and end on 10 March 2018. These dates were chosen to correspond to the Sunday through Saturday weekly schedule, which will be more thoroughly explained in Chapter 4. This subset of the data left 117,446 data points for analysis - 95,906 for forecasting and 21,540 for validation. Data after 10 Mar 2018 remained unavailable to the researcher due to a system software change; however, it is still being collected and should ultimately be used for further validation and updates to the presented forecast models.

Data Transformations

Many time series forecasting techniques depend on a stationary time series, meaning that the series is not dependent on the specific observed time (Hyndman & Athanasopoulos, 2018). Therefore, a time series with a significant trend or seasonal component would be non-stationary as observed at different times of the year. A series with no predictable pattern like “white noise” would be stationary and should have “constant statistical properties such as mean, variance, and autocorrelation over time” (Nau, 2018). This is important because it can make the series easier to forecast if the statistical properties remain constant.

One technique to achieve a stationary series is to take the difference of the time series. This is accomplished by subtracting the value of the time Y at time period t from Y at time period t – 1; if this operation produces a stationary series that is random and not autocorrelated, then it is also called the random walk model described below (Nau, 2018). It is also possible to take the second difference of a time series by taking the first difference again of the first difference just taken; however, this can easily cause an unintended error if not properly analyzed. Other data transformations include mathematical operations to stabilize the variance in a time series, thereby increasing the ability to predict future events as the subsequent observances would have a smaller range of highs and lows from the mean. For example, taking the logarithm (Log 10) or natural log (LN) of the data set can reduce variance as differencing can stabilize the mean thus reducing data trend and seasonality (Hyndman & Athanasopoulos, 2018).

Time Series Forecasting Methods

There are several different methods for forecasting data over time. Those in this section are some of the most commonly used and are presented from the simplest to the more complex; several concepts are reliant on one-another.

The Naïve Method

Naïve forecasts simply predict future outcomes based on the last observation recorded (Hyndman & Athanasopoulos, 2018). Simple forecasts such as these can serve as an important benchmark for comparison against more complex forecasting methods that may be prone to more error. The naïve model should represent the worst case of an error a researcher is willing to accept in any given forecast. These models are quick and require little to no computing power to produce as they assume that the next data point is equal to the last observed point (Singh, 2018). The naïve method can be surprisingly adequate when applied to a stable or limited data series. The equation for the naïve method is:

Equation 1: The Naïve Method

$$\boldsymbol{\hat{Y}_t = Y_{t-1}} \tag{1}$$

Where

$\hat{Y}_t$ = the forecast value at time t

The Random Walking Method

The random walking method is the first difference of the naïve method, which outputs the difference between successive data points from the original observations, as shown below (Hyndman & Athanasopoulos, 2018).

Equation 2: First Difference

$$\boldsymbol{\acute{Y}_t = Y_t - Y_{t-1}} \tag{2}$$

Assuming this output is stationary, the random walking model can be re-written as:

Equation 3: The Random Walking Method

$$\boldsymbol{\hat{Y}_t = Y_{t-1} + \varepsilon_t} \tag{3}$$

Where

ε_t = accounts for random error or white noise

This random walking model has the same benefit presented for the naïve method – a benchmark for the worst case of error – yet it is designed for non-stationary data.

The Simple Average Method

Instead of using the last value to predict the next value, the simple average method uses the average of all previous values observed in order to predict the expected value of a future value (Singh, 2018). This method is best suited for series that maintain a nearly constant mean and will not account well for data with an upward or downward trend. The equation for the simple average method is below and encompasses a simple arithmetic average of all observed values that number a total of p.

Equation 4: The Simple Average Method

$$\hat{Y}_{t+1} = \frac{1}{p}\sum_{t=1}^{p} Y_t \tag{4}$$

Where

p = previous values

The Simple Moving Average Method

A potentially more useful application building upon this method is the simple moving average method, which allows for the use of a specified subset of the data often but not always the more recent values to account for variation in earlier data. Another way to think about this method is that it considers a "sliding window" of data that is specified by k for a stationary data series in the below equation (Singh, 2018).

Equation 5: The Simple Moving Average Method

$$\hat{Y}_{t+1} = \frac{Y_t + Y_{t-1} + Y_{t-2} + \ldots + Y_{t-k+1}}{k} \tag{5}$$

Where

$\hat{Y}_{t+1}$ = the forecast value at time t + 1

It is important to note that if k = 1, the output is the same as the naïve method above; the other extreme is that if k = p or all previous observations, the output would be the same as the simple average or mean method above. Furthermore, it is possible to weigh past observations differently, resulting in the weighted moving average method (Singh, 2018).

The Simple Exponential Smoothing (SES) Method

In an effort to bridge the gap between the simple average and weighted moving average methods, simple exponential smoothing takes into account all data by using weighted averages that decrease exponentially to account more for recent, potentially relevant data, and less for older, potentially irrelevant, data (Singh, 2018). This process weighs recent data the most, a benefit of a simple moving average, yet considers all data in the set, unlike moving average. In order to define the equation for the simple exponential smoothing method several new terms must be introduced. First, α is a "smoothing constant" that is between 0 and 1 (Nau, 2014a, p. 8). It is important to note that, once again, if α=1 then the SES model will equal the naïve model (or random walking if differenced) and if α=0 then the SES will equal the simple average or mean model. The second issue is to define the series L that represents a level or local mean of the series "computed recursively" from the previous data as shown in equation 6 (Nau, 2014a, p. 7).

Equation 6: Level Mean of the Series

$$L_t = \alpha Y_t + (1-\alpha)L_{t-1} \qquad (6)$$

Where

$$0 \leq \alpha \leq 1$$

The estimated local mean at time t is calculated by "interpolating between the just-observed value and the previous estimated level, with weights of α and 1-α, respectively" (Nau, 2014a, p. 8). Larger values of α will increase the weight of more recent observations while, conversely, lower values will increase the weight of older observations. A significant benefit of the SES model is that α can be varied with new data input to minimize forecast error. Assuming the forecast series has no trend, the forecast for time t+1 can be estimated by the local mean at time t in equation 7.1 (Nau, 2014a, p. 8). This leads to the simple exponential smoothing equation by substituting equation 6 in for L_t:

Equation 7: The Simple Exponential Smoothing (SES) Method

$$\hat{Y}_{t+1} = L_t \tag{7.1}$$

$$\hat{Y}_{t+1} = \alpha Y_t + (1 - \alpha)\hat{Y}_t \tag{7.2}$$

Where

$$0 \leq \alpha \leq 1$$

Double (Brown's) Linear Exponential Smoothing Method

While simple moving average and simple exponential smoothing methods rely on data without a trend, Brown's linear exponential smoothing model computes both a level and trend denoted as L_t and T_t, respectively (Nau, 2014a, p. 16). To accomplish this, the series S at time t is exponentially smoothed using SES and the same α, once for S´ and a second time for S''.

Equation 8: Double (Brown's) Linear Exponential Smoothing Method

$$S'_t = \alpha Y_t + (1 - \alpha)S'_{t-1} \tag{8.1}$$

$$S''_t = \alpha S'_t + (1 - \alpha)S''_{t-1} \tag{8.2}$$

$$L_t = 2S'_t - S''_{t-1} \tag{8.3}$$

$$T_t = (\alpha / (1 - \alpha))(S'_t - S''_{t-1}) \tag{8.4}$$

$$\hat{Y}_{t+k} = L_t + kT_t \tag{8.5}$$

Where

S'_t = Single smoothed series at time t

S''_t = Double smoothed series at time t

L_t = Estimated level at time t

T_t = Estimated trend at time t

α = Smoothing constant $0 \leq \alpha \leq 1$

Linear (Holt's) Exponential Smoothing Method

Holt's linear exponential smoothing method is similar to Brown's in that it also estimates level and trend. However, it employs two different smoothing parameters, one for level (α) and one for trend (β), allowing those estimates to vary at independent rates thus fitting more data patterns (Nau, 2014a, p. 16).

Equation 9: Linear (Holt's) Exponential Smoothing Method

$$L_t = \alpha Y_t + (1-\alpha)(L_{t-1} + T_{t-1}) \tag{9.1}$$

$$T_t = \beta(L_t - L_{t-1}) + (1-\beta)\, T_{t-1} \tag{9.2}$$

$$\hat{Y}_{t+k} = L_t + kT_t \tag{9.3}$$

Where

L_t = Estimated level at time t

T_t = Estimated trend at time t

α = Smoothing constant $0 \leq \alpha \leq 1$

β = Trend smoothing constant $0 \leq \beta \leq 1$

Damped Trend Linear Exponential Smoothing Method

A downside of Holt's linear exponential smoothing method is that it forecasts a constant trend infinitely into the future resulting in over or under forecasting over longer periods of time (Hyndman & Athanasopoulos, 2018). In order to combat this forecasting problem, the damped trend method includes a parameter (ϕ) to adjust the trend line toward 0 slope as time periods increase. If $\phi = 1$ the damped trend method will equal the Holt's linear method, while possible values for ϕ range from $0 < \phi < 1$.

Equation 10: Damped Trend Linear Exponential Smoothing Method

$$L_t = \alpha Y_t + (1 - \alpha)(L_{t-1} + \phi T_{t-1}) \tag{10.1}$$

$$T_t = \beta(L_t - L_{t-1}) + (1 - \beta)\, \phi T_{t-1} \tag{10.2}$$

$$\hat{Y}_{(k)} = L_t + \sum_{i=1}^{k} \phi^i\, T_t \tag{10.3}$$

Where

$$\mathrm{L_t} = \text{ Estimated level at time t}$$

$$\mathrm{T_t} = \text{ Estimated trend at time t}$$

α = Smoothing constant $0 \leq \alpha \leq 1$

β = Trend smoothing constant $0 \leq \beta \leq 1$

ϕ = Dampening parameter $0 < \phi < 1$

Seasonal Exponential Smoothing Method

Up to this point, methods have built on one another to address mean, moving average, and exponential smoothing. In particular, the exponential smoothing models have increased in complexity to accommodate series with different trends. The seasonal exponential smoothing addresses those series with a seasonal or cyclical component but no overall trend increasing or decreasing by utilizing a level and seasonal term (Hyndman & Athanasopoulos, 2018).

Equation 11: Seasonal Exponential Smoothing Method

$$L_t = \alpha(Y_t - S_{t-p}) + (1-\alpha)L_{t-1} \quad (11.1)$$

$$S_t = \delta(Y_t - L_t) + (1-\delta)S_{t-p} \quad (11.2)$$

$$\hat{Y}_{(k)} = L_t + S_{t-p+k} \quad (11.3)$$

Where

$\mathrm{L_t} = \text{ Estimated level at time t}$

$\mathrm{S_t} = \text{ Estimated length of seasonality at time t}$

α = Smoothing constant $0 \leq \alpha \leq 1$

δ = Seasonality smoothing constant $0 \leq \delta \leq 1$

The Holt-Winters Seasonal Method

In an effort to put all of these components together into one method, Holt and Winters developed a triple exponential smoothing method with a term for level (L_t), trend (T_t), and seasonal (S_t) components, along with respective independent smoothing parameters (α, β, δ) (Hyndman & Athanasopoulos, 2018). There are two different representations depending on the series behavior of the seasonal component. "The additive method is preferred when the seasonal variations are roughly consistent through the series, while the multiplicative method is preferred when the seasonal variations are changing proportional to the level of the series" (Hyndman & Athanasopoulos, 2018). Data samples with little trend or seasonality will result in similar outputs via either method. The difference as presented in the equations below is that the additive form is expressed in absolute terms while the multiplicative is expressed in relative terms (Hyndman & Athanasopoulos, 2018)

Equation 12: The Holt-Winters Seasonal Method - Additive

Additive Seasonality – Constant Amplitude

$$L_t = \alpha(Y_t - S_{t-p}) + (1-\alpha)(L_{t-1} + T_{t-1}) \quad (12.1)$$

$$T_t = \beta(L_t - L_{t-1}) + (1-\beta)\, T_{t-1} \quad (12.2)$$

$$S_t = \delta(Y_t - L_t) + (1-\delta)S_{t-p} \quad (12.3)$$

$$\hat{Y}_{(k)} = L_t + kT_t + S_{t-p+k} \quad (12.4)$$

Where

$\mathrm{L_t} =$ Estimated level at time t

$\mathrm{T_t} =$ Estimated trend at time t

$\mathrm{S_t} =$ Estimated length of seasonality at time t

α = Smoothing constant $0 \leq \alpha \leq 1$

β = Trend smoothing constant $0 \leq \beta \leq 1$

δ = Seasonality smoothing constant $0 \leq \delta \leq 1$

p = Number of seasons per year

Equation 13: The Holt-Winters Seasonal Method - Multiplicative

Multiplicative Seasonality – Constant Cycles (in overall percentage teams)

$$L_t = \alpha \frac{Y_t}{S_{t-p}} + (1-\alpha)(L_{t-1} + T_{t-1}) \quad (13.1)$$

$$T_t = \beta(L_t - L_{t-1}) + (1-\beta)\, T_{t-1} \quad (13.2)$$

$$S_t = \delta \frac{Y_t}{L_t} + (1-\delta)S_{t-p} \quad (13.3)$$

$$\hat{Y}_{(k)} = L_t S_{t-p+k} \quad (13.4)$$

Where

$$L_t = \text{Estimated level at time t}$$

$$T_t = \text{Estimated trend at time t}$$

$$S_t = \text{Estimated length of seasonality at time t}$$

α = Smoothing constant $0 \leq \alpha \leq 1$

β = Trend smoothing constant $0 \leq \beta \leq 1$

δ = Seasonality smoothing constant $0 \leq \delta \leq 1$

p = Number of seasons per year

ARIMA (The Box-Jenkins Methodology)

Auto-Regressive Integrated Moving Average (ARIMA) models are a variation of discrete-time filtering methods created by electrical engineers in the 1930s and 40s. Statisticians George Box and Gwilym Jenkins popularized these models starting in the 1970s for application to business and economic data (Nau, 2014b). Almost all of the previously covered methods can be described by an ARIMA model, as shown in Table 3 below. The application of the ARIMA model is broken down into three parts – auto-regressive (AR), integrated (I), and moving average (MA), often with the standard notation p, d, q to explain the order and degree of each segment of the model (Nau, 2014b). An autoregressive (AR) term is a forecast of the indicated variable utilizing a linear regression of a number of past values of the variable (Hyndman & Athanasopoulos, 2018). Therefore, the AR term is focused on fitting the forecast to previous observations in the series. The integrated (I) term is a factor when the series is not stationary and must be differenced to stabilize the mean or reduce trend and seasonality. ARIMA models rarely need to be differenced more than twice, and over-differencing can output strong negative autocorrelation along with a strong moving average signature (Nau, 2014b). Lastly, the moving average (MA) term utilizes past forecast errors in the series and regression for forecasting future values (Hyndman & Athanasopoulos, 2018). Application of a specific ARIMA model is an iterative process to develop the best model, involving trial and error in order to minimize forecast error displayed in the autocorrelation (ACF) and partial autocorrelation (PACF) residual plots (Nau, 2014b).

Table 3: Forecasting Method and ARIMA Equivalence

Forecasting Method	Equivalent ARIMA
Mean Model / White Noise	ARIMA (0,0,0) + c
Random Walk Model	ARIMA (0,1,0)
Simple Exponential Smoothing	ARIMA (0,1,1)
Double (Brown) Linear Exponential Smoothing	ARIMA (0,2,2)
Linear (Holt) Exponential Smoothing	ARIMA (0,2,2)
Damped Trend Linear Exponential Smoothing	ARIMA (1,1,2)
Seasonal Exponential Smoothing	ARIMA (1,1,p+1)(0,1,0)p
Holt-Winters Additive Seasonal Method	ARIMA (1,1,p+1)(0,1,0)p

Source: (Jones & Arnold, 2019)

Autoregressive (AR)

The AR part of the model is often written as *p* and "describes how each observation is a function of the previous *p* observations" (Hyndman, 2001, p. 1). The equation for $p > 1$ is:

Equation 14: Autoregressive Standard Equation

$$Y_t = c + \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \ldots + \phi_p Y_{t-p} + \varepsilon_t \tag{14}$$

Where

$$\mathrm{Y_t} = \text{Observed value at time t}$$

$$Y_{t-1} = \text{The previous value at time t}$$

$$\varepsilon_t = \text{Random error}$$

c = Constant

ϕ_p = Constant

Integrated (I)

The I part of the model is often written as *d* and "determines whether the observed values are modelled directly, or whether the differences between consecutive observations are modelled instead" (Hyndman, 2001, p. 1). This term is used to correct for non-stationary series, that is those with trend or seasonality, which is quite common in business and economic data. "Forecasting has been of particular importance to industry, business, and economics where many times series are often represented as non-stationary and, in particular, as having no natural mean" (Box & Jenkins, 1970, p. 7). Series are rarely differenced more than twice to obtain stationarity.

Moving Average (MA)

The MA part of the model is often written as *q* and "describes how each observation is a function of the previous *q* errors" (Hyndman, 2001, p. 1). The equation for $q > 1$ is:

Equation 15: Moving Average Standard Equation

$$Y_t = c + \Theta_1 \varepsilon_{t-1} + \Theta_2 \varepsilon_{t-2} + \ldots + \Theta_q \varepsilon_{t-q} + \varepsilon_t \tag{15}$$

Where

$\mathrm{Y_t} = \text{Observed value at time t}$

$\varepsilon_{t-q} = \text{Random error at a previous time t} - \text{q}$

ε_t = Random error at time t

c = Constant

Θ_q = Constant

Seasonal ARIMA

For data series that exhibit seasonal fluctuations, ARIMA can be used to model the data in a way similar to the seasonal exponential smoothing or the Holt-Winters methods, which include a seasonal component. In order to apply seasonal ARIMA, or SARIMA, three more terms are added - seasonal auto-regressive (P), seasonal integrated (D), and seasonal moving average (Q). The seasonal function thus is expressed as ARIMA(p,d,q)(P,D,Q) where the uppercase and lowercase terms are independent of one another and have recommended values based on residual ACF/PACF outputs (Nau, 2014b).

Time Series Forecasting Model Approach

The researcher utilized JMP 12.0 statistical software developed by SAS to analyze the 2012 – 2018 ARST statuses of Confirmed, Bought, Filled, Published, and Cancelled grouped by seven-day weeks (Sunday thru Saturday). Weekly analysis was chosen for seasonal reasons apparent in the data set. Some days resulted in zero requests (major holidays, etc.) and daily forecasting with a zero value for the variable of interest results in software calculation errors amongst the various formulas. An initial exploration into monthly forecasting resulted in a less than optimal sample size with only 72 months of data, 12 of which were to be withheld for validation. Furthermore, there are some monthly cycles and some seasonality, but the true cycles appeared to be weekly, as requests on the weekends were much less than mid-week (and often non-existent thus skewing data and making an analysis of a smaller time scale not useful with the techniques evaluated). Therefore, weekly aggregation of data allowed for 309 observations, 52 of which were withheld for validation.

Each of the five series were plotted in JMP, log (ln) transformed and differenced to achieve stationary series data, before applying the above exponential smoothing and ARIMA methods in order to determine the best models to compare to the withheld 52-weeks of data and to determine the best forecast model moving forward for each request status. An explanation of the output results and how they were used to determine the best possible models for validation is included in the next chapter. Furthermore, analysis and comparison of error for each status is included with a graph depicting the forecast, validation data, and confidence interval for each request status.

Measures of Error and Model Comparison

The output from the JMP program produces multiple measures to compare against one another by defining error in multiple ways. For example, the JMP model comparison output includes degrees of freedom, variance, Akaike's 'A" Information Criterion (AIC), Schwartz's Bayesian Information Criterion (SBC), RSquare, -2LogLikelihood (-2LogLH), Mean Absolute Percentage Error (MAPE), and Mean Absolute Error (MAE). Whereas the AIC and SBC are used primarily for model comparison, the MAPE and MAE are measures of error to determine how well the model actually represents the historical observations. Each of these measures will be discussed below in more detail; for now, it is important to note that none of these measures can be certain in predicting future values. For that, the model predictions must be compared to data reserved for model validation.

Degrees of Freedom (DF)

Degrees of Freedom are measured from n (the number of observations in the series) – k (the “number of fitted parameters in the model”) (JMP Support, 2018). The model degrees of freedom are normally reduced primarily by the number of forecast periods which are included in the value of k. However, other data alterations like differencing, AR, and MA terms also increase model parameters and reduce the degrees of freedom.

Variance

An estimate of variance in the model is calculated by dividing the sum squared error (SSE), which is calculated by summing the squared residuals by the degrees of freedom, SSE / (n-k). This results in a sample estimate of the variance or the random changes in the model (JMP Support, 2018).

RSquare

Also known as the coefficient of determination (R^2) is calculated by 1-SSE/SST where SSE is the “sum of the squares of the residuals” and SST is the total sum of the squares (JMP Support, 2018).

Equation 16: Sum Squared Error

$$\mathrm{SSE} = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \tag{16}$$

Equation 17: Total Sum of Squares

$$\mathrm{SST} = \sum_{i=1}^{n} (y_i - \bar{y}_i)^2 \tag{17}$$

RSquare is presented as a correlation between the explained variance and total variance of Y values in a squared form to provide a proportion of the variation that can be explained (Makridakis et al., 1998, pp. 199–200). This makes RSquare an especially useful tool when using regression and understanding. For example, it shows how linear variance in y is explained by x. However, it can be a more difficult tool with time series forecasting; which can have an extremely low RSquare's that must be taken into context with other forecasting methods as one piece of the puzzle. One aspect to watch out for is a negative Rsquare resulting from a SSE that is larger than the SST; this result may indicate a poorly fitting model (JMP Support, 2018).

-2LogLikelihood

The above measures focus primarily on variance while the next three focus on model comparison. -2LogLikelihood, or -2LogLH as it is depicted in JMP, is an iterative optimization method to find the maximum likelihood of the probability density function (PDF) by observing the actual sample observations given a defined set of parameters (Makridakis et al., 1998). The likelihood function is applied to time series forecasting which also uses historical observational data modeled after a known probability distribution function for a given set of parameters or variables (Myung, 2003, p. 92). The process uses a method called maximum likelihood estimation (MLE) which derives a probability distribution that best fits the data observation in order to determine other selection criteria such as AIC and SBC (Myung, 2003, p. 93). The actual calculation varies depending on the projected PDF and potential for local and global maximums. In practice, -2LogLikelihood is actually taking negative two times the natural log of the likelihood function which is then evaluated "at the best fit parameter estimates with the smallest value representing a better fit model" (JMP Support, 2018). The major drawback of this parameter estimation technique is that -2LogLikelihood fails to penalize models with excessive parameters. This can result in 'overfitting' a model with additional parameters in order to result in a low -2LogLikelihood value, yet there is no corresponding value in how well the model fits future data; this is where AIC and SBC have value.

Akaike's "A" Information Criterion (AIC)

The Akaike's "A" Information Criterion (AIC) is a measure of a model's goodness-of-fit which can be used to decide between competing forecast models (Makridakis et al., 1998, p. 589). This is because the AIC estimates the quality of the model by including twice the number of model parameters in the equation, thereby reducing the chance of 'overfitting' or 'underfitting' the model with the lowest AIC representing the best fit model.

Equation 18: Akaike's "A" Information Criterion (AIC)

$$AIC = -2\text{LogLikelihood} + 2\text{k} \tag{18}$$

Where

$k =$ Number of estimated parameters in the model

Schwartz's Bayesian Information Criterion (SBC)

The Schwartz's Bayesian Information Criterion (SBC), like AIC, is a measure of a model's goodness-of-fit which can be used to decide between competing forecast models; it is also called an order selection criteria (Makridakis et al., 1998, p. 592). Like AIC, SBC estimates model quality while reducing the chance of 'overfitting' or 'underfitting' with a more complex model. SBC includes both the number of parameters and number of observations in the equation, thereby producing a slightly different result than AIC and often an overall less complex model which should be prone to less error.

Equation 19: Schwartz's Bayesian Information Criterion (SBC)

$$SBC = -2\mathrm{LogLikelihood} + \mathrm{k} * \ln(\mathrm{n}) \tag{19}$$

Where

$$\mathrm{k} = \text{ Number of estimated parameters in the model}$$

n = Number of observations in the model

With a complete understanding of the above measures of model goodness-of-fit, the researcher utilized JMP software for the model comparisons below. The software output is rank ordered by AIC followed by SBC, both computed using an MLE approach for calculating best fit model parameters using -2LogLikelihood. This approach was used for both forecast model selection and order selection criteria when deciding between multiple ARIMA models.

Model Validation

The above measures of goodness-of-fit are focused on how well each model compares to the historical data it is based upon and not how well the model might fit future data. 52-weeks of data were reserved from each of the five statuses being forecasted for the purpose of forecast validation. The error statistics MAE and MAPE can be used to see which forecasting technique resulted in the least error and should be considered for modeling future requests for that type of status. Each of the five demand statuses was modeled independently and, therefore, have slightly different models for best fit of the data. It also should be noted that all five statuses were transformed by taking the natural log of the data in order to reduce variance for each of the forecasting methods.

Mean Absolute Error (MAE)

Although there are several ways of communicating forecast error, this research will focus on two of the most common, mean absolute error (MAE) and mean absolute percentage error (MAPE). Each of these measures of error represents the difference between the actual observation and the prediction or forecast for that observation. The first, mean absolute error, is a simpler representation of error as a measure of the difference between the observation and prediction without regard to the sign of the error (Makridakis et al., 1998, p. 605). The absolute value of the value's difference results in the positive and negative errors not being cancelled out. However, MAE is not particularly good at comparing different sized data sets or unit values; for that MAPE is a better option.

Equation 20: Mean Absolute Error (MAE)

$$MAE = \frac{1}{n} \sum_{i=1}^{n} |y_i - \hat{y}_i| \tag{20}$$

Where

n = Number of observations in the model

Mean Absolute Percentage Error (MAPE)

Mean absolute percentage error is the average of all the percentage errors of a data set without regard to the sign of the errors (Makridakis et al., 1998, p. 605). MAPE includes an absolute value so that the positive and negative errors are not cancelled out. Moreover, the percentage of errors allows for better comparison across multiple forecasts of different magnitudes or even units as the percentage error can be compared equally. The drawback is that observations of zero result in an undefined MAPE. Both MAE and MAPE are common measures for forecasting error and are utilized throughout this analysis for their simplicity and commonality.

Equation 21: Mean Absolute Percentage Error (MAPE)

$$MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{21}$$

Where

n = Number of observations in the model

Summary

After evaluating the 618 AOC's ARST data provided from Apr 2012 to Mar 2018, the researcher saw patterns and seasonality in the data when separated by request status. The five statuses of Confirmed (the receiver and tanker have agreed to the details of the request which is supported), Bought (the tanker unit has agreed to support the request by clicking "buy" in ARST and, therefore, is supporting the request), Filled (the request is entered with an identified tanker unit in ARST assuming some initial coordination; however, the request may or may not be supported), Published (the request was not entered with an identified tanker unit and was not supported in the system, an unsupported air refueling request), and Cancelled (the request that was cancelled by either the tanker or receiver unit and may or may not have been able to be supported).

These five statuses were then aggregated by weeks from 8 Apr 2012 to 10 Mar 2018 in order to develop 309 weeks of data for the 5 statuses. The researcher then explored the major time series forecasting methods, including exponential smoothing and ARIMA, to understand how they applied to one another and to determine which might be the most valuable in forecasting future demand by each of the five statuses designated in the data set. The next chapter will apply this methodology along with measures of the error to determine the best potential models to compare to the withheld validation data.

IV. Analysis and Results

"I have seen the future and it is very much like the present, only longer."
- Kehlog Albran

Chapter Overview

This chapter presents the results of the researcher's time series forecasting for each of the five identified statuses represented in the ARST data set described in Chapter 1. The time series forecasting was accomplished in JMP to create models that would best represent the data while forecasting the next 52-weeks of air refueling requests. This analysis was grouped by seven-day weeks starting on the Sunday of each week. Upon initial review, the researcher determined that neither monthly groupings of data nor daily groupings of data were feasible for the analysis method due to small sample size for monthly groupings (72 total) and zero observations on certain days for particular daily groupings. Therefore, the weekly method was chosen for an adequate sample size of 309 observations with no weekly status groupings resulting in zero observations over the seven-day week. Once the data was grouped by week, multiple forecasting methods were applied to the data sets. Each forecasting method results in some error. Therefore, this chapter will begin with a discussion of measures of error and a discussion as to how different methods best fit a forecast for the source data. The best models were then used to forecast 52-weeks of data that was reserved from the initial forecast; this allowed for model validation and analysis of how error as a measure of best fit applied to the forecast data over the validation period.

Data Preparation

In order to utilize Box-Jenkins ARIMA Methodology for the time series forecasting models, the data had to be first analyzed for stationarity and seasonality. The researcher expected seasonal fluctuations over monthly and weekly time periods, which were visible in the data with peaks of requests surrounding April and October and a lack of requests overall in December and January.

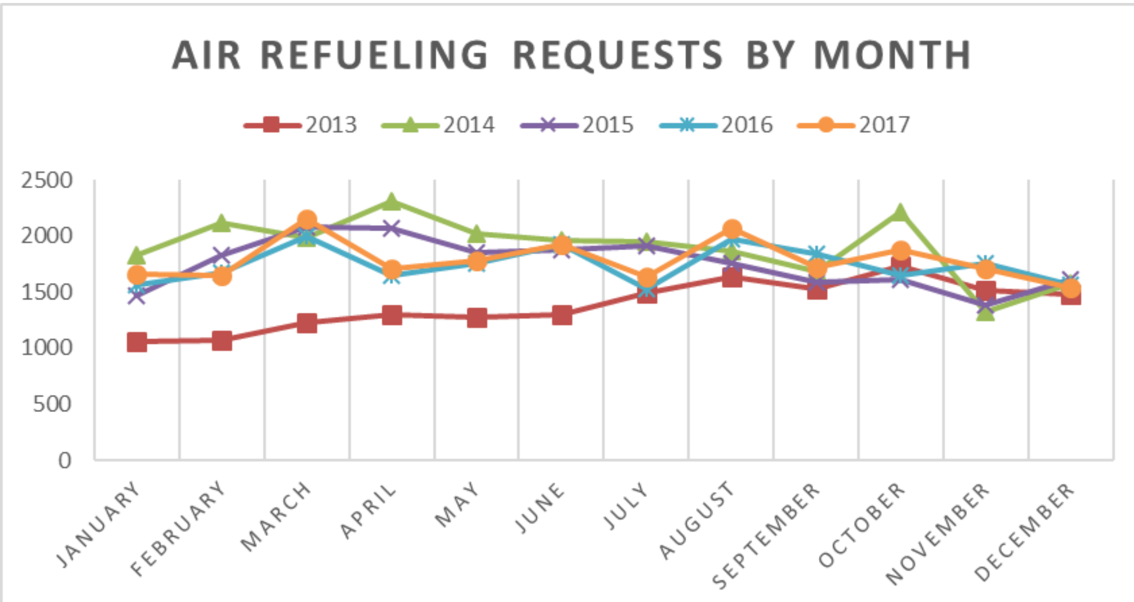

Figure 3: Monthly Distribution of Air Refueling Requests

The most apparent cyclical or seasonal fluctuations are seen within the weekly period. Graphing the data by the day of the week clearly shows that most requests are for Tuesday through Thursday, with the fewest for Saturday and Sunday, as shown in Figure 4. This figure represents the daily distribution of all 309 weeks of request data, each represented by a different color line in order to show highs and lows of the requests by the day of the week. Therefore, although there is some monthly seasonal fluctuation in the data, the real cyclical or seasonal activity happens at the weekly level. This observation was confirmed with forecasting models with less error over a 52-week period versus a 12-month period. This drives the 52-period calculations in the forecasting techniques with a seasonal component seen below.

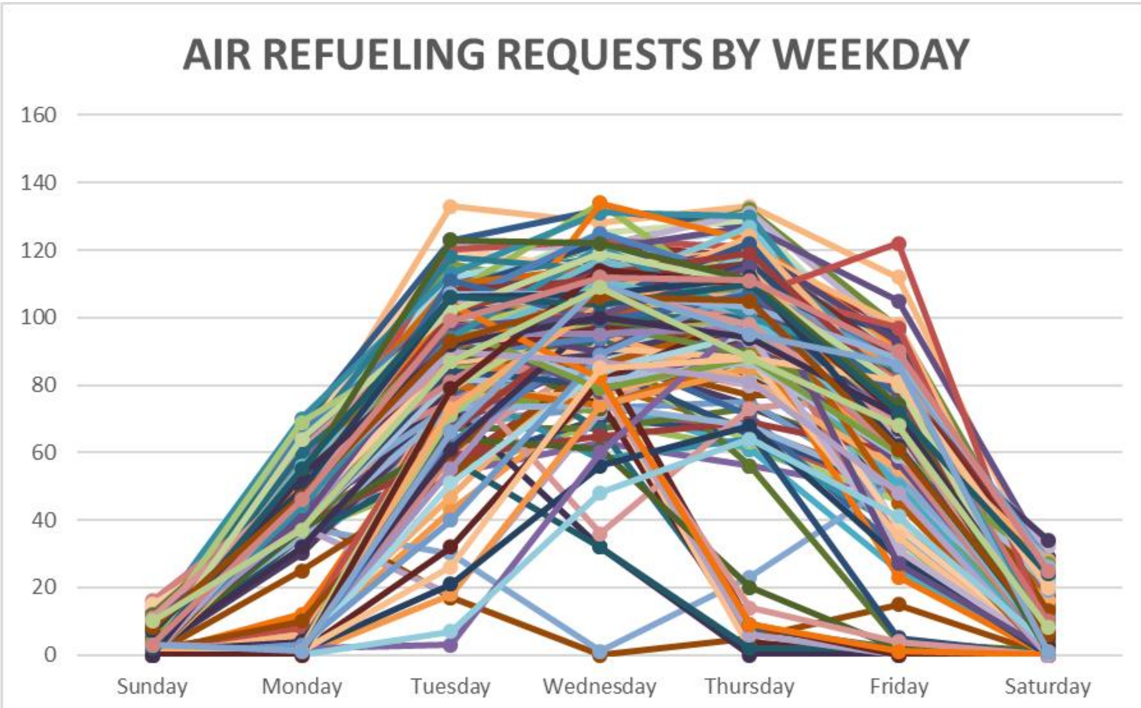

Figure 4: Daily Distribution of Air Refueling Requests

The next aspect to look for in the data was stationarity, which means that the data set's mean and variance are constant throughout the observed time periods in order to apply the ARIMA forecasting techniques. This is important because the larger the variance, the more difficult it will be to predict future occurrences. Therefore, the smaller the variance, the more accurate the forecast should be. One of the easiest ways to check for stationarity is to look at the autocorrelation function (ACF); if the ACF quickly decreases toward zero, the data set is stationary. Alternatively, if the values decrease slowly, as shown below, the data set is non-stationary (Ngo, 2013, p. 1). This example is from the Confirmed status and depicts the slow decrease in ACF values in the bottom left column, only decreasing within an acceptable margin of error at the 10^{th} lag point.

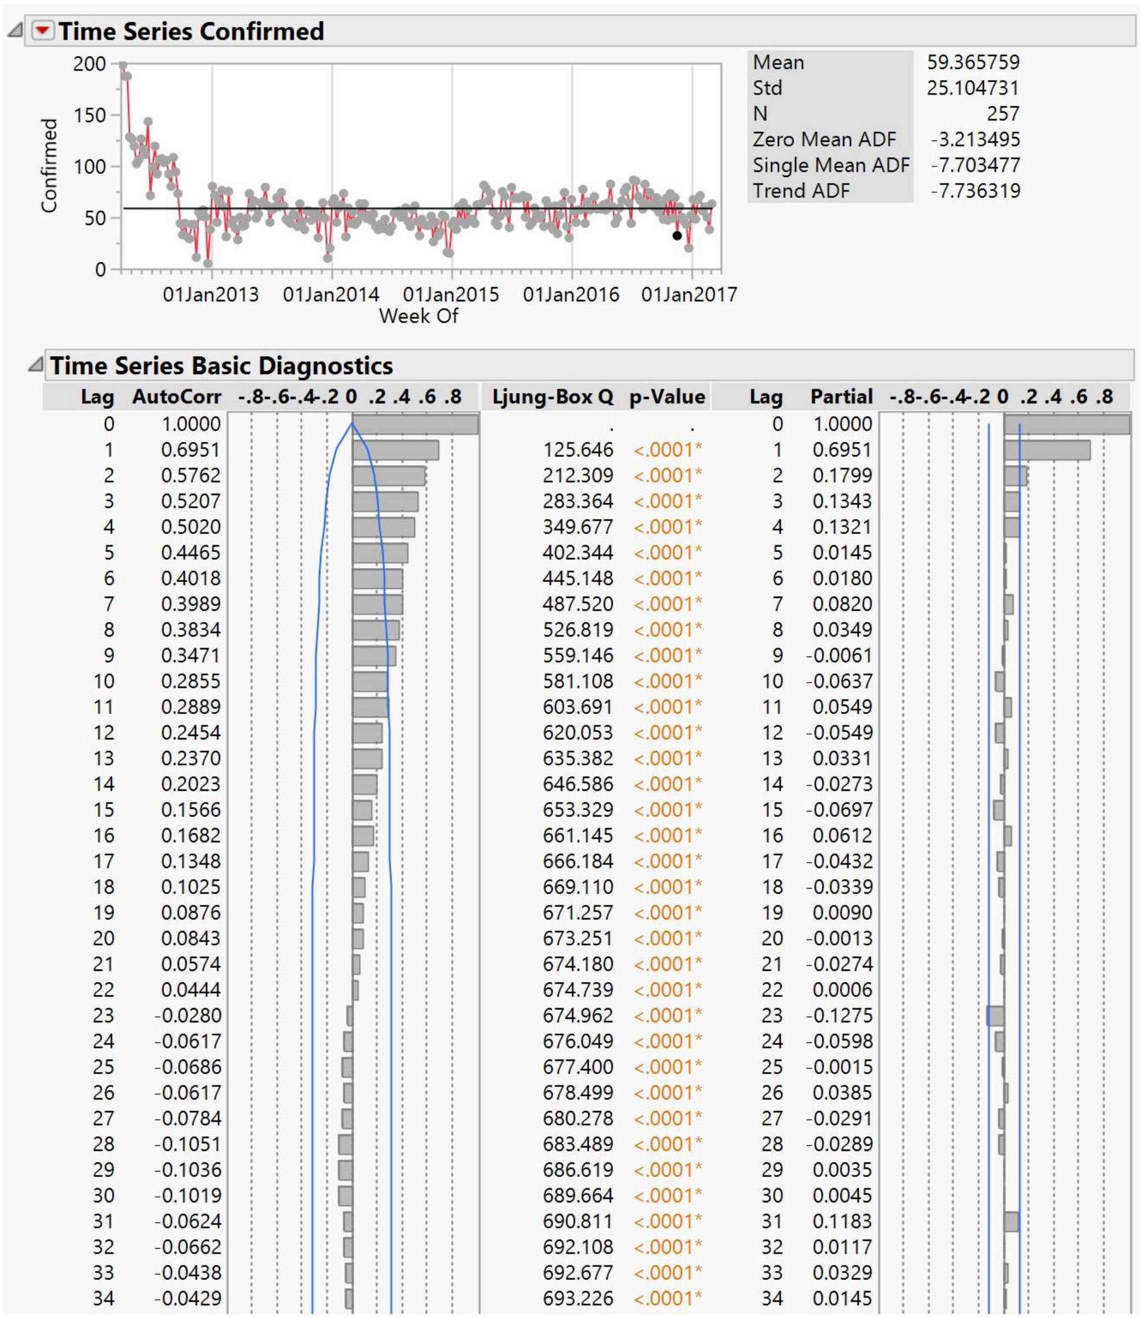

Lag	AutoCorr	Ljung-Box Q	p-Value	Lag	Partial
0	1.0000	.	.	0	1.0000
1	0.6951	125.646	<.0001*	1	0.6951
2	0.5762	212.309	<.0001*	2	0.1799
3	0.5207	283.364	<.0001*	3	0.1343
4	0.5020	349.677	<.0001*	4	0.1321
5	0.4465	402.344	<.0001*	5	0.0145
6	0.4018	445.148	<.0001*	6	0.0180
7	0.3989	487.520	<.0001*	7	0.0820
8	0.3834	526.819	<.0001*	8	0.0349
9	0.3471	559.146	<.0001*	9	-0.0061
10	0.2855	581.108	<.0001*	10	-0.0637
11	0.2889	603.691	<.0001*	11	0.0549
12	0.2454	620.053	<.0001*	12	-0.0549
13	0.2370	635.382	<.0001*	13	0.0331
14	0.2023	646.586	<.0001*	14	-0.0273
15	0.1566	653.329	<.0001*	15	-0.0697
16	0.1682	661.145	<.0001*	16	0.0612
17	0.1348	666.184	<.0001*	17	-0.0432
18	0.1025	669.110	<.0001*	18	-0.0339
19	0.0876	671.257	<.0001*	19	0.0090
20	0.0843	673.251	<.0001*	20	-0.0013
21	0.0574	674.180	<.0001*	21	-0.0274
22	0.0444	674.739	<.0001*	22	0.0006
23	-0.0280	674.962	<.0001*	23	-0.1275
24	-0.0617	676.049	<.0001*	24	-0.0598
25	-0.0686	677.400	<.0001*	25	-0.0015
26	-0.0617	678.499	<.0001*	26	0.0385
27	-0.0784	680.278	<.0001*	27	-0.0291
28	-0.1051	683.489	<.0001*	28	-0.0289
29	-0.1036	686.619	<.0001*	29	0.0035
30	-0.1019	689.664	<.0001*	30	0.0045
31	-0.0624	690.811	<.0001*	31	0.1183
32	-0.0662	692.108	<.0001*	32	0.0117
33	-0.0438	692.677	<.0001*	33	0.0329
34	-0.0429	693.226	<.0001*	34	0.0145

Figure 5: Non-Stationary Data Set

In order to correct for non-stationary data, the first or second difference can be taken to reduce the variance and make the data more constant around the mean, thereby transforming the data into a stationary time series (Ngo, 2013, p. 1). Below is the first difference of the same data set shown above. It now decreases quickly after the first lag and is centered around the mean of -0.53 with a smaller standard deviation.

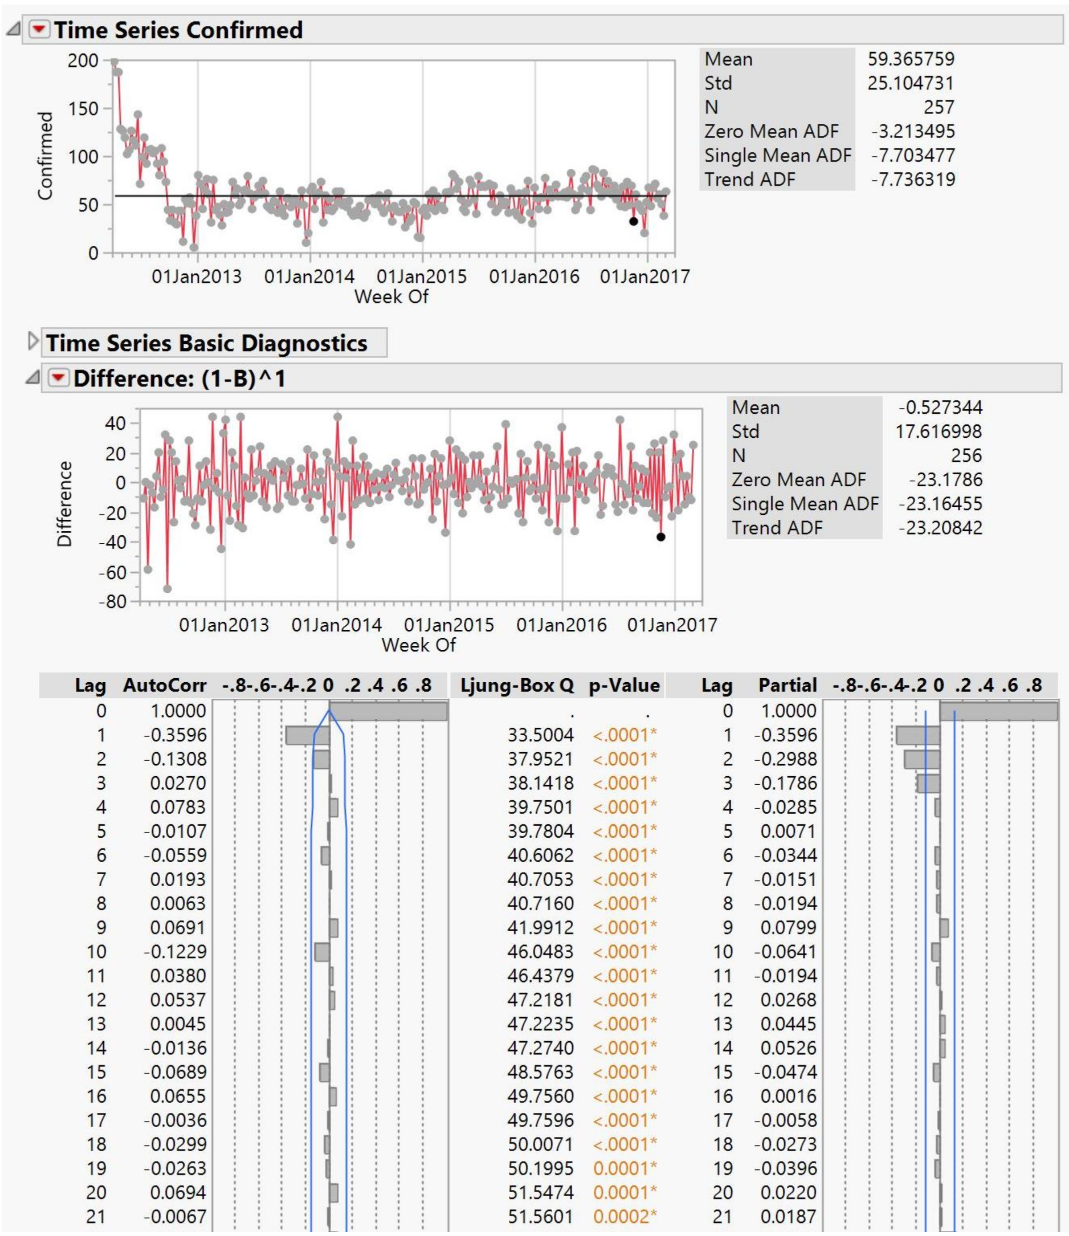

Lag	AutoCorr	-.8-.6-.4-.2 0 .2 .4 .6 .8	Ljung-Box Q	p-Value	Lag	Partial	-.8-.6-.4-.2 0 .2 .4 .6 .8
0	1.0000		.	.	0	1.0000	
1	-0.3596		33.5004	<.0001*	1	-0.3596	
2	-0.1308		37.9521	<.0001*	2	-0.2988	
3	0.0270		38.1418	<.0001*	3	-0.1786	
4	0.0783		39.7501	<.0001*	4	-0.0285	
5	-0.0107		39.7804	<.0001*	5	0.0071	
6	-0.0559		40.6062	<.0001*	6	-0.0344	
7	0.0193		40.7053	<.0001*	7	-0.0151	
8	0.0063		40.7160	<.0001*	8	-0.0194	
9	0.0691		41.9912	<.0001*	9	0.0799	
10	-0.1229		46.0483	<.0001*	10	-0.0641	
11	0.0380		46.4379	<.0001*	11	-0.0194	
12	0.0537		47.2181	<.0001*	12	0.0268	
13	0.0045		47.2235	<.0001*	13	0.0445	
14	-0.0136		47.2740	<.0001*	14	0.0526	
15	-0.0689		48.5763	<.0001*	15	-0.0474	
16	0.0655		49.7560	<.0001*	16	0.0016	
17	-0.0036		49.7596	<.0001*	17	-0.0058	
18	-0.0299		50.0071	<.0001*	18	-0.0273	
19	-0.0263		50.1995	0.0001*	19	-0.0396	
20	0.0694		51.5474	0.0001*	20	0.0220	
21	-0.0067		51.5601	0.0002*	21	0.0187	

Figure 6: Stationary Data Set

Differencing is a more complex process when applied to seasonal data. The five data sets being analyzed are non-stationary and seasonal data. Therefore, a data transformation must be applied prior to differencing and application of a forecasting technique in that order to reduce the data variance such as log, exp, square root, etc. (Ngo, 2013, p. 3). The researcher applied all 17 potential data transformation techniques in JMP to each of the five data sets and determined that natural log of the data reduced the variance the most for each of the five seasonal data sets. As an example, below is the first difference of the same Confirmed data set showing the data centered around a mean of effectively zero and a standard deviation of 0.4.

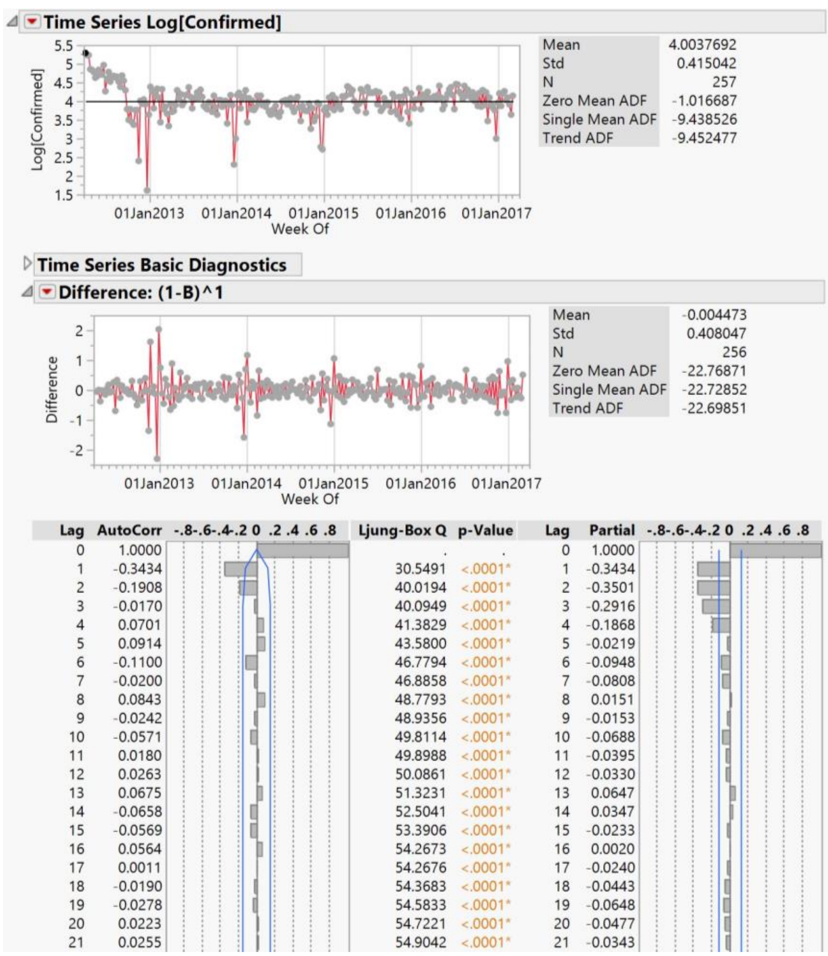

Figure 7: Seasonal Stationary Data Set

Model Selection, Estimation, and Diagnostics

After the above analysis, the researcher determined that the five air refueling request statuses were seasonal data on a weekly basis to be approximated by 52 equal time periods per year. The air refueling requests were disaggregated into the five statuses of Confirmed, Bought, Filled, Published, and Cancelled because each type of request has its own mean, variance, and specific input variables. A single forecast model could be fit for overall request into the ARST system; however, it would have additional error embedded into the model. Furthermore, the focus of this research was to highlight unsupported air refueling requests. The researcher is assuming that Confirmed and Bought requests are supported, that Filled and Cancelled include both supported and unsupported requests at unknown percentages until the system is updated with that fidelity, and that the Publish status includes all known unsupported requests. Therefore, disaggregation was required to get at the known unsupported requests.

Furthermore, all data sets required seasonal data transformation via taking the natural log of the original data for analysis and raising that number to the number e in order to transform the result back to the original form after the forecasting was complete. Also, all data sets with the exception of the Cancelled time series required differencing to develop stationary data sets. A full breakdown and validation of each data set are set out below.

The Confirmed Status

As discussed in Chapter 1, air refueling requests that have been supported by both the receiver unit and the tanker unit have the Confirmed status and were likely completed air refueling events. Prior to early April 2012, the only statuses shown in the ARST data were Confirmed and Published which is why the number of requests is very high in the early weeks of data; correspondingly, it is also why the other three statuses are low and take some time to catch up.

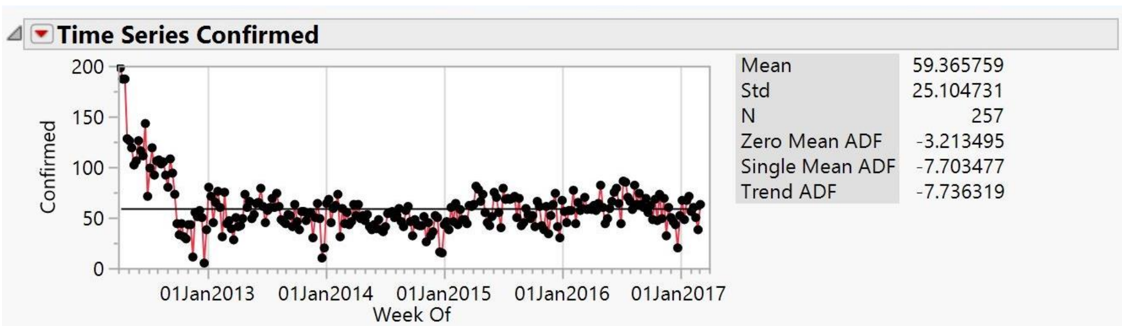

Figure 8: Time Series of Weekly Confirmed Status, 08 Apr 12 – 11 Mar 17

The above JMP graph of the Confirmed status over weekly intervals from Sunday through Saturday shows high request numbers in the first few months of 2012 moderating into a more cyclical fashion around the mean as time progresses. Note that this representation has not been transformed in any way and yet still shows a high standard deviation along with seasonality. In order to run multiple forecasting techniques, the data set was transformed with the natural log of the original observation, thus decreasing the variance of the data about the mean.

Model Selection

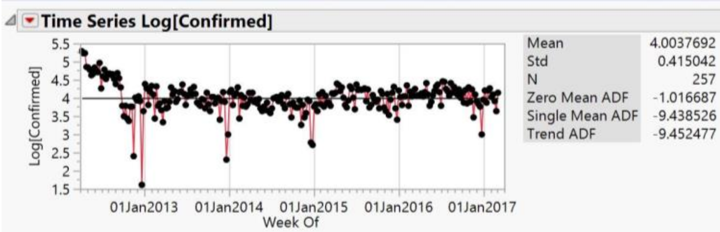

Time Series Basic Diagnostics

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal ARIMA(2, 1, 2)(1, 1, 0)52 No Intercept	199	0.0854796	89.275002	105.86560	0.008	79.275002	1	1	5.596895	0.215690
Winters Method (Additive)	201	0.0904561	104.38180	114.33616	-0.09	98.381801	2	3	5.766066	0.221448
Seasonal Exponential Smoothing(52, Zero to One)	202	0.0929329	107.12551	113.76175	-0.11	103.12551	3	2	5.833960	0.224379
Simple Exponential Smoothing(Zero to One)	255	0.1177171	180.61843	184.16361	0.295	178.61843	4	4	6.557148	0.234531
Damped-Trend Linear Exponential Smoothing	253	0.1186477	184.61843	195.25397	0.295	178.61843	5	5	6.557148	0.234531
Linear (Holt) Exponential Smoothing	253	0.1190704	190.43030	197.51283	0.254	186.4303	6	7	6.752387	0.245660
Double (Brown) Exponential Smoothing	254	0.1214103	193.08005	196.62132	0.244	191.08005	7	6	6.807101	0.247176
I(1) No Intercept	256	0.1665223	267.58430	267.58430	-0.00	267.5843	8	8	7.541781	0.273352
ARIMA(0, 0, 0)	256	0.1729327	279.33534	282.88441	0.000	277.33534	9	9	7.682062	0.283234

Figure 9: Confirmed Status Model Comparison

The model comparison chart above allows for the assessment of different forecasting techniques in JMP. Each of the forecasting techniques described in Chapter 3 was applied to the Confirmed data. Moreover, a minimum of 1250 different iterations of seasonal ARIMA permutations was applied in accordance with the Box-Jenkins Methodology (available in Appendix C) in order to arrive at the lowest AIC/SBC values while verifying significant P values for each parameter. In this case, the seasonal ARIMA(2,1,2)(1,1,0)52 with no intercept resulted in the lowest AIC/SBC values along with the lowest MAPE and MAE error values for the data used to develop the forecast from 08 Apr 12 – 11 Mar 17. This indicates that the seasonal ARIMA model is the best fit for the forecast data with the least number of parameters and the least overall error.

Time Series Log[Confirmed]

Model: Seasonal ARIMA(2, 1, 2)(1, 1, 0)52 No Intercept

Model Summary

DF	199	Stable	Yes
Sum of Squared Errors	17.0104367	Invertible	Yes
Variance Estimate	0.08547958		
Standard Deviation	0.29236891		
Akaike's 'A' Information Criterion	89.2750021		
Schwarz's Bayesian Criterion	105.865602		
RSquare	0.00842266		
RSquare Adj	-0.0115085		
MAPE	5.59689474		
MAE	0.21569047		
-2LogLikelihood	79.2750021		

Parameter Estimates

Term	Factor	Lag	Estimate	Std Error	t Ratio	Prob>|t|
AR1,1	1	1	-1.001788	0.1063419	-9.42	<.0001*
AR1,2	1	2	-0.224531	0.1039438	-2.16	0.0320*
AR2,52	2	52	-0.335018	0.0836814	-4.00	<.0001*
MA1,1	1	1	-0.291071	0.0973296	-2.99	0.0031*
MA1,2	1	2	0.543818	0.0958664	5.67	<.0001*

Forecast

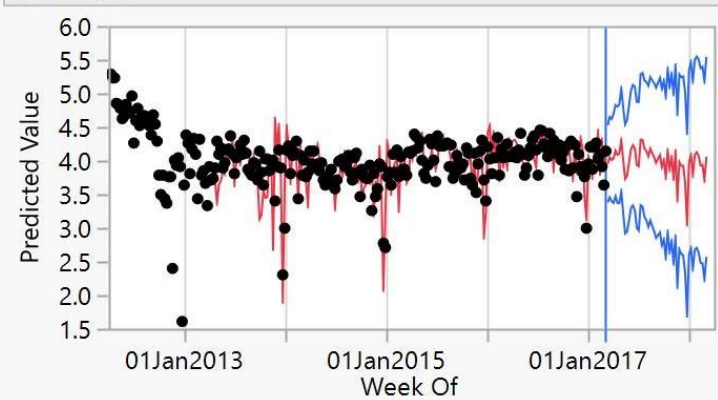

Residuals

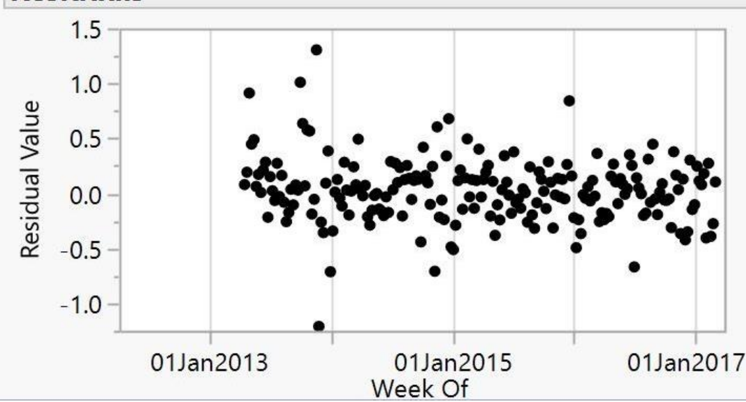

Figure 10: Confirmed Status Seasonal ARIMA(2,1,2)(1,1,0)52 No Intercept

Looking closer at the selected seasonal ARIMA model, the basic information about the model is available, and each of the parameters is significant. The AR1,2 parameter is the least of all at just above 3%. However, it is not possible to exclude one AR or MA parameter and keep the other parameters with this type of ARIMA forecasting. Nevertheless, all meet a 95% or greater significance. Additionally, the depicted residuals appear to be evenly distributed about zero and show no clear patterns or skew.

Model Validation

Based on the JMP output and forecast measures such as AIC/SBC, the seasonal ARIMA model appears to be the best fit for the data used during the forecast period. However, the only true test of a forecast model is to compare the forecast, for 52-weeks in this case, to actual data not used in the development of the model. The researcher utilized the MAPE and MAE error statistics in order to show which model represents the least error and, therefore, is the most accurate forecast for the subsequent 52-weeks, from 12 March 17 to 10 March 18.

Table 4: Confirmed Status – 52-Week Validation Set Error Statistics

Model	MAPE	MAE
Mean Model / White Noise (0,0,0) + c	4.687855	0.187497
Random Walk Model (0,1,0)	4.605422	0.178326
Simple Exponential Smoothing (0 ≤ α ≤ 1)	4.946746	0.199302
Double (Brown) Linear Exponential Smoothing	5.722216	0.234371
Linear (Holt) Exponential Smoothing	5.359548	0.218370
Damped Trend Linear Exponential Smoothing	4.946746	0.199302
Seasonal Exponential Smoothing (52 weeks, (0 ≤ δ ≤ 1))	6.622882	0.268338
Winters Method (Additive)	5.120940	0.205673
Seasonal ARIMA (2,1,2)(1,1,0)52 No Intercept	4.310800	0.173742

As shown above, the seasonal ARIMA model both provided the best model of the forecast data and resulted in the least error for the 52-week validation data set, leading the researcher to conclude that it was the best model analyzed to forecast future Confirmed air refueling requests.

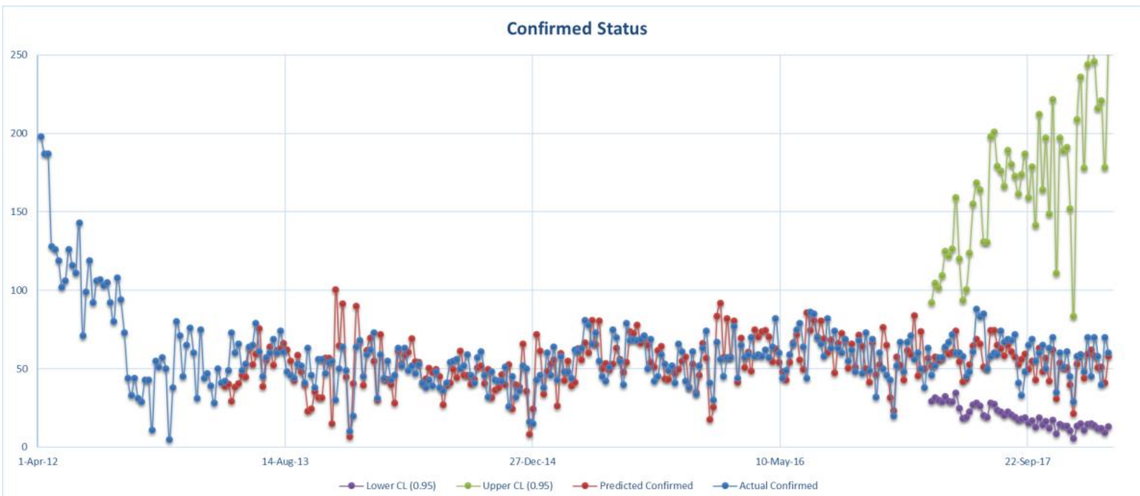

Figure 11: Confirmed Status Forecast with 95% Confidence Interval

The above graph depicts the entire data set over the 309 weeks, including a 95% confidence interval for the forecast versus the actual request values in the last 52-weeks. In reality, there were 3058 Confirmed requests with a forecast for 2863. Therefore, the forecast was a conservative estimate by 195 requests or 6% less than actual.

The Bought Status

Air refueling requests with the Bought status have been supported by the tanker unit, but not yet Confirmed by the receiver unit. For the purposes of understanding the tanker capacity and ability to support receiver demand, these requests were considered as likely completed air refueling events.

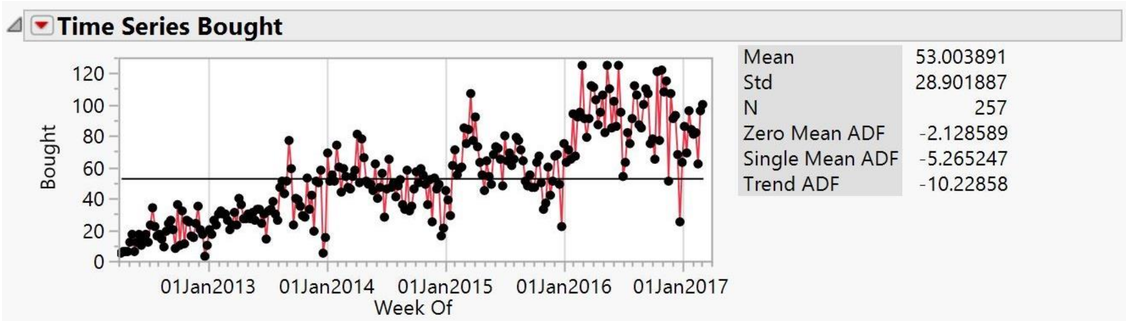

Figure 12: Time Series of Weekly Bought Status, 08 Apr 12 – 11 Mar 17

The above JMP graph of the Bought status over weekly intervals from Sunday through Saturday shows low request numbers through late 2013 as this new status was only included in the ARST since the spring of 2012. Also apparent is some cyclical seasonality. This graph of the Bought data has not been transformed in any way and yet still shows a high standard deviation, along with seasonality, to illustrate the difference from the figure below. In order to run multiple forecasting techniques, this data set was transformed with the natural log of the original observation, thereby decreasing the variance of the data about the mean.

Model Selection

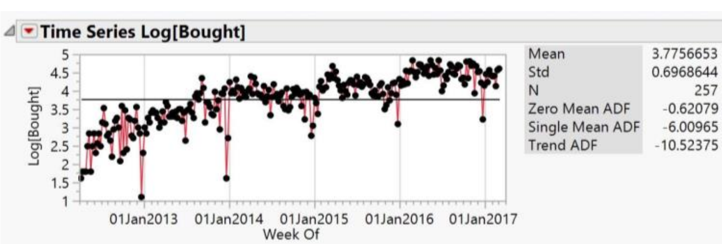

Time Series Basic Diagnostics

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal ARIMA(0, 1, 1)(0, 1, 1)52 No Intercept	202	0.1272945	165.77693	172.41317	0.424	161.77693	1	1	7.300150	0.277830
Seasonal Exponential Smoothing(52, Zero to One)	202	0.1241088	167.03699	173.67323	0.419	163.03699	2	2	7.365317	0.280770
Winters Method (Additive)	201	0.123676	167.52052	177.47488	0.422	161.52052	3	3	7.335935	0.279841
Simple Exponential Smoothing(Zero to One)	255	0.1488443	240.62547	244.17064	0.683	238.62547	4	4	8.667453	0.277161
Damped-Trend Linear Exponential Smoothing	253	0.1500209	244.62547	255.26100	0.683	238.62547	5	5	8.667453	0.277161
Linear (Holt) Exponential Smoothing	253	0.1490995	248.47556	255.55808	0.663	244.47556	6	6	8.891393	0.275433
Double (Brown) Exponential Smoothing	254	0.1542407	254.20592	257.74719	0.654	252.20592	7	7	9.105281	0.285873
I(1) No Intercept	256	0.208965	325.70593	325.70593	0.555	325.70593	8	8	9.973805	0.317796
ARIMA(0, 0, 0)	256	0.487517	545.69591	549.24498	-0.00	543.69591	9	9	17.551592	0.548101

Figure 13: Bought Status Model Comparison

The model comparison chart above allows for the assessment of different forecasting techniques in JMP. Each of the forecasting techniques described in Chapter 3 was also applied to the Bought data. Moreover, a minimum of 1250 different iterations of seasonal ARIMA permutations was applied in accordance with the Box-Jenkins Methodology (available in Appendix D) in order to arrive at the lowest AIC/SBC values while verifying significant P values for each parameter. In this case, the seasonal ARIMA(0,1,1)(0,1,1)52 with no intercept resulted in the lowest AIC/SBC values along with the lowest MAPE, although three models resulted in a slightly lower MAE for the data used to develop the forecast from 08 Apr 12 – 11 Mar 17. The lower MAEs are all remarkably close to the seasonal ARIMA's value; this represents something to consider during model validation. Moreover, the model comparison indicates that the seasonal ARIMA model is the best fit for the forecast data with the least number of parameters and the least overall error.

Time Series Log[Bought]

Model: Seasonal ARIMA(0, 1, 1)(0, 1, 1)52 No Intercept

Model Summary

DF	202	Stable	Yes
Sum of Squared Errors	25.7134916	Invertible	Yes
Variance Estimate	0.12729451		
Standard Deviation	0.35678357		
Akaike's 'A' Information Criterion	165.776928		
Schwarz's Bayesian Criterion	172.413168		
RSquare	0.42422153		
RSquare Adj	0.42137114		
MAPE	7.30015018		
MAE	0.27783025		
-2LogLikelihood	161.776928		

Parameter Estimates

Term	Factor	Lag	Estimate	Std Error	t Ratio	Prob>|t|
MA1,1	1	1	0.69825594	0.0624390	11.18	<.0001*
MA2,52	2	52	0.29358845	0.0832284	3.53	0.0005*

Forecast

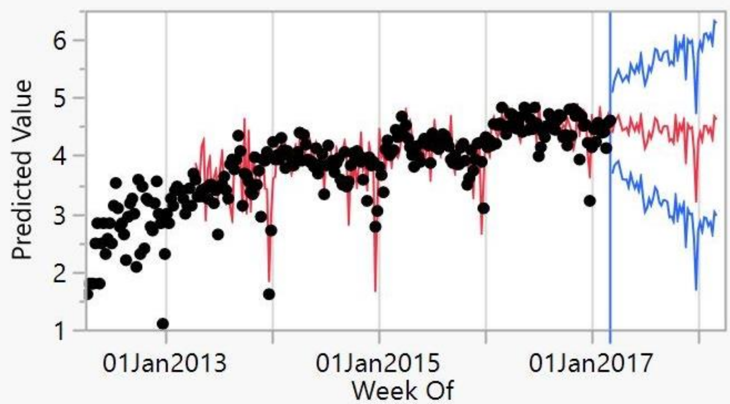

Residuals

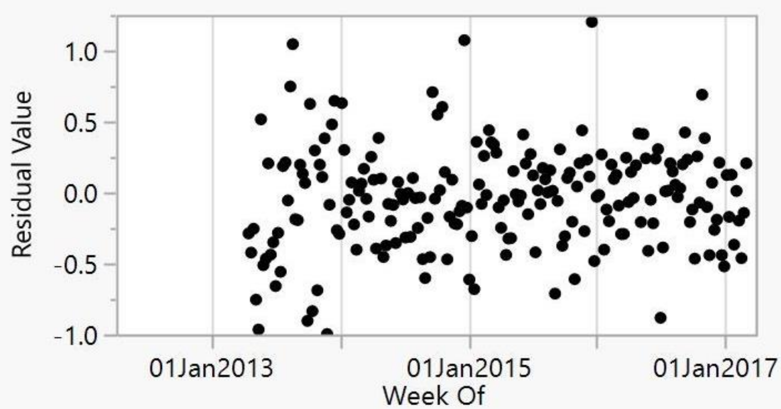

Figure 14: Bought Status Seasonal ARIMA(0,1,1)(0,1,1)52 No Intercept

Looking closer at the selected seasonal ARIMA model, the basic information about the model is available. Each of the parameters is significant and well above a 95% significance level. Additionally, the depicted residuals appear to be evenly distributed about zero and show no clear patterns or skew.

Model Validation

Based on the JMP output and forecast measures such as AIC/SBC, the seasonal ARIMA model appears to be the best fit for the data used during the forecast period. Again, the 52-weeks of validation data will be used to calculate the lowest MAPE and MAE error statistics. The resulting model represents the least error and, therefore, is the most accurate forecast for the subsequent 52-weeks (12 March 17 to 10 March 18).

Table 5: Bought Status – 52-Week Validation Set Error Statistics

Model	MAPE	MAE
Mean Model / White Noise (0,0,0) + c	12.347367	0.535522
Random Walk Model (0,1,0)	7.735098	0.316859
Simple Exponential Smoothing (0 ≤ α ≤ 1)	4.487647	0.180085
Double (Brown) Linear Exponential Smoothing	4.002118	0.160925
Linear (Holt) Exponential Smoothing	9.651085	0.398166
Damped Trend Linear Exponential Smoothing	4.487640	0.180085
Seasonal Exponential Smoothing (52 weeks, (0 ≤ δ ≤ 1)	6.040294	0.256403
Winters Method (Additive)	4.033481	0.170903
Seasonal ARIMA (0,1,1)(0,1,1)52 No Intercept	3.918795	0.164768

As shown above, the seasonal ARIMA model both provided the best model of the forecast data and resulted in the least MAPE and second lowest MAE error for the 52-week validation data set, leading the researcher to conclude that it was the best model analyzed to forecast future Bought air refueling requests.

Figure 15: Bought Status Forecast with 95% Confidence Interval

The above graph depicts the entire data set over the 309 weeks, including a 95% confidence interval for the forecast versus the actual request values in the last 52-weeks. In reality, there were 3882 Bought requests with a forecast for 4392. Therefore, the forecast overestimated by 510 requests or 13%.

The Filled Status

Air refueling requests submitted with tanker details and the presumption of pre-coordination between the receiver and tanker unit have the Filled status. The challenge with this category is that receivers can enter tanker unit information from which they intend to receive air refueling service without any knowledge about whether the tanker unit can actually support the AR event. Therefore, some of the requests in the Filled status were probably supported by the tanker units and others were not. It is just not possible with the current database information to determine how much demand was unsupported.

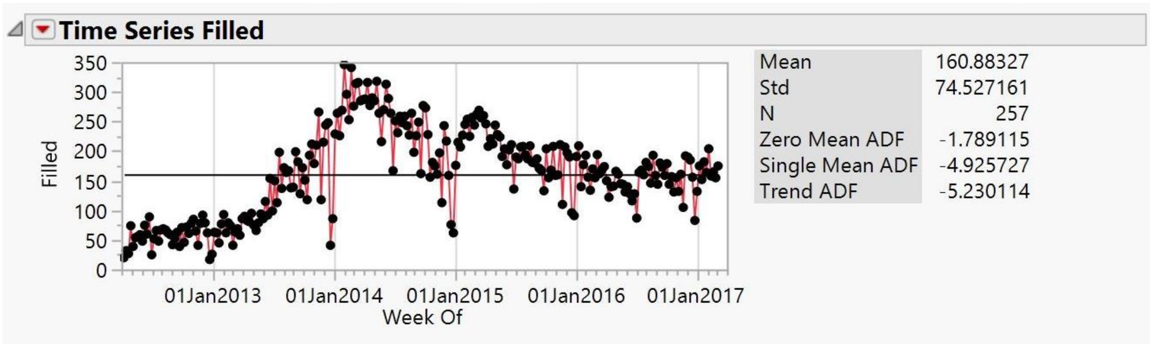

Figure 16: Time Series of Weekly Filled Status, 08 Apr 12 – 11 Mar 17

The above JMP graph of the Filled status over weekly intervals from Sunday through Saturday shows low request numbers through late 2013 as this new status was only included in the ARST since the spring of 2012. Some cyclical seasonality is apparent, along with very high numbers from 2014 to 2015 as users gained an understanding of the status. This graph of the Filled data has not been transformed in any way and yet still shows a high standard deviation along with seasonality to illustrate the difference from the figure below. In order to run multiple forecasting techniques, this data set was transformed with the natural log of the original observation decreasing the variance of the data about the mean.

Model Selection

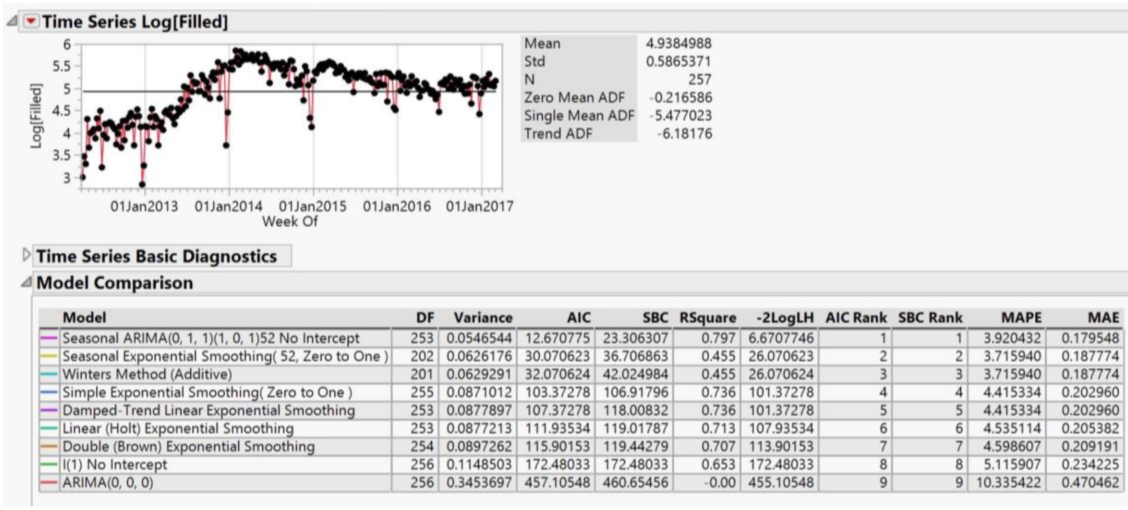

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal ARIMA(0, 1, 1)(1, 0, 1)52 No Intercept	253	0.0546544	12.670775	23.306307	0.797	6.6707746	1	1	3.920432	0.179548
Seasonal Exponential Smoothing(52, Zero to One)	202	0.0626176	30.070623	36.706863	0.455	26.070623	2	2	3.715940	0.187774
Winters Method (Additive)	201	0.0629291	32.070624	42.024984	0.455	26.070624	3	3	3.715940	0.187774
Simple Exponential Smoothing(Zero to One)	255	0.0871012	103.37278	106.91796	0.736	101.37278	4	4	4.415334	0.202960
Damped-Trend Linear Exponential Smoothing	253	0.0877897	107.37278	118.00832	0.736	101.37278	5	5	4.415334	0.202960
Linear (Holt) Exponential Smoothing	253	0.0877213	111.93534	119.01787	0.713	107.93534	6	6	4.535114	0.205382
Double (Brown) Exponential Smoothing	254	0.0897262	115.90153	119.44279	0.707	113.90153	7	7	4.598607	0.209191
I(1) No Intercept	256	0.1148503	172.48033	172.48033	0.653	172.48033	8	8	5.115907	0.234225
ARIMA(0, 0, 0)	256	0.3453697	457.10548	460.65456	-0.00	455.10548	9	9	10.335422	0.470462

Figure 17: Filled Status Model Comparison

The model comparison chart above allows for the assessment of different forecasting techniques in JMP. Each of the forecasting techniques described in Chapter 3 was also applied to the Filled data. Moreover, a minimum of 1250 different iterations of seasonal ARIMA permutations was applied in accordance with the Box-Jenkins Methodology (available in Appendix E) in order to arrive at the lowest AIC/SBC values while verifying significant P values for each parameter. In this case, the seasonal ARIMA(0,1,1)(1,0,1)52 with no intercept resulted in the lowest AIC/SBC values along with the lowest MAE, although two other models resulted in a slightly lower MAPE for the data used to develop the forecast from 08 Apr 12 to 11 Mar 17. The lower MAPE values remained relatively close to the seasonal ARIMA's MAPE error value. This indicates that the seasonal ARIMA model is the best fit for the forecast data with the least number of parameters and the least overall error.

Time Series Log[Filled]

Model: Seasonal ARIMA(0, 1, 1)(1, 0, 1)52 No Intercept

Model Summary

DF	253	Stable	Yes
Sum of Squared Errors	13.827561	Invertible	Yes
Variance Estimate	0.05465439		
Standard Deviation	0.23378279		
Akaike's 'A' Information Criterion	12.6707746		
Schwarz's Bayesian Criterion	23.3063069		
RSquare	0.79698626		
RSquare Adj	0.7953814		
MAPE	3.92043186		
MAE	0.17954776		
-2LogLikelihood	6.6707746		

Parameter Estimates

Term	Factor	Lag	Estimate	Std Error	t Ratio	Prob>|t|
AR2,52	2	52	0.77030373	0.0663125	11.62	<.0001*
MA1,1	1	1	0.66653419	0.0494409	13.48	<.0001*
MA2,52	2	52	0.26649630	0.1069835	2.49	0.0134*

Forecast

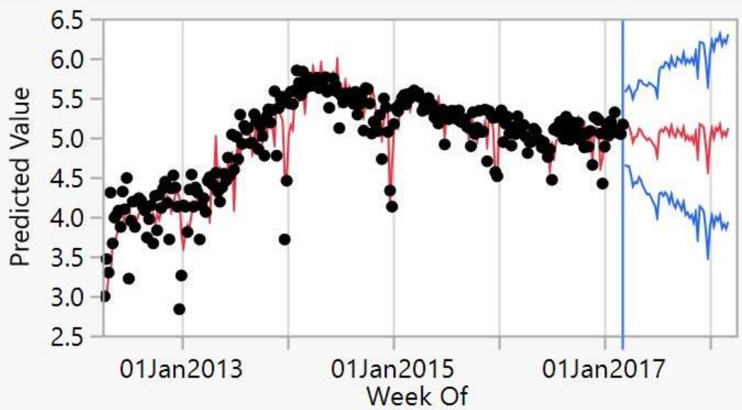

Residuals

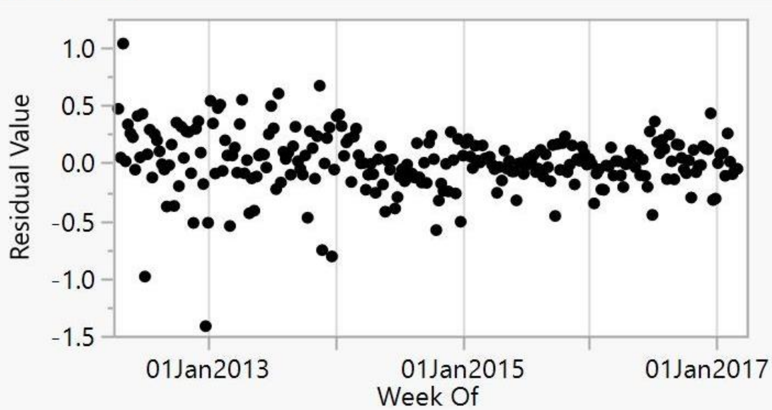

Figure 18: Filled Status Seasonal ARIMA(0,1,1)(1,0,1)52 No Intercept

Again, looking closer at the selected Filled seasonal ARIMA model, the basic information about the model is available. Each of the parameters is significant and well above a 95% significance level; the lowest is the MA2,52 parameter, which is only 1.3%, still significant and required to maintain the MA1,1 parameter. Additionally, the depicted residuals appear to be evenly distributed about zero and show no clear patterns or skew.

Model Validation

Based on the JMP output and forecast measures such as AIC/SBC, the Filled seasonal ARIMA model appears to be the best fit for the data used during the forecast period. Again, the 52-weeks of validation data will be used to calculate the lowest MAPE and MAE error statistics. The resulting model represents the least error and, therefore, is the most accurate forecast for the subsequent 52-weeks (12 March 17 to 10 March 18).

Table 6: Filled Status – 52-Week Validation Set Error Statistics

Model	MAPE	MAE
Mean Model / White Noise (0,0,0) + c	3.936391	0.195276
Random Walk Model (0,1,0)	3.355093	0.159267
Simple Exponential Smoothing (0 ≤ α ≤ 1)	3.102410	0.147881
Double (Brown) Linear Exponential Smoothing	4.418292	0.211397
Linear (Holt) Exponential Smoothing	3.319933	0.157722
Damped Trend Linear Exponential Smoothing	3.102410	0.147881
Seasonal Exponential Smoothing (52 weeks, (0 ≤ δ ≤ 1)	3.734782	0.185105
Winters Method (Additive)	3.734782	0.185105
Seasonal ARIMA (0,1,1)(1,0,1)52 No Intercept	2.222276	0.108962

As shown above, the seasonal ARIMA model both provided the best model of the forecast data and resulted in the least error for the 52-week validation data set, leading the researcher to conclude that it was the best model analyzed to forecast future Filled air refueling requests.

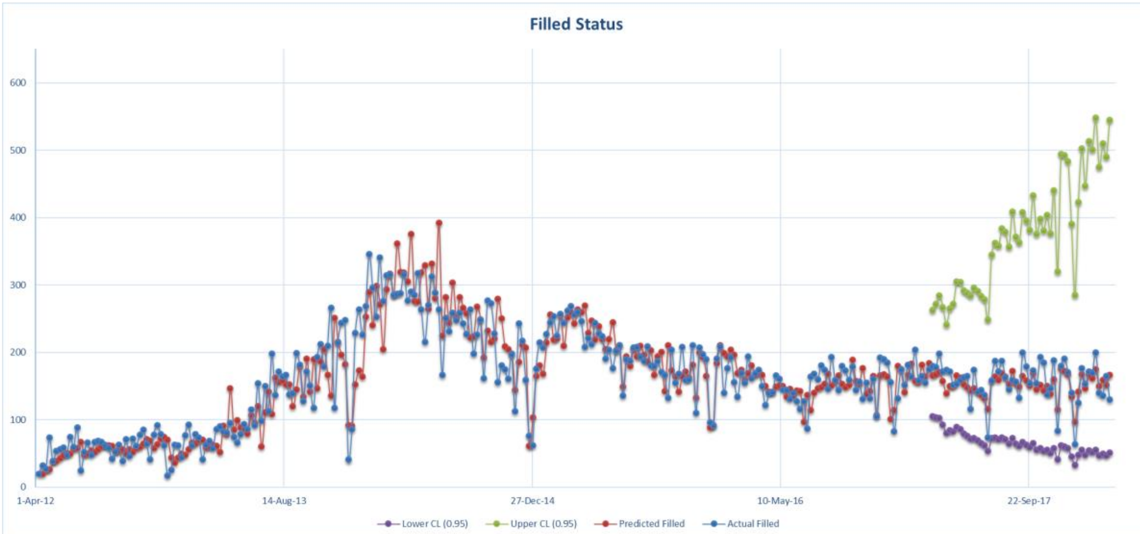

Figure 19: Filled Status Forecast with 95% Confidence Interval

The above graph depicts the entire data set over the 309 weeks, including a 95% confidence interval for the forecast versus the actual request values in the last 52-weeks. Of note, there were 8223 Filled requests with a forecast for 7970. Therefore, the forecast was a conservative estimate by 253 requests or 3% less than actual.

The Published Status

Air refueling requests that remain in the Published status are those requests with no tanker details and that are understood to have been unsupported in the system. This category is in short term focus for this research because, unlike the Filled status and the Cancelled status yet to come, where it was unclear how many requests went unsupported, all of those requests that remained in the Published status were unsupported by tanker units.

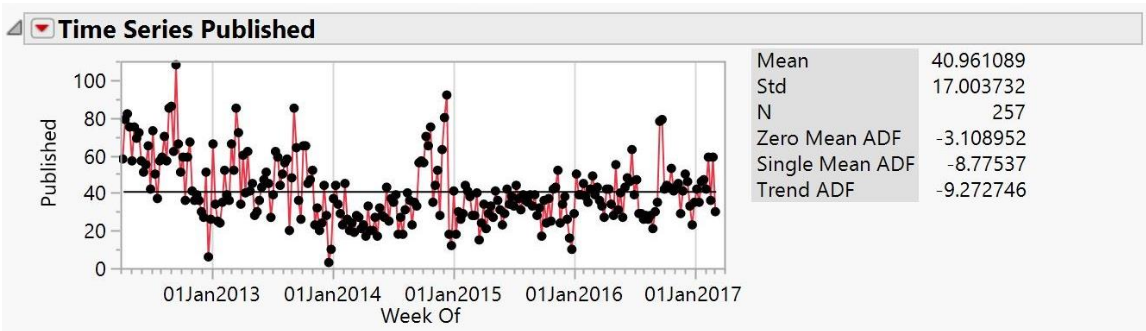

Figure 20: Time Series of Weekly Published Status, 08 Apr 12 – 11 Mar 17

The above JMP graph of the Published status over weekly intervals from Sunday through Saturday shows some higher request numbers in 2012 as the system moved away from Confirmed and Published being the only possible categories. The data is fairly consistent about the mean, with some cyclical seasonality. The above graph of the Published data has not been transformed in any way and yet still shows a high standard deviation along with seasonality to illustrate the difference from the figure below. In order to run multiple forecasting techniques, this data set was transformed with the natural log of the original observation decreasing the variance of the data about the mean.

Model Selection

Time Series Log[Published]

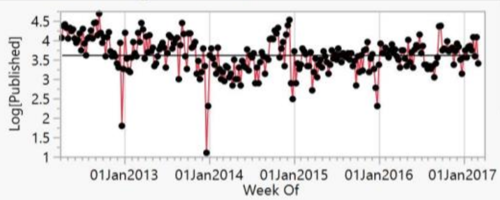

Mean	3.6215343
Std	0.4511277
N	257
Zero Mean ADF	-1.104706
Single Mean ADF	-9.524027
Trend ADF	-9.764479

Time Series Basic Diagnostics

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal ARIMA(1, 1, 2)(1, 1, 0)52	199	0.1435461	208.09099	224.68159	0.093	198.09099	1	1	9.246431	0.316599
Seasonal Exponential Smoothing(52, Zero to One)	202	0.1405198	220.52234	227.15858	0.002	216.52234	2	2	9.351976	0.319951
Winters Method (Additive)	201	0.1412189	222.52234	232.47670	0.002	216.52234	3	3	9.351976	0.319951
Simple Exponential Smoothing(Zero to One)	255	0.1596962	258.66317	262.20834	0.218	256.66317	4	4	8.994626	0.288300
Damped-Trend Linear Exponential Smoothing	253	0.1609586	262.66317	273.29870	0.218	256.66317	5	5	8.994626	0.288300
Linear (Holt) Exponential Smoothing	253	0.1609765	268.75299	275.83552	0.193	264.75299	6	6	9.130169	0.295633
Double (Brown) Exponential Smoothing	254	0.1694224	278.14542	281.68669	0.159	276.14542	7	7	9.263584	0.298679
ARIMA(0, 0, 0)	256	0.2043111	322.18787	325.73695	-0.00	320.18787	8	8	10.329647	0.335868
I(1) No Intercept	256	0.2137264	331.47354	331.47354	-0.05	331.47354	9	9	10.450418	0.338557

Figure 21: Published Status Model Comparison

The model comparison chart above allows for the assessment of different forecasting techniques in JMP. Each of the forecasting techniques described in Chapter 3 was also applied to the Filled data. Moreover, a minimum of 1250 different iterations of seasonal ARIMA permutations was applied in accordance with the Box-Jenkins Methodology (available in Appendix F) in order to arrive at the lowest AIC/SBC values while verifying significant P values for each parameter. In this case, the seasonal ARIMA(1,1,2)(1,1,0)52 resulted in the lowest AIC/SBC values but not the lowest MAPE and MAE error statistics. In fact, the seasonal ARIMA model was 4^{th} and 5^{th} lowest for the data used to develop the forecast from 08 Apr 12 to 11 Mar 17. This indicates that the seasonal ARIMA model is the best fit for the forecast data with the least number of parameters but with slightly higher error values than represented by other models. The true test of the better model will be a comparison of the error values that result from the validation set presented below.

Time Series Log[Published]

Model: Seasonal ARIMA(1, 1, 2)(1, 1, 0)52

Model Summary

DF	199	Stable	Yes
Sum of Squared Errors	28.5656664	Invertible	Yes
Variance Estimate	0.14354606		
Standard Deviation	0.37887473		
Akaike's 'A' Information Criterion	208.090987		
Schwarz's Bayesian Criterion	224.681587		
RSquare	0.09271058		
RSquare Adj	0.07447361		
MAPE	9.24643061		
MAE	0.31659882		
-2LogLikelihood	198.090987		

Parameter Estimates

Term	Factor	Lag	Estimate	Std Error	t Ratio	Prob>|t|	Constant Estimate
AR1,1	1	1	0.771906	0.0996602	7.75	<.0001*	0.00109487
AR2,52	2	52	-0.511351	0.0616322	-8.30	<.0001*	
MA1,1	1	1	1.426364	0.1444692	9.87	<.0001*	
MA1,2	1	2	-0.426381	0.1435263	-2.97	0.0033*	
Intercept	1	0	0.003176	0.0009278	3.42	0.0008*	

Forecast

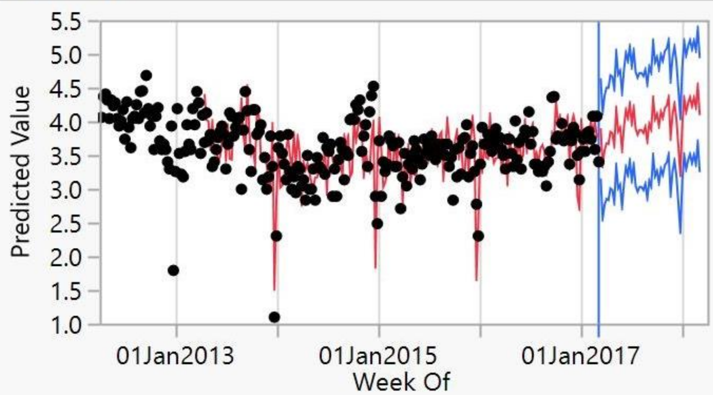

Residuals

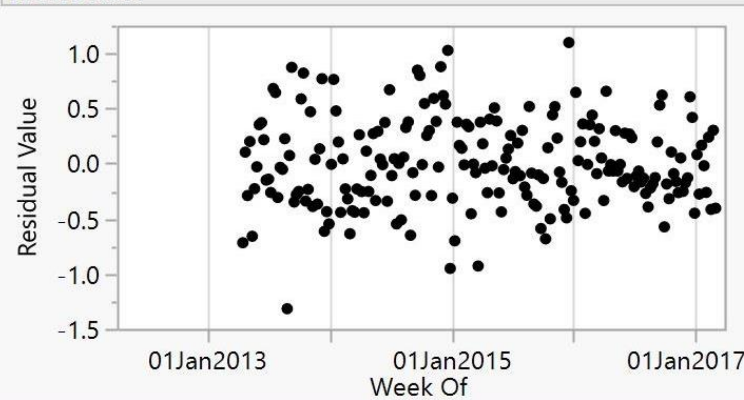

Figure 22: Published Status Seasonal ARIMA(1,1,2)(1,1,0)52

Again, looking closer at the selected Filled seasonal ARIMA model, the basic information about the model is available. Each of the parameters is significant and well above a 95% significance level. Additionally, the depicted residuals appear to be evenly distributed about zero and show no clear patterns or skew.

Model Validation

Based on the JMP output and forecast measures such as AIC/SBC, the Published seasonal ARIMA model again appears to be the best fit for the data used during the forecast period. Again, the 52-weeks of validation data will be used to calculate the lowest MAPE and MAE error statistics. The resulting model represents the least error and, therefore, is the most accurate forecast for the subsequent 52-weeks from 12 March 17 to 10 March 18.

Table 7: Published Status – 52-Week Validation Set Error Statistics

Model	MAPE	MAE
Mean Model / White Noise (0,0,0) + c	11.397139	0.473538
Random Walk Model (0,1,0)	16.486559	0.682855
Simple Exponential Smoothing (0 ≤ α ≤ 1)	9.216920	0.383172
Double (Brown) Linear Exponential Smoothing	7.459565	0.308530
Linear (Holt) Exponential Smoothing	10.420547	0.433338
Damped Trend Linear Exponential Smoothing	9.216920	0.383172
Seasonal Exponential Smoothing (52 weeks, (0 ≤ δ ≤ 1)	9.062291	0.373491
Winters Method (Additive)	9.062291	0.373491
Seasonal ARIMA (1,1,2)(1,1,0)52	6.762654	0.276404

As shown above, the seasonal ARIMA model both provided the best model of the forecast data and resulted in the least error for the 52-week validation data set, leading the researcher to conclude that it was the best model analyzed to forecast future Published air refueling requests.

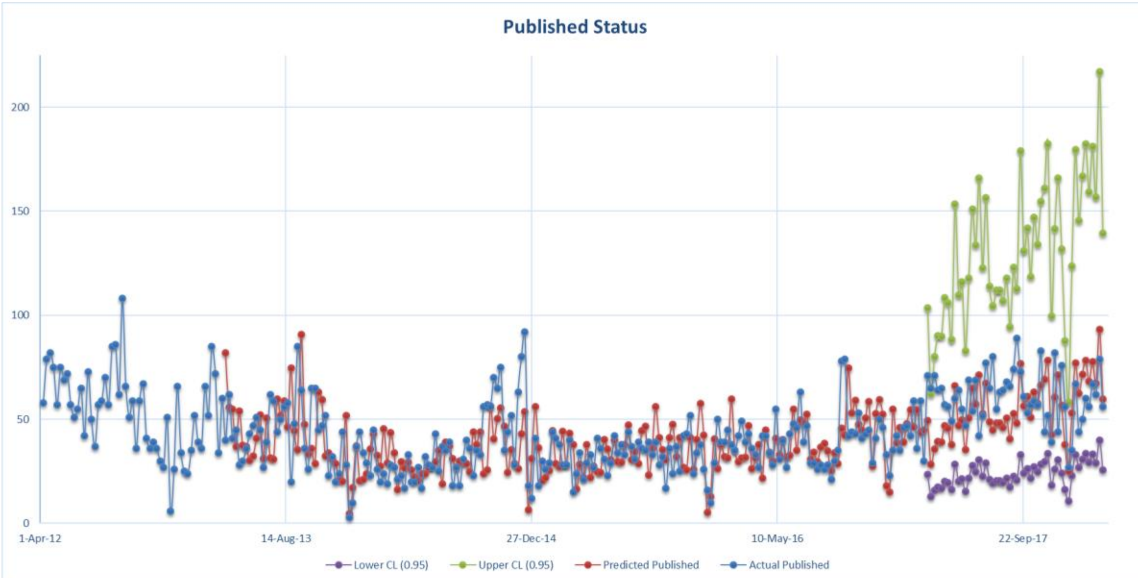

Figure 23: Published Status Forecast with 95% Confidence Interval

The above graph depicts the entire data set over the 309 weeks, including a 95% confidence interval for the forecast versus the actual request values in the last 52-weeks. In reality, there were 3150 Filled requests with a forecast for 2897. Therefore, the forecast was a conservative estimate by 253 requests or 8% less than actual.

The Cancelled Status

Air refueling requests that remain in the Cancelled status are those requests that were cancelled by either the tanker or receiver unit. Unfortunately, the system fails to account for why an AR event was cancelled, which could be for a number of reasons to include the inability of the tanker unit to support the air refueling event. Therefore, there are certainly unsupported requests in the Cancelled status that cannot be accounted for at this time.

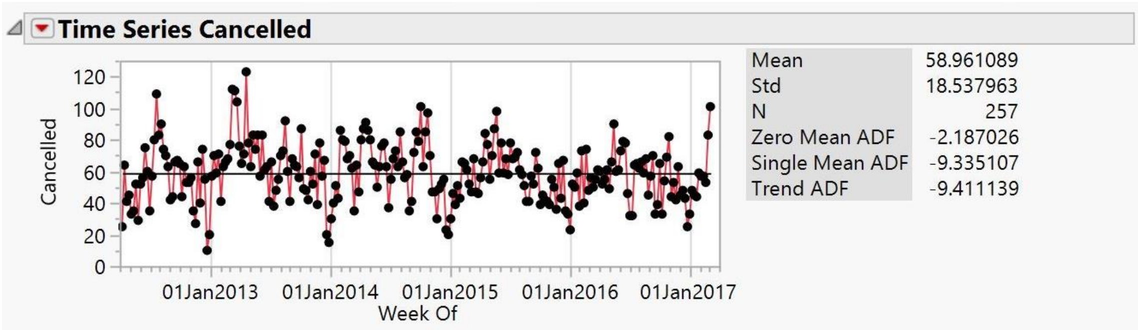

Figure 24: Time Series of Weekly Cancelled Status, 08 Apr 12 – 11 Mar 17

The above JMP graph of the Cancelled status over weekly intervals from Sunday through Saturday shows a fairly consistent distribution of data about the mean with some cyclical seasonality. The above graph of the Cancelled data has not been transformed in any way and yet still shows a high standard deviation along with seasonality to illustrate the difference from the figure below. In order to run multiple forecasting techniques, this data set was transformed with the natural log of the original observation decreasing the variance of the data about the mean.

Model Selection

Time Series Log[Cancelled]

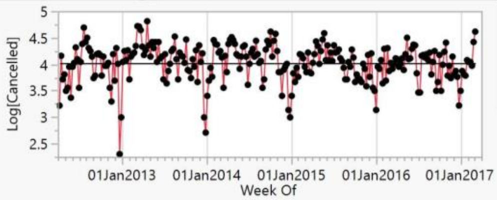

Mean	4.0208277
Std	0.3533596
N	257
Zero Mean ADF	-0.484238
Single Mean ADF	-9.67892
Trend ADF	-9.704816

Time Series Basic Diagnostics

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal ARIMA(2, 0, 0)(2, 0, 0)52	252	0.0819538	108.75940	126.50478	0.284	98.759397	1	1	5.799628	0.222581
Seasonal Exponential Smoothing(52, Zero to One)	202	0.0757818	138.22695	144.86319	-0.14	134.22695	2	2	6.666302	0.262050
Winters Method (Additive)	201	0.0761588	140.22695	150.18131	-0.14	134.22695	3	3	6.666302	0.262050
Simple Exponential Smoothing(Zero to One)	255	0.1106493	164.39601	167.94119	0.095	162.39601	4	4	6.729038	0.258706
Damped-Trend Linear Exponential Smoothing	253	0.1115246	168.39608	179.03161	0.095	162.39608	5	5	6.729088	0.258708
Double (Brown) Exponential Smoothing	254	0.120478	189.64748	193.18874	-0.03	187.64748	6	6	7.192615	0.276289
ARIMA(0, 0, 0)	256	0.1253508	196.63618	200.18525	-0.00	194.63618	7	7	7.034649	0.268104
I(1) No Intercept	256	0.1313543	206.85320	206.85320	-0.07	206.8532	8	8	7.183811	0.276975
Linear (Holt) Exponential Smoothing	253	0.126337	217.81752	224.90004	-0.15	213.81752	9	9	7.410497	0.279131

Figure 25: Cancelled Status Model Comparison

The model comparison chart above allows for the assessment of different forecasting techniques in JMP. Each of the forecasting techniques described in Chapter 3 was also applied to the Filled data. Moreover, a minimum of 1250 different iterations of seasonal ARIMA permutations was applied in accordance with the Box-Jenkins Methodology (available in Appendix G) in order to arrive at the lowest AIC/SBC values while verifying significant P values for each parameter. In this case, the seasonal ARIMA(2,0,0)(2,0,0)52 resulted in the lowest AIC/SBC values and the lowest MAPE and MAE error statistics used to develop the forecast from 08 Apr 12 to 11 Mar 17. This indicates that the seasonal ARIMA model is the best fit for the forecast data with the least number of parameters and the least overall error. Of note, the Cancelled status was the only one of the five that did not require differencing in order to reduce the variance for the forecasting techniques.

Time Series Log[Cancelled]

Model: Seasonal ARIMA(2, 0, 0)(2, 0, 0)52

Model Summary

DF	252	Stable	Yes
Sum of Squared Errors	20.6523561	Invertible	Yes
Variance Estimate	0.08195379		
Standard Deviation	0.28627573		
Akaike's 'A' Information Criterion	108.759397		
Schwarz's Bayesian Criterion	126.504777		
RSquare	0.28424204		
RSquare Adj	0.2728808		
MAPE	5.79962801		
MAE	0.22258126		
-2LogLikelihood	98.7593967		

Parameter Estimates

Term	Factor	Lag	Estimate	Std Error	t Ratio	Prob>|t|	Constant Estimate
AR1,1	1	1	0.3685796	0.0649041	5.68	<.0001*	
AR1,2	1	2	0.1406660	0.0646559	2.18	0.0305*	0.95609356
AR2,52	2	52	0.1628757	0.0644140	2.53	0.0121*	
AR2,104	2	104	0.3519093	0.0768860	4.58	<.0001*	
Intercept	1	0	4.0151520	0.0564173	71.17	<.0001*	

Forecast

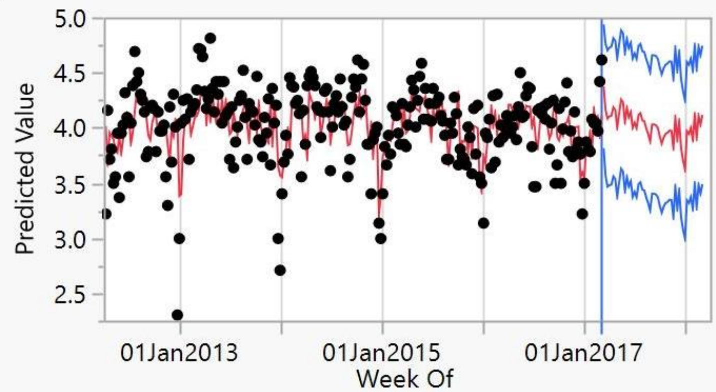

Residuals

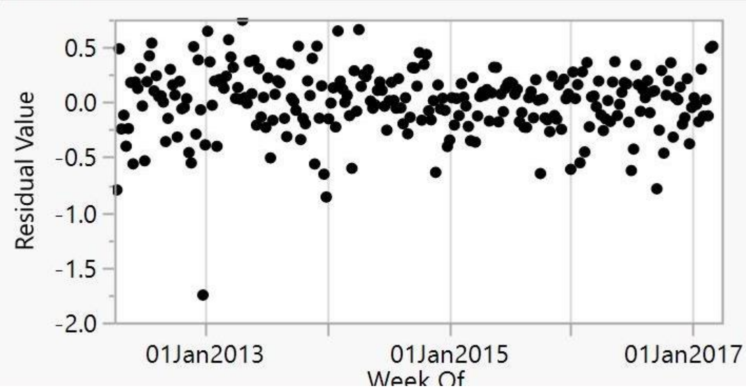

Figure 26: Cancelled Status Seasonal ARIMA(2,0,0)(2,0,0)52

Again, looking closer at the selected Filled seasonal ARIMA model, the basic information about the model is available. Each of the parameters is significant and well above a 95% significance level. Of note, the AR,1,2 and AR2,52 are the highest at 3.1% and 1.2%, respectively, although both still significant at the 95% level. Additionally, the depicted residuals appear to be evenly distributed about zero with a notable low outlier around Christmas 2012.

Model Validation

Based on the JMP output and forecast measures such as AIC/SBC, the Cancelled seasonal ARIMA model again appears to be the best fit for the data used during the forecast period. Again, the 52-weeks of validation data will be used to calculate the lowest MAPE and MAE error statistics. The resulting model represents the least error and, therefore, is the most accurate forecast for the subsequent 52-weeks from 12 March 17 to 10 March 18.

Table 8: Cancelled Status – 52-Week Validation Set Error Statistics

Model	MAPE	MAE
Mean Model / White Noise (0,0,0) + c	5.992025	0.236018
Random Walk Model (0,1,0)	13.984956	0.539294
Simple Exponential Smoothing (0 ≤ α ≤ 1)	8.129361	0.307461
Double (Brown) Linear Exponential Smoothing	30.313196	1.197508
Linear (Holt) Exponential Smoothing	6.412654	0.255705
Damped Trend Linear Exponential Smoothing	8.138256	0.307804
Seasonal Exponential Smoothing (52 weeks, (0 ≤ δ ≤ 1)	7.216771	0.283724
Winters Method (Additive)	7.216771	0.283724
Seasonal ARIMA (2,0,0)(2,0,0)52	5.727990	0.228003

As shown above, the seasonal ARIMA model both provided the best model of the forecast data and resulted in the least error for the 52-week validation data set, leading the researcher to conclude that it was the best model analyzed to forecast future Cancelled air refueling requests.

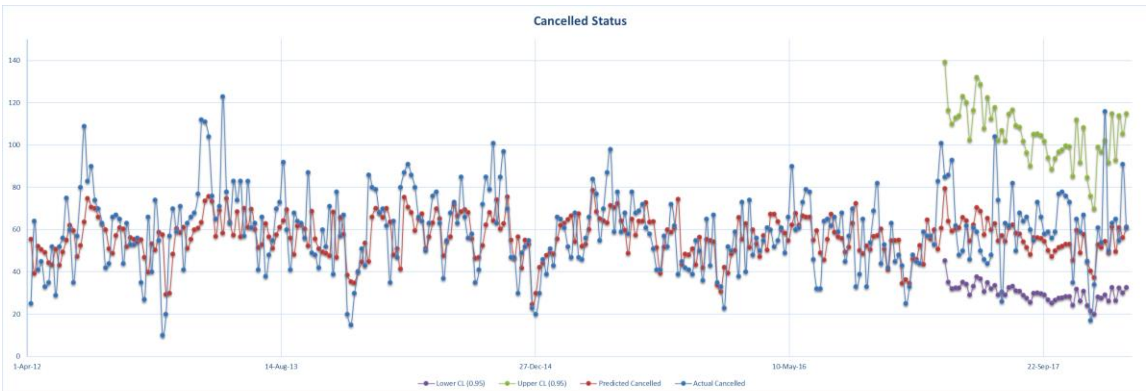

Figure 27: Cancelled Status Forecast with 95% Confidence Interval

The above graph depicts the entire data set overall 309 weeks, including a 95% confidence interval for the forecast versus the actual request values in the last 52-weeks. Of note, there were 3227 Filled requests with a forecast for 2931. Therefore, the forecast was a conservative estimate by 296 requests or 9% less than actual. Significantly, the Cancelled status was much more volatile than the other statuses due to increased unpredictability resulting in more forecasting error.

Forecasting

As stated above, each of the five ARST request statuses was best modeled by seasonal ARIMA models as verified by the lowest overall MAPE and MAE error when applied to the validation data set. There were differences in the parameters of the seasonal model between each of the statuses but, overall, the seasonal ARIMA proved to be the best model for forecasting future air refueling requests in the ARST as shown in the consolidated Table 9 below. The number of air refueling requests on a weekly or annual basis is helpful for long-term planners and answers the researcher's first research question, about which forecasting technique is best to predict future demand, seasonal ARIMA in this case. Since air refueling is accounted for financially by the number of airframe hours utilized in the expenditure of the service, flight hours executed is a useful metric for communicating future demand in financial terms. Therefore, the researcher's second question applies to methods to convert the number of forecast requests into the number of requested hours for air refueling with a reasonable degree of certainty.

Table 9: Air Refueling Requests by ARST Status, Difference, and Percent Change

Confirmed Status		Bought Status		Filled Status		Published Status		Cancelled Status	
Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted
3058	2863	3882	4392	8223	7970	3150	2897	3227	2931
195		-510		253		253		296	
6%		-13%		3%		8%		9%	
Total Actual Requests			21540		Total Predicted Requests			21053	
Difference (Actual - Predicted)				487					
Percent Change				2%					

To attempt to answer this second research question, the researcher analyzed multiple statistical measures of the hours requested in the ARST. The use of statistical measures as a method was chosen because a number of factors go into a specific hourly request for AR, to include the size or number of aircraft to be refueled, length of the track, operational or training requirements, etc. Furthermore, these varied requirements do not necessarily have anything to do with when a request is made in the time series or what status it is assigned, unlike the requests by the statuses that were forecast above. Therefore, multiple mean and median statistical figures were analyzed to determine the best hourly predictor based on the number of forecast requests by status. The researcher analyzed the statistical measures of one year, two years, three years, four years, and all data prior to the forecast period beginning 12 Mar 2017. These distributions and summary statistics are available in their entirety in Appendixes H – L. Furthermore, the researcher also provided hourly duration requests for the 52-week validation period from 12 Mar 17 – 10 Mar 18 in Appendix M for comparison. By plotting the hourly duration requests and reviewing the summary statistics for each category, it is apparent that there is a wide range of values from zero to several thousand. This results from input error or potentially from users not understanding or not seriously considering the actual duration of air refueling they are requesting. Whatever the case, although each range has significant outliers, the data is tightly grouped around the mean and median in each case. Interestingly, the median value is 60 minutes in every case except the one-year confirmed and validation confirmed ranges where the median value is 66 and 70 minutes, respectively. This is most likely a result of the request for one hour of air refueling service being a common practice for multiple aircraft and a fairly standard air refueling track length. In order to determine the best statistical measure for predicting hourly duration based on the number of forecasted air refueling requests, the researcher computed six test cases utilizing the median value of 60 minutes for all five statuses, followed by the average number of minutes requested by each requested status for all of the test data (08 Apr – 11 Mar 17), four years of test data (17 Mar 13 – 11 Mar 17), three years of test data (16 Mar 14 – 11 Mar 17), two years of test data (15 Mar 15 – 11 Mar 17), and one year of test data (13 Mar 16 – 11 Mar 17). Lastly, the average of the 52-weeks of validation data was computed as a mark of comparison with the six computed estimates. The comparison is shown below in Table 10. Interestingly, although the 60-

minute median in each category resulted in the closest hourly estimation for the Published category with an underestimate of 231 hours, the median technique also resulted in the worst overall estimation of hourly duration with an underestimation of 2,897 over the 52-week validation period. Even if the 66-minute Confirmed status median is used, as was seen in the 1-year distribution, the total hourly duration is still underestimated by 2,610 hours. Therefore, it is apparent that utilizing the mean over some range of values will result in a closer estimate of future hourly duration given a forecast of future requests.

Table 10: Air Refueling Hourly Duration Estimates

	Confirmed		Bought		Filled		Published		Cancelled		Total Hours	
Data	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted
Mean (60)	3,540	2,863	4,827	4,392	9,044	7,970	3,128	2,897	3,411	2,931	23,950	21,053
Difference	677		435		1074		231		480			2,897
All Test	3,540	2,675	4,827	4,633	9,044	9,426	3,128	2,658	3,411	2,677	23,950	22,068
Difference	865		194		-382		470		733			1,882
4 Years	3,540	3,060	4,827	4,707	9,044	9,594	3,128	2,777	3,411	2,787	23,950	22,925
Difference	480		120		-550		351		623			1,025
3 Years	3,540	3,166	4,827	4,806	9,044	9,869	3,128	2,768	3,411	2,889	23,950	23,498
Difference	374		21		-825		360		521			451
2 Years	3,540	3,249	4,827	4,980	9,044	10,588	3,128	2,791	3,411	3,138	23,950	24,745
Difference	291		-152		-1544		338		273			(795)
1 Year	3,540	3,282	4,827	5,439	9,044	9,534	3,128	2,874	3,411	3,074	23,950	24,203
Difference	258		-612		-490		255		337			(253)
Validation	3,540	3,314	4,827	5,461	9,044	8,766	3,128	2,877	3,411	3,098	23,950	23,515
Difference	226		-634		278		252		313			434

Reviewing the five distributions of the mean hourly requested duration of air refueling service showed that the averages over a shorter time span closer to the reserved data resulted in a better overall predictor of the total number of hours that would be requested. With the exception of the two years of data distribution that drastically overestimated the Filled status, the difference between the actual hours requested and predicted number of hours requested based on previous averages decreased as the data range decreased toward the preceding one year prior. It is worth noting that the one year prior only overestimated the number of hours requested by 253 or about 1%, which is much less than the control of the average validation data that underestimated by 434 hours. Additional testing in future years will be required to see if the previous year's distribution data remains the best estimator for hourly air refueling duration requests, although this technique is certainly a starting point in estimating air refueling hours given a forecast number of requests. The full table for one year of air refueling hourly duration request distribution is shown below in Table 11.

Table 11: Air Refueling Estimates (Averages Calculated in Minutes)

	Confirmed (68.8)		Bought (74.3)		Filled (71.8)		Published (59.5)		Cancelled (62.9)	
Data	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted
Requests	3058	2863	3882	4392	8223	7970	3150	2897	3227	2931
Hours	3,540	3,282	4,827	5,439	9,044	9,534	3,128	2,874	3,411	3,074
Difference	258		-612		-490		255		337	
% Change	7%		-13%		-5%		8%		10%	
Total Actual Hour Requests			23950		Total Predicted Hour Requests				24203	
Difference (Actual - Predicted)					-253					
Percent Change					-1%					

Research Question Analysis

This research focused on specific time series forecasting techniques addressed in Chapter 3 in order to determine which technique would best model the five identified air refueling statuses presented in the air refueling scheduling tool (ARST). The researcher utilized the MAPE and MAE error statistics in order to validate the best fit forecast model for each of the five statues by minimizing the error between the forecasting model and the 52-weeks of validation data from 12 March 17 to 10 March 18. In each of the five cases, the seasonal ARIMA model was the best overall fit of a forecasting model, resulting in the lowest MAPE and MAE error overall for the nine models evaluated.

The second research question focused on how the forecast number of requests could then be used as a predictor for the number of hours of air refueling service requested. Multiple methods were experimented with including numbers distributions and summary statistics. The mean of the previous one-year of each defined request status multiplied by the number of requests represented the best model evaluated. This method resulted in 90% accuracy or better for each status with the exception of the Bought status that was 87% accurate and a total hour overestimate of 253 hours or 1% more than actual. No other method evaluated resulted in a closer estimation or with all five categories greater than 90% accuracy.

Summary

This chapter presented multiple measures of forecasting comparison to include -2Log Likelihood, AIC, and SBC, along with the statistical error measures of MAPE and MAE. The researcher presented the initial data analysis and determination of weekly seasonality along with the concepts of differencing and data transformation. Next, each of the five ARST request statuses were presented, including their model selection in JMP and model validation including error analysis, and then graphed to fully understand how close the prediction was to the actual reserved data. In each case, the seasonal ARIMA model was the best fit for the forecast data and was validated with the lowest overall error utilizing the MAPE and MAE error statistics. Lastly, in accordance with the second research question, the total number of hours requested was estimated utilizing the previous one year's mean hourly duration requests by ARST status multiplied by the forecast number of requests in each status. This method resulted in the best estimation overall with regard to the actual number of hours requested and represented a baseline for future analysis along with a mathematically supported hourly estimation of forecast air refueling demand that could then be used for forthcoming decision making. A discussion of the financial implications, conclusions, and implications of this research will be included in the next chapter.

V. Conclusions and Recommendations

"All models are wrong, some are useful" – George E. P. Box

Chapter Overview

This research into tanker training demand forecasting not only fills a void in the current tools available to USTRANSCOM, AMC, and the 618th AOC, but also allows for forecasts that can be updated to obtain more precise results as more data is gathered in future years. Additionally, the forecast models can be adapted for both training and operational requirements in order to build a more holistic view of the tanker force supply and demand. This research sought to define the most accurate of nine forecasting techniques evaluated for predicting future air refueling requests in five different statuses as defined by the 618th AOC's air refueling scheduling tool. Furthermore, based on the most precise forecasting model, the research pursued a calculation by which the number of forecast air refueling flight hours could be predicted with up to 90% accuracy annually.

This chapter will focus on the specific conclusions of the above research, concentrating on the significant applications of the forecast models and specifically what they could indicate for future demand. Specific recommendations for action are presented to include an analysis and application of DoD fixed wing hourly reimbursement rates and their application for future air refueling service cost comparisons. Lastly, various avenues for future research into the tanker supply and demand problem will be explored, to include 1) a review of request prioritization, 2) the multitude of barriers to air refueling service market entry, and 3) a lack of a DoD tanker pricing strategy. This last issue not only results in the DoD/USAF being unable to participate in other air refueling markets but also has the potential to cause the DoD to be ill-prepared to work with commercial air refueling organizations if/when the latter bring a service to market.

Conclusions of Research

The first research question focused on identifying the best forecasting model for the various air refueling training requests for each of the five designated statuses defined by the air refueling scheduling tool. After evaluating nine different forecasting models, multiple goodness-of-fit measures demonstrated that the seasonal ARIMA forecasting models were the best fit for the modeled data. Furthermore, the MAPE and MAE error measures used to validate the forecast models resulted in the lowest error in every case for the seasonal ARIMA models with one exception. The MAE for the Bought status was the second lowest by four one-thousandths of a point with a much lower MAPE for comparison. Overall, these low error statistics confirm that the seasonal ARIMA models are the best models to represent future air refueling demand.

Secondly, the research evaluated six different statistical measures of hourly duration data in order to predict future requested flying hour duration based on the forecast number of requests by ARST status. The results showed that in four of the five categories the future requested hours could be predicted within 90 percent accuracy. The fifth category (Bought) was predicted with 87 percent accuracy. Overall, using the previous year's data resulted in an overestimation of 253 flight hours or 1% off from the actual 23,950 hours requested, as shown in Table 11.

Significance of Research

Not only was the seasonal ARIMA model the most accurate model evaluated, but it also highlighted a significant trend with regard to the number of air refueling requests. The researcher noticed that in four of the five status categories, the seasonal ARIMA model under forecasted the future number of requests; only the Bought status was over forecast as noted in Table 9. This could be the result of an increasing trend in the number of air refueling requests. Therefore, the researcher re-examined the request data over the period used for the forecast. Five sets of 52-weeks of request data are represented below in Figure 28 with the last set of data 3/12/17 – 3/10/18 being the validation data set.

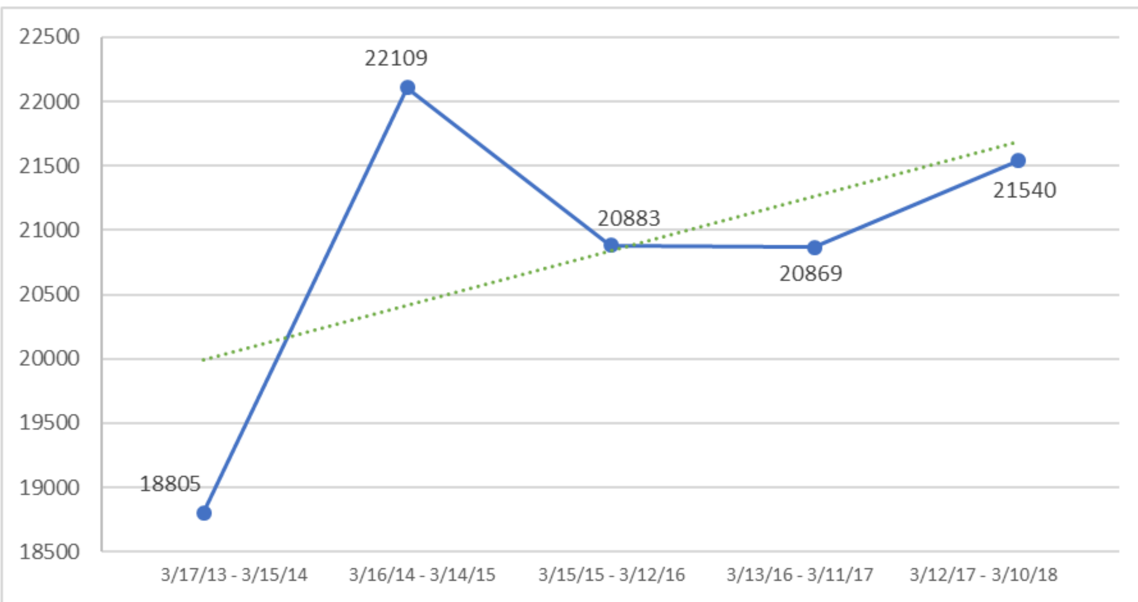

Figure 28: Total Number of ARST Requests in 52-Week Periods

This data represents a historical low in the 2013 timeframe due in part to sequestration and also potentially to the ARST program being new to the tanker and receiver community as of 2010, which could have resulted in fewer requests being captured. In 2014 – 2015 the number of requests spiked, again potentially due to increased training that was delayed or suspended due to sequestration in 2013. March 2015 – 2017 showed very stable total requests with a slight jump in the following year. This increase in the number of total air refueling requests from Mar 2017 to Mar 2018 is also the period of time used to validate the forecast models and could explain why most of the categories were under forecasted. It also appears that there is an upward trend in the number of air refueling requests represented by the green trendline. Additional future years of data will be required to determine if this upward trend will continue; however, it seems logical as more air refueling requests are being captured in the ARST system along with greater overall training requirements for the Joint Force. The total number of air refueling requests by ARST category over the same range are displayed in Figure 29.

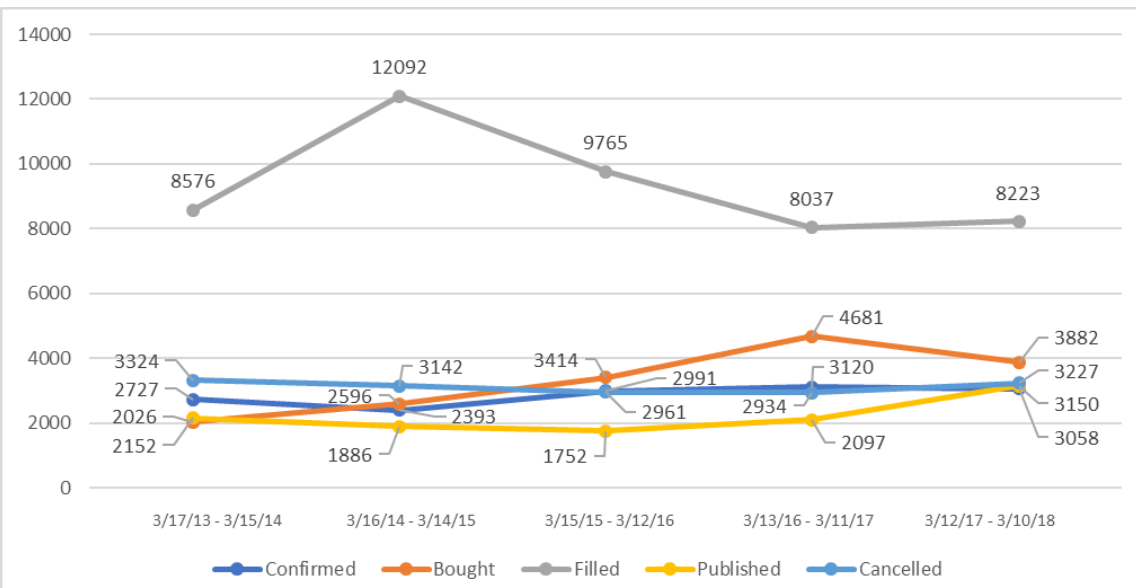

Figure 29: Total Number of Requests by ARST Status in 52-Week Periods

This representation starts to explain some of the overestimation in the Bought status as it was on an upward trend until March of 2017, leading the model to predict a higher than actual value for the Bought Status. Significantly, this representation reveals that the Confirmed and Cancelled requests are fairly consistent across the data rage and the Filled requests have leveled out in the last year. The main upward trend is the Published requests that represent unsupported air refueling requests. Again, additional data for future years will be required to see if this upward trend continues. Moreover, years of study could refine the above forecasting and continue to build upon the time series data set. Lastly, it represents an opportunity to execute more sorties with a potentially growing unsupported market, thus increasing readiness of the Joint Force or for outside support from a future commercial contractor.

Recommendations for Action

Planners at USTRANSCOM, AMC, and the 618th AOC have an opportunity to utilize these forecasting models not only for tanker training but also for operational use outside of training and other larger data sets that were not available to this researcher. Moreover, in order to refine the forecasts and, thereby, increase the accuracy of the output, several more years of data need to be gathered and analyzed.

The second part of this research focused on being able to communicate, in financial terms, what the air refueling requests and, in particular, the unsupported requests mean or could mean in the future. The objective, given an ability to forecast the number of requests by status, was to convert those requests into a predicted number of flight hours of service. That was accomplished with a reasonable degree of certainty with the results depicted in Table 11. The researcher chose flight hours as the units of analysis because the DoD publishes fixed wing reimbursement rates for all aircraft each fiscal year, which equates to the hourly rate of operating the aircraft for one additional hour.

To put that another way, the DoD already has a quantitative financial measure for aircraft flight hours. Several Air Mobility Command (AMC) aircraft have alternate rates that they can charge other uniformed services or agencies for operations such as airlift under the Defense Working Capital Fund or Transportation Working Capital Fund (TWCF), but USAF tankers completing a refueling mission do not fall under a separate working capital fund at this time. Therefore, the researcher chose the DoD component of the fixed wing reimbursable rate published each fiscal year by the Office of the Secretary Defense (OSD) for the KC-135 and KC-10 as a baseline for this discussion. The rate for the fiscal year 2019 equates to $13,419 for the KC-135R, $13,463 for the KC-135T, and $16,078 for the KC-10A (McAndrew, 2018, p. 4). Moreover, it is important to understand that this rate does not include many sustainment, logistics, operational, ownership, or other equipment capitalization costs; it is only a function of Aviation Fuel, Other Operational Material, Consumable Materials and Repair Parts, Depot Level Repairs (DLR), Intermediate Level Maintenance, Depot Maintenance, and Contract Maintenance Services as depicted in Figure 30 below (AFCAA/FMCY, 2018, p. 10).

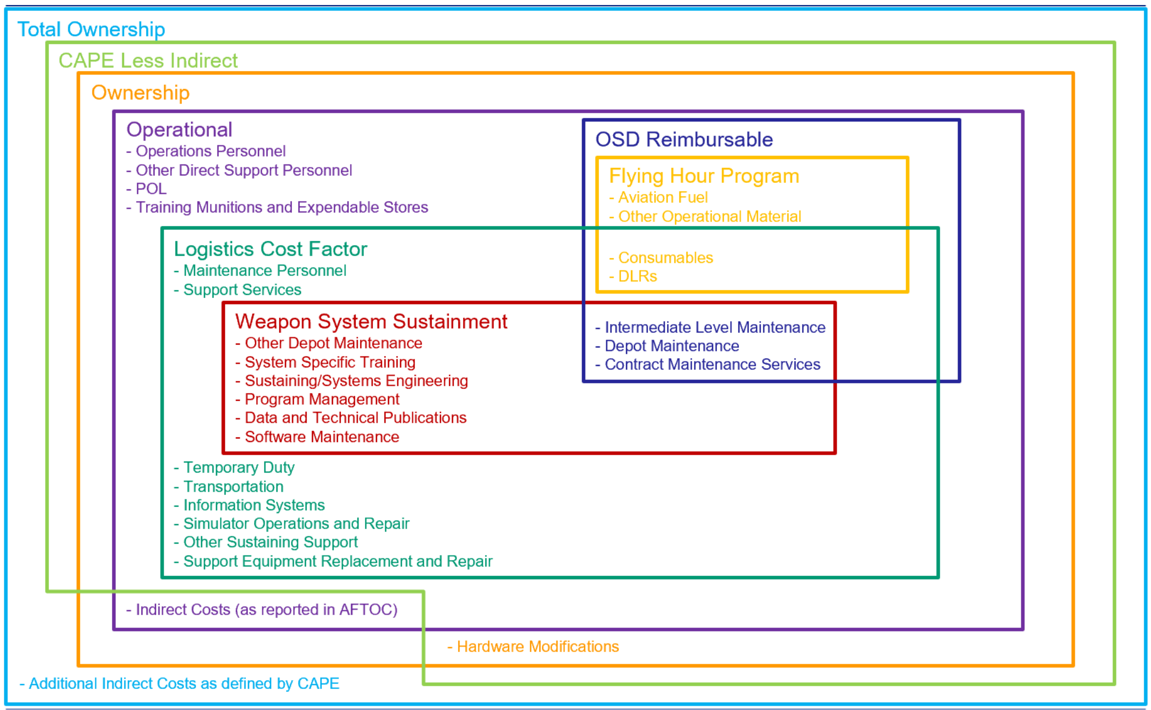

Figure 30: DoD Cost Metric Categories

Source: (AFCAA/FMCY, 2018, p. 12)

As is apparent, there are a number of other costs that are not included in the fixed wing reimbursable rate that would have to be considered if the DoD choose to compare costs directly with a commercial provider at some later date, yet this comparison cannot be accomplished until a commercial organization develops or acquires a boom tanker and, moreover, is outside of the scope of this research. What can be addressed at this time is what it would cost the DoD financially to execute the unsupported flight hours under the Published status in the ARST utilizing existing aircraft, sustainment, and support facilities and personnel. It is important to note in any such analysis that the flight hours predicted as part of the research and analysis in Chapter 4 only equates to the flying hours used during the process of aerial refueling with a receiver and does not include the flight time required for takeoff, setup en route to an air refueling track, recovery back to an airfield, and approach and landing. All of these events add additional flight time to the tanker aircraft. Therefore, again as a baseline for discussion, the researcher assumed a conservative forty-five minutes for takeoff and flight to the air refueling track and another forty-five minutes to return to an appropriate airfield, approach, and landing. The process of adding an additional one hour and thirty minutes to each request will more fairly account for the total aircraft flight time required for each air refueling request and therefore a more appropriate cost. An example of this calculation is depicted in Table 12.

Table 12: Actual and Forecast Reimbursement Cost for Unsupported Requests

Published ARST Requests for Air Refueling (Unsupported)						
Actual	**3150**					
Forecast	**2897**					
Published Air Refueling Hour Estimates (Unsupported)						
Actual	**3128**		**Actual + 1.5**		**7853**	
Forecast	**2874**		**Forecast + 1.5**		**7219.5**	
Estimated and Actual Fixed Wing Reimbursement Cost for Unsupported Requests						
Airframe	**KC-135R**		**KC-135T**		**KC-10A**	
FY19 Rate	**$13,419**		**$13,463**		**$16,078**	
	Actual	**Estimated**	**Actual**	**Estimated**	**Actual**	**Estimated**
	$105,379,407	**$96,878,471**	**$105,724,939**	**$97,196,129**	**$126,260,534**	**$116,075,121**
Difference	**$8,500,937**		**$8,528,811**		**$10,185,413**	
Percent Change	**8%**		**8%**		**8%**	

The above baseline for calculating reimbursement costs would need to be modified based on a more appropriate mix of tankers available, but the concept is the same. For example, tanker support would not come solely from KC-10As or KC-135Ts, as both of these aircraft represent a smaller portion of the fulfillment of any unsupported requests as these airframes represent a smaller portion of the total tanker force. However, even accounting for the KC-10A and KC-135T as approximately 13% and 12% of the tanker force, respectively, it is still approximately $100M per year using this cost metric to execute known unsupported requests with the above assumptions. Furthermore, again because it is important, this all assumes that the USAF has the aircraft, crews, and other operations and support capacity to execute with this cost metric. At the end of the day, it is more likely that additional tanker capacity will be required in the future. It is also likely that, when a commercial tanker is brought to the market, calculations such as those in Table 12 will be helpful in not only forecasting the air refueling hours required by the DoD but also helpful in making a fair cost comparison between USAF air refueling costs and any proposed commercial service.

Lastly, after reviewing several years of DoD fixed wing reimbursable rate for the tanker force, the researcher noted that the costs appear to be increasing. This trend is likely due to the increased cost of aviation fuel and the steadily rising maintenance costs as depicted in Figure 31.

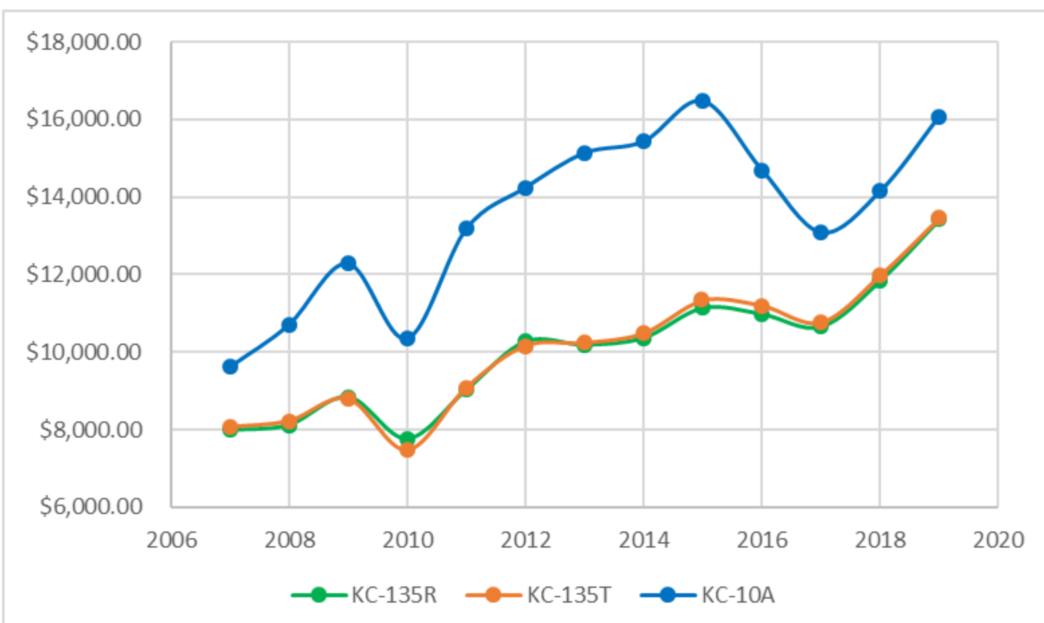

Figure 31: DoD Fixed Wing Reimbursement Rates (FY 2007 – 2019)

Source: (McAndrew, 2018)

In order to truly understand the rate trend, the researcher corrected the values for inflation via the Bureau of Labor and Statistics Consumer Price Index (CPI) as seen in Figure 32 below.

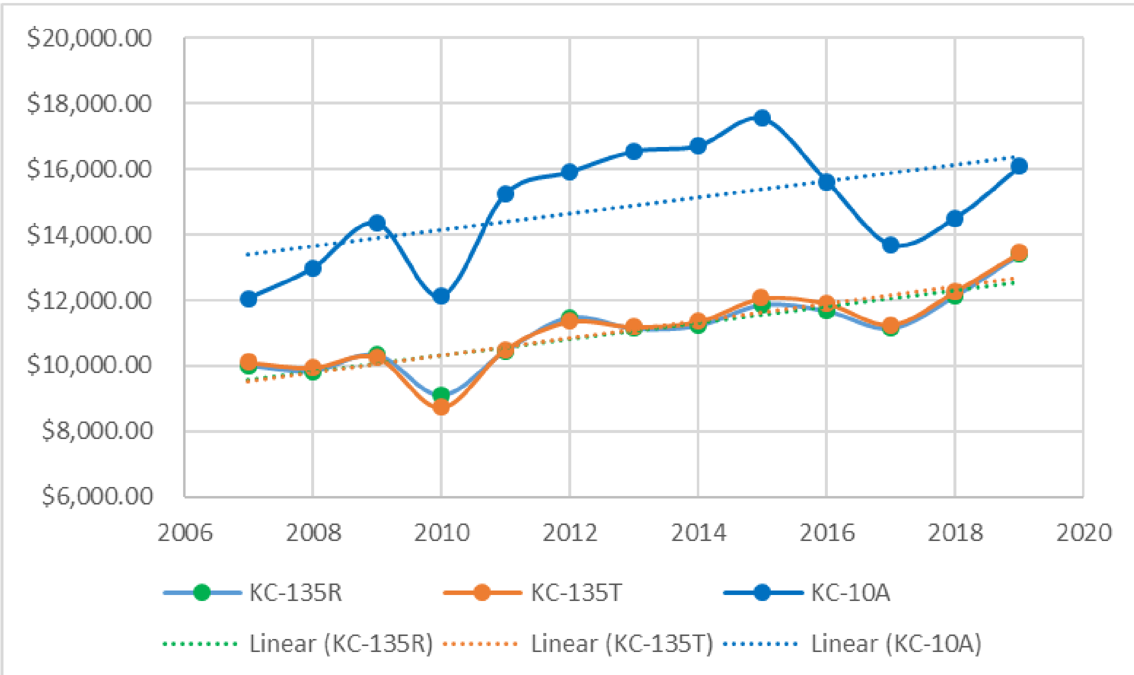

Figure 32: CPI Inflation Corrected Reimbursement Rates

Source: (McAndrew, 2018)

Even with the reimbursement rates corrected for inflation, there are similar positive linear trendlines for each of the three airframes studied; and the annual costs by specific tanker airframe are increasing over time. Moreover, the actual number of tankers available is not projected to increase in the near-term and may even decrease as requirements have the potential to increase; requests for service are clearly showing an overall increase. All of these trends point to an opportunity for the DoD to leverage commercial air refueling service to fill near-term gaps and moderate spikes in demand for specific peacetime training opportunities.

Recommendations for Future Research

Future research in the USAF air refueling tanker supply and demand problem should continue to evaluate the feasibility of innovative and non-traditional ideas not only to increase supply but also to understand demand signals from receiver units better. This research primarily focused on better understanding demand from a forecasting perspective; however, there are multiple other elements that could help planners better understand the underlying demand signals. Starting with the data this researcher analyzed in the ARST, the database should be expanded to include prompts for users to explain why a sortie is cancelled (lack of availability of aircraft, lack of aircrew, priority, etc.) or if the sortie is not "Bought" from the Filled status. Multiple feedback mechanisms built into the database could allow the 618th AOC a great deal more specific information on which sorties were unsupported and why, for future research and analysis.

Secondly, the DoD Air Refueling Support Priority System is governed by CJCSI 4120.02D and is the system through which USTRANSCOM validates and manages air refueling requirements with the support of AMC and the 618th AOC. This system is required because there needs to be a fair process to allocate limited resources to the highest priority users. The process needs to be more transparent, and users of the system, particularly non-USAF users, need to be better trained in the benefits of the system. This lack of knowledge has probably led to some receiver units not receiving needed service and ultimately no longer requesting service because of false perceptions that they will not receive service even if requested. There is an opportunity for USTRANSCOM, AMC, and the 618th AOC to build transparency in the system, some of which already is happening through the ARST system to advertise air refueling opportunities, build knowledge about the request system, and result in a more successful tanker and receiver matches. Moreover, aspects of this research could be continued to include more research into the weekly request rate. For example, more capacity appears available during the weekends, Mondays, and Fridays than midweek, according to Figure 4. Simple concepts such as this, packaged together for our Joint Force, could result in less frustration about the process and more efficient use of our limited resources.

While the above future research strategies are designed to understand more clearly and to moderate the demand, the supply side of the equation could be supplemented via commercial contracted air refueling service. The likelihood of this concept succeeding is directly related to the support received from the DoD, USTRANSCOM, and AMC. This support is required not only because the DoD owns 100% of the current market but also because of enormous barriers to market entry for a commercial air refueling service. These barriers include everything from a lack of a Federal Aviation Regulation (FAR) for commercial air refueling to likely interference from large aviation industry contractors like the Boeing Corporation, as well as numerous government policies that limit the sale or lease of military equipment. Many of these issues are policy related and could be solved if the commercial industry had DoD support; to date, such support has failed to materialize despite commercial interest since the mid-1990s. Each of these issues could be researched in detail as interest continues to grow throughout the DoD.

As previously stated, USAF air refueling service is not currently governed by TWCF and only charges receiver units for the fuel transferred and not the flight hours expended as they are considered training for the aircrew. The costs of the tanker flight hours are covered by operational and maintenance (O&M) funds. Future research in the benefits and consequences of a TWCF-like program for the USAF tanker fleet could be very beneficial on several fronts. First of all, it would modify demand signals as some fighter units may reduce requests to what they actually need versus what is convenient due to actually having to use command funds for the transit. For example, an overseas fighter deployment from Luke AFB to Aviano AB Italy may only be requested from the East coast to Lajes Field, Azores, (only the long distance over water portion where landing fields for refueling are not available) versus the entire distance. This has the potential to reduce overall demand and use the tanker aircraft more efficiently, where they are needed. Moreover, assigning an appropriate value for air refueling service would allow the USAF to fully participate in air refueling exchanges such as the Movement Coordination Centre Europe's (MCCE) Air Transport, Air to Air Refueling, and other Exchanges of Services (ATARES) program. These programs could be an important first step into laying the groundwork for the DoD to be prepared to incorporate commercial air refueling service once it is brought to the market.

Summary

This research established the best forecasting method as the seasonal ARIMA model out of the nine evaluated for predicting future air refueling requests by ARST status. Six methods then were evaluated for estimating the number of hours of air refueling duration expected annually for each ARST status with a focus on those requests that were not supported and, therefore, presented an opportunity to increase readiness training and potentially be accomplished by a commercial contractor if and when that capability is finally available to the market. The researcher then presented the significance of an increasing trend of unsupported requests along with a consistently increasing trend of the fixed wing reimbursement rate corrected for inflation. While more analysis of future years of data will be required to continue to validate the forecasting models, an example of how to perform a cost comparison of multiple tanker options was presented with the intent for an easier comparison against commercial air refueling services when they come available. Lastly, the researcher presented multiple directions for future research aimed both at better understanding the demand signals and also at supporting the supply of air refueling tankers through research into both policy efforts and funding changes. Ultimately, the current USAF air refueling fleet is facing likely increased requirements, more requests for service, and similar or fewer tankers to accomplish the task. Continued research into the supply and demand problems can and will result in solutions to these problems and potentially create alternatives, rather than just failing to support requests, or worse yet, making our aircrews work longer and harder. Having the means to consider new opportunities to address this supply and demand problem will hopefully allow the USAF to fulfill its air refueling mission to the Joint Force better.

Appendix A: Quad Chart

Appendix B: ARIMA Model Comparison

Daily, Monthly, and Weekly Forecasting Comparison

The following daily and monthly forecasts were the researcher's original attempt at defining the interval upon which to forecast. Note the daily forecasts compute extremely high AIC/SBC and no MAPR due to zero observations on particular dates. The monthly forecasts have a limited number of observations and produce negative AIC/SBC values.

Daily

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
✓		Seasonal ARIMA(0, 1, 18)(0, 1, 18)52	1371	9.3570054	7556.4601	7750.7074	0.687	7482.4601	1.000000		.	2.731424
✓		Seasonal ARIMA(0, 1, 7)(0, 1, 7)52	1393	14.396294	8047.0126	8125.7615	0.553	8017.0126	0.000000		.	3.361691
✓		Seasonal ARIMA(0, 1, 5)(0, 1, 5)52	1397	17.508744	8202.2015	8259.9506	0.498	8180.2015	0.000000		.	3.559857
✓		Seasonal ARIMA(0, 0, 0)(0, 1, 7)12	1441	23.327236	8765.0099	8807.2389	0.370	8749.0099	0.000000		.	4.084376
✓		Seasonal Exponential Smoothing(52, Zero to One)	1406	37.207405	9263.1875	9273.6873	-0.07	9259.1875	0.000000		.	5.173543
✓		Winters Method (Additive)	1405	37.233887	9265.1875	9280.9372	-0.07	9259.1875	0.000000		.	5.173543
✓		Seasonal Exponential Smoothing(12, Zero to One)	1446	36.802269	9389.7101	9400.2660	0.035	9385.7101	0.000000		.	4.821486
✓		Winters Method (Additive)	1445	40.010519	9518.6060	9534.4398	-0.06	9512.606	0.000000		.	5.554903
✓		Simple Exponential Smoothing(Zero to One)	1459	39.687893	9522.3887	9527.6749	0.002	9520.3887	0.000000		.	5.434729
✓		Damped-Trend Linear Exponential Smoothing	1457	39.742372	9526.3887	9542.2472	0.002	9520.3887	0.000000		.	5.434729
✓		Double (Brown) Exponential Smoothing	1458	39.824815	9534.4110	9539.6965	-0.01	9532.411	0.000000		.	5.450732
✓		Linear (Holt) Exponential Smoothing	1457	39.818562	9535.2028	9545.7738	-0.01	9531.2028	0.000000		.	5.450894

Monthly

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
✓		IMA(1, 1) No Intercept	46	0.0307241	-28.81429	-26.96414	0.088	-30.81429	0.396306		2.562784	0.139840
✓		Simple Exponential Smoothing(Zero to One)	46	0.0307241	-28.81429	-26.96414	0.088	-30.81429	0.396306		2.562784	0.139840
✓		IMA(1, 1)	45	0.0314033	-26.81602	-23.11572	0.088	-30.81602	0.145919		2.562362	0.139798
✓		Damped-Trend Linear Exponential Smoothing	44	0.0324741	-24.80350	-19.25305	0.083	-30.8035	0.053346		2.574044	0.140408
✓		Linear (Holt) Exponential Smoothing	44	0.0322999	-19.56648	-15.90920	-0.03	-23.56648	0.003890		2.753312	0.150514
✓		Double (Brown) Exponential Smoothing	45	0.0342284	-19.31509	-17.48645	-0.08	-21.31509	0.003430		2.760636	0.150893
✓		Seasonal Exponential Smoothing(12, Zero to One)	33	0.032016	-15.78551	-12.67481	-0.37	-19.78551	0.000587		2.772228	0.152875
✓		Winters Method (Additive)	32	0.0330165	-13.78551	-9.119462	-0.37	-19.78551	0.000216		2.772228	0.152875

The following outputs are for the same data; the first was without a logarithmic data transformation, while the second was transformed via the natural log. Note the difference in variance along with much lower AIC/SBC values.

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
✓		Seasonal ARIMA(1, 1, 1)(1, 1, 0)52 No Intercept	201	234.91852	1713.9702	1723.9246	-0.31	1707.9702	0.119921		37.999911	12.537200
✓		Seasonal ARIMA(1, 1, 1)(1, 1, 1)52 No Intercept	200	231.9521	1714.6207	1727.8932	-0.30	1706.6207	0.086625		37.642359	12.476097
✓		Seasonal ARIMA(2, 1, 1)(1, 1, 0)52 No Intercept	200	234.47737	1714.8390	1728.1114	-0.30	1706.839	0.077670		37.137134	12.443354
		Seasonal ARIMA(1, 1, 1)(2, 1, 0)52 No Intercept	200	233.33847	1714.9705	1728.2430	-0.30	1706.9705	0.072725		37.752986	12.503639
		Seasonal ARIMA(2, 1, 1)(1, 1, 1)52 No Intercept	199	232.07498	1715.6703	1732.2609	-0.30	1705.6703	0.051254		36.918045	12.395629
		Seasonal ARIMA(1, 1, 1)(3, 1, 0)52 No Intercept	199	230.09033	1715.9683	1732.5589	-0.30	1705.9683	0.044158		37.429446	12.378324
		Seasonal ARIMA(2, 1, 1)(2, 1, 0)52 No Intercept	199	233.44256	1716.0167	1732.6073	-0.30	1706.0167	0.043102		37.028825	12.425991
		Seasonal ARIMA(1, 1, 1)(2, 1, 1)52 No Intercept	199	209.66977	1716.0419	1732.6325	-0.30	1706.0419	0.042564		37.489186	12.408913
		Seasonal ARIMA(3, 1, 1)(1, 1, 0)52 No Intercept	199	234.9173	1716.1800	1732.7706	-0.30	1706.18	0.039723		37.215997	12.450609
		Seasonal ARIMA(1, 1, 1)(0, 1, 1)52 No Intercept	201	228.84567	1716.5523	1726.5066	-0.32	1710.5523	0.032977		37.285776	12.367590
		Seasonal ARIMA(2, 1, 1)(3, 1, 0)52 No Intercept	198	228.92917	1716.7439	1736.6526	-0.29	1704.7439	0.029965		36.579324	12.274908
		Seasonal ARIMA(2, 1, 1)(2, 1, 1)52 No Intercept	198	209.72647	1716.9386	1736.8473	-0.30	1704.9386	0.027185		36.727706	12.329608
		Seasonal ARIMA(3, 1, 1)(1, 1, 1)52 No Intercept	198	232.66911	1717.0892	1736.9979	-0.29	1705.0892	0.025213		37.049808	12.430996
		Seasonal ARIMA(3, 1, 1)(2, 1, 0)52 No Intercept	198	234.12836	1717.4517	1737.3604	-0.29	1705.4517	0.021033		37.131878	12.449769
		Seasonal ARIMA(2, 1, 1)(0, 1, 1)52 No Intercept	200	228.84599	1717.6236	1730.8961	-0.32	1709.6236	0.019301		36.584798	12.287037
		Seasonal ARIMA(3, 1, 1)(3, 1, 0)52 No Intercept	197	228.26506	1717.8966	1741.1235	-0.29	1703.8966	0.016838		36.727480	12.305810
		Seasonal ARIMA(1, 1, 1)(4, 1, 0)52 No Intercept	198	1.3126201	1717.9583	1737.8671	-0.30	1705.9583	0.016326		37.427960	12.377655
		Seasonal ARIMA(1, 1, 1)(3, 1, 1)52 No Intercept	198	205.89825	1717.9672	1737.8759	-0.30	1705.9672	0.016254		37.429126	12.378188
		Seasonal ARIMA(4, 1, 1)(1, 1, 0)52 No Intercept	198	236.11665	1718.1356	1738.0443	-0.30	1706.1356	0.014942		37.144556	12.421381
		Seasonal ARIMA(3, 1, 1)(2, 1, 1)52 No Intercept	197	209.91629	1718.2295	1741.4564	-0.29	1704.2295	0.014256		36.870185	12.364426
		Seasonal ARIMA(2, 1, 1)(4, 1, 0)52 No Intercept	197	0.3203738	1718.7326	1741.9595	-0.29	1704.7326	0.011085		36.585435	12.275387

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
☑	☐	Seasonal ARIMA(1, 1, 2)(1, 1, 0)52	199	0.1435461	208.09099	224.68159	0.093	198.09099	0.040876		9.246431	0.316599
☐	☐	Seasonal ARIMA(1, 1, 2)(1, 1, 2)52	197	0.1275094	209.38064	232.60748	0.101	195.38064	0.021450		9.148346	0.312746
☐	☐	Seasonal ARIMA(2, 1, 1)(1, 1, 0)52	199	0.144668	209.38066	225.97126	0.087	199.38066	0.021450		9.244203	0.316767
☐	☐	Seasonal ARIMA(1, 1, 2)(1, 1, 1)52	198	0.1434427	209.64837	229.55709	0.094	197.64837	0.018762		9.244105	0.316268
☐	☐	Seasonal ARIMA(1, 1, 1)(1, 1, 0)52 No Intercept	201	0.1488412	209.69939	219.65375	0.073	203.69939	0.018290		9.285754	0.318511
☐	☐	Seasonal ARIMA(2, 1, 2)(1, 1, 0)52	198	0.1441446	209.71891	229.62763	0.096	197.71891	0.018112		9.219734	0.315748
☐	☐	Seasonal ARIMA(1, 1, 2)(2, 1, 0)52	198	0.1438471	209.81466	229.72338	0.094	197.81466	0.017265		9.248561	0.316521
☐	☐	Seasonal ARIMA(1, 0, 1)(1, 1, 0)52	201	0.1477288	210.04406	223.33610	0.084	202.04406	0.015394		9.262517	0.314895
☐	☐	Seasonal ARIMA(1, 1, 2)(1, 1, 0)52 No Intercept	200	0.1482357	210.12460	223.39708	0.078	202.1246	0.014787		9.218173	0.316846
☐	☐	Seasonal ARIMA(2, 1, 1)(1, 1, 0)52 No Intercept	200	0.1486082	210.34407	223.61655	0.077	202.34407	0.013250		9.226350	0.316780
☐	☐	Seasonal ARIMA(2, 0, 1)(1, 1, 0)52	200	0.1475591	210.43806	227.05311	0.094	200.43806	0.012642		9.143014	0.311310
☐	☐	Seasonal ARIMA(1, 1, 2)(0, 1, 1)52	199	0.1410218	210.60250	227.19310	0.078	200.6025	0.011644		9.259274	0.316468
☐	☐	Seasonal ARIMA(1, 1, 2)(2, 1, 1)52	197	0.1303883	210.62257	233.84941	0.097	196.62257	0.011528		9.200971	0.314583
☐	☐	Seasonal ARIMA(2, 1, 1)(1, 1, 2)52	197	0.1274447	210.65521	233.88205	0.095	196.65521	0.011341		9.149128	0.313007
☐	☐	Seasonal ARIMA(1, 1, 2)(0, 1, 2)52	198	0.1439475	210.71358	230.62230	0.089	198.71358	0.011015		9.263495	0.316855
☐	☐	Seasonal ARIMA(1, 0, 1)(1, 1, 2)52	199	0.1283309	210.76423	230.70229	0.094	198.76423	0.010739		9.230328	0.313292
☐	☐	Seasonal ARIMA(0, 1, 3)(1, 1, 0)52 No Intercept	200	0.1489018	210.85022	224.12270	0.075	202.85022	0.010287		9.213806	0.316004
☐	☐	Seasonal ARIMA(1, 0, 2)(1, 1, 0)52	200	0.1477274	210.87729	227.49234	0.092	200.87729	0.010149		9.160656	0.311710
☐	☐	Seasonal ARIMA(1, 1, 2)(2, 1, 2)52	196	0.1206916	210.90249	237.44745	0.099	194.90249	0.010022		9.130501	0.312135
☐	☐	Seasonal ARIMA(1, 1, 1)(1, 1, 0)52	200	0.148947	210.97806	224.25054	0.076	202.97806	0.009650		9.380163	0.320728
☐	☐	Seasonal ARIMA(2, 1, 1)(1, 1, 1)52	198	0.1446743	210.98634	230.89506	0.089	198.98634	0.009611		9.244413	0.316549

Appendix C: Confirmed Time Series: Best Seasonal ARIMA Values from JMP

The following JMP output includes the best seasonal ARIMA values computed for the Confirmed status sorted by lowest AIC followed by lowest SBC. The researcher explored all options in accordance with the Box-Jenkins Methodology in order to produce a model with significant parameters that was the least complex possible. This required at least 625 permutations for an equation with an intercept and another 625 without an intercept. Additional permutations were calculated to verify the best result.

Time Series Log[Confirmed]

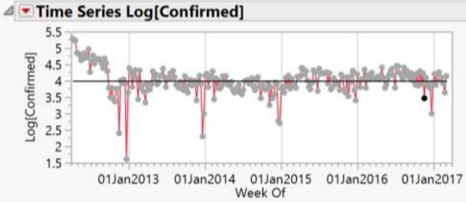

Mean	4.0037692
Std	0.415042
N	257
Zero Mean ADF	-1.016687
Single Mean ADF	-9.438526
Trend ADF	-9.452477

Time Series Basic Diagnostics

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
		Seasonal ARIMA(2, 1, 2)(1, 1, 0)52 No Intercept	199	0.0854796	89.275002	105.86560	0.008	79.275002	0.022616		5.596895	0.215690
		Seasonal ARIMA(1, 1, 3)(1, 1, 0)52 No Intercept	199	0.0856568	89.581056	106.17166	0.008	79.581056	0.019407		5.580605	0.214817
		Seasonal ARIMA(2, 1, 2)(0, 1, 2)52 No Intercept	198	0.0842005	89.842351	109.75107	0.012	77.842351	0.017030		5.561535	0.214371
		Seasonal ARIMA(1, 1, 3)(0, 1, 2)52 No Intercept	198	0.0844253	90.197673	110.10639	0.011	78.197673	0.014258		5.546106	0.213530
		Seasonal ARIMA(2, 1, 2)(1, 1, 0)52	198	0.0854809	90.395477	110.30420	0.013	78.395477	0.012915		5.590353	0.215168
		Seasonal ARIMA(2, 1, 4)(1, 1, 0)52 No Intercept	197	0.0854712	90.753717	113.98056	0.023	76.753717	0.010797		5.545271	0.213571
		Seasonal ARIMA(1, 1, 3)(1, 1, 0)52	198	0.0856924	90.785796	110.69452	0.011	78.785796	0.010625		5.575611	0.214386
		Seasonal ARIMA(3, 1, 4)(0, 1, 2)52 No Intercept	195	0.0828478	90.829368	120.69245	0.035	72.829368	0.010397		5.459281	0.210052
		Seasonal ARIMA(2, 1, 4)(0, 1, 2)52 No Intercept	196	0.0838033	90.929210	117.47417	0.028	74.92921	0.009890		5.518595	0.212506
		Seasonal ARIMA(2, 1, 2)(2, 1, 0)52 No Intercept	198	0.0857294	91.009632	110.91835	0.010	79.009632	0.009500		5.592037	0.215445
		Seasonal ARIMA(1, 1, 4)(1, 1, 0)52 No Intercept	198	0.0858719	91.078729	110.98745	0.010	79.078729	0.009178		5.598588	0.215737
		Seasonal ARIMA(2, 1, 2)(1, 1, 1)52 No Intercept	198	0.0858143	91.107794	111.01651	0.009	79.107794	0.009045		5.590280	0.215405
		Seasonal ARIMA(2, 1, 2)(0, 1, 2)52	197	0.0843725	91.119989	114.34683	0.016	77.119989	0.008990		5.556096	0.213971
		Seasonal ARIMA(3, 1, 4)(1, 1, 0)52 No Intercept	196	0.085157	91.175293	117.72025	0.029	75.175293	0.008745		5.485315	0.211105
		Seasonal ARIMA(0, 1, 4)(0, 1, 2)52 No Intercept	198	0.0844716	91.213604	111.12232	0.006	79.213604	0.008579		5.562268	0.213352
		Seasonal ARIMA(3, 1, 2)(1, 1, 0)52 No Intercept	198	0.0859266	91.267266	111.17599	0.009	79.267266	0.008352		5.594095	0.215534
		Seasonal ARIMA(2, 1, 4)(1, 1, 0)52	196	0.0852114	91.346950	117.89191	0.030	75.34695	0.008026		5.525043	0.212380
		Seasonal ARIMA(1, 1, 3)(2, 1, 0)52 No Intercept	198	0.0859344	91.351882	111.26060	0.009	79.351882	0.008006		5.577020	0.214628
		Seasonal ARIMA(1, 1, 3)(1, 1, 1)52 No Intercept	198	0.0860076	91.437334	111.34605	0.008	79.437334	0.007671		5.574792	0.214570
		Seasonal ARIMA(1, 1, 2)(1, 1, 0)52 No Intercept	200	0.0866752	91.445537	104.71802	-0.02	83.445537	0.007640		5.594718	0.215155
		Seasonal ARIMA(1, 1, 3)(0, 1, 2)52	197	0.0846133	91.541002	114.76784	0.015	77.541002	0.007284		5.540459	0.213165

Appendix D: Bought Time Series: Best Seasonal ARIMA Values from JMP

The following JMP output includes the best seasonal ARIMA values computed for the Bought status sorted by lowest AIC followed by lowest SBC. The researcher explored all options in accordance with the Box-Jenkins Methodology in order to produce a model with significant parameters that was the least complex possible. This required at least 625 permutations for an equation with an intercept and another 625 without an intercept. Additional permutations were calculated to verify the best result.

Time Series Log[Bought]

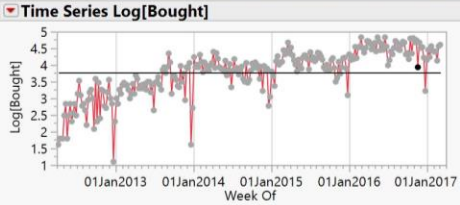

Mean	3.7756653
Std	0.6968644
N	257
Zero Mean ADF	-0.62079
Single Mean ADF	-6.00965
Trend ADF	-10.52375

Time Series Basic Diagnostics

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
		Seasonal ARIMA(4, 1, 3)(0, 1, 1)52	195	0.1231879	165.31132	195.17440	0.463	147.31132	0.018522		7.057395	0.270268
		Seasonal ARIMA(0, 1, 1)(0, 1, 1)52 No Intercept	202	0.1272945	165.77693	172.41317	0.424	161.77693	0.014675		7.300150	0.277830
		Seasonal ARIMA(0, 1, 1)(0, 1, 1)52	201	0.1265867	166.06797	176.02233	0.429	160.06797	0.012688		7.218054	0.275882
		Seasonal ARIMA(4, 1, 3)(0, 1, 1)52 No Intercept	196	0.124554	166.08612	192.63108	0.457	150.08612	0.012573		7.104412	0.270947
		Seasonal ARIMA(0, 1, 1)(1, 1, 0)52 No Intercept	202	0.1277333	166.17811	172.81435	0.423	162.17811	0.012008		7.312729	0.278493
		Seasonal ARIMA(4, 1, 3)(1, 1, 0)52 No Intercept	196	0.1246785	166.20508	192.75004	0.456	150.20508	0.011847		7.118623	0.271660
		Seasonal ARIMA(1, 1, 2)(0, 1, 1)52	199	0.1257293	166.37907	182.96967	0.438	156.37907	0.010860		7.283186	0.276950
		Seasonal ARIMA(4, 1, 1)(0, 1, 1)52	197	0.1250124	166.61118	189.83802	0.450	152.61118	0.009670		7.162760	0.273223
		Seasonal ARIMA(0, 1, 1)(1, 1, 0)52	201	0.1272361	166.63330	176.58766	0.427	160.6333	0.009564		7.231134	0.276551
		Seasonal ARIMA(2, 1, 3)(1, 1, 1)52 No Intercept	197	0.1231225	166.66087	189.88771	0.447	152.66087	0.009433		7.055339	0.268787
		Seasonal ARIMA(2, 1, 3)(1, 1, 0)52 No Intercept	198	0.1239974	166.69013	186.59885	0.442	154.69013	0.009296		7.121077	0.271332
		Seasonal ARIMA(2, 1, 3)(1, 1, 1)52	196	0.1224331	166.88995	193.43491	0.451	150.88995	0.008412		6.991753	0.267476
		Seasonal ARIMA(4, 1, 1)(1, 1, 0)52	197	0.1255328	167.07033	190.29717	0.449	153.07033	0.007686		7.158477	0.273243
		Seasonal ARIMA(1, 1, 2)(1, 1, 0)52	199	0.1266035	167.16487	183.75547	0.436	157.16487	0.007332		7.288929	0.277359
		Seasonal ARIMA(2, 1, 3)(1, 1, 0)52	197	0.1235834	167.20812	190.43496	0.446	153.20812	0.007175		7.063122	0.270239
		Seasonal ARIMA(4, 1, 4)(0, 1, 1)52	194	0.1237463	167.21317	200.39437	0.464	147.21317	0.007157		7.053165	0.270022
		Seasonal ARIMA(0, 1, 1)(2, 1, 0)52 No Intercept	201	0.1272509	167.24199	177.19635	0.425	161.24199	0.007054		7.287401	0.277230
		Seasonal ARIMA(4, 1, 3)(0, 1, 2)52	194	0.1237598	167.25163	200.43283	0.463	147.25163	0.007020		7.047761	0.269871
		Seasonal ARIMA(1, 1, 2)(0, 1, 1)52 No Intercept	200	0.1272626	167.26157	180.53405	0.431	159.26157	0.006986		7.345619	0.278172
		Seasonal ARIMA(4, 1, 1)(0, 1, 1)52 No Intercept	198	0.1265055	167.47021	187.37893	0.443	155.47021	0.006294		7.199875	0.273511
		Seasonal ARIMA(0, 1, 1)(0, 1, 2)52 No Intercept	201	0.1275682	167.49205	177.44641	0.424	161.49205	0.006225		7.297651	0.277720

Appendix E: Filled Time Series: Best Seasonal ARIMA Values from JMP

The following JMP output includes the best seasonal ARIMA values computed for the Filled status sorted by lowest AIC followed by lowest SBC. The researcher explored all options in accordance with the Box-Jenkins Methodology in order to produce a model with significant parameters that was the least complex possible. This required at least 625 permutations for an equation with an intercept and another 625 without an intercept. Additional permutations were calculated to verify the best result.

Time Series Log[Filled]

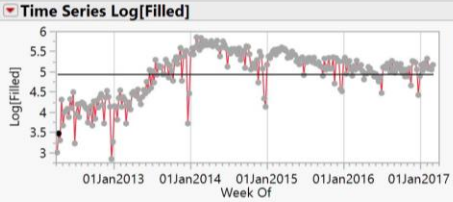

Mean	4.9384988
Std	0.5865371
N	257
Zero Mean ADF	-0.216586
Single Mean ADF	-5.477023
Trend ADF	-6.18176

Time Series Basic Diagnostics

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
		Seasonal ARIMA(0, 1, 1)(1, 0, 1)52 No Intercept	253	0.0546544	12.670775	23.306307	0.797	6.6707746	0.184212		3.920432	0.179548
		Seasonal ARIMA(0, 1, 1)(1, 1, 0)52	201	0.0586211	13.872191	23.826551	0.493	7.8721912	0.101026		3.560509	0.180183
		Seasonal ARIMA(0, 1, 1)(1, 0, 1)52	252	0.054812	14.126467	28.307177	0.798	6.1264672	0.088965		3.903508	0.178601
		Seasonal ARIMA(0, 1, 1)(1, 1, 0)52 No Intercept	202	0.0592151	14.153103	20.789343	0.489	10.153103	0.087788		3.592006	0.181625
		Seasonal ARIMA(1, 1, 1)(1, 0, 1)52 No Intercept	252	0.0549243	14.635211	28.815921	0.797	6.6352108	0.068983		3.919199	0.179524
		Seasonal ARIMA(1, 1, 1)(1, 1, 0)52	200	0.0586031	14.856382	28.128862	0.495	6.8563824	0.061762		3.570698	0.180862
		Seasonal ARIMA(0, 1, 1)(1, 1, 1)52 No Intercept	201	0.0588266	14.914678	24.869037	0.491	8.9146775	0.059987		3.611322	0.182683
		Seasonal ARIMA(1, 1, 1)(1, 1, 0)52 No Intercept	201	0.0591318	14.992039	24.946399	0.492	8.9920389	0.057711		3.595188	0.181967
		Seasonal ARIMA(0, 1, 1)(1, 1, 1)52	200	0.0584968	15.062259	28.334739	0.494	7.0622588	0.055720		3.575951	0.180987
		Seasonal ARIMA(1, 1, 1)(1, 1, 1)52 No Intercept	200	0.0587716	15.801633	29.074113	0.494	7.8016331	0.038500		3.613810	0.182964
		Seasonal ARIMA(1, 1, 1)(1, 1, 1)52	199	0.0584796	16.050486	32.641086	0.497	6.0504856	0.033996		3.581014	0.181402
		Seasonal ARIMA(1, 1, 1)(1, 0, 1)52	251	0.0550929	16.076708	33.802595	0.798	6.0767078	0.033553		3.901745	0.178517
		Seasonal ARIMA(0, 1, 1)(1, 0, 0)52 No Intercept	254	0.0563598	16.593256	23.683611	0.794	12.593256	0.025916		3.980039	0.182673
		Seasonal ARIMA(0, 1, 1)(1, 0, 0)52	253	0.0564882	17.957293	28.592826	0.795	11.957293	0.013103		3.957674	0.181397
		Seasonal ARIMA(1, 0, 1)(1, 1, 0)52 No Intercept	202	0.0594507	17.990994	27.960024	0.421	11.990994	0.012884		3.690269	0.185924
		Seasonal ARIMA(1, 1, 1)(1, 0, 0)52 No Intercept	253	0.0566564	18.398159	29.033692	0.795	12.398159	0.010511		3.977472	0.182636
		Seasonal ARIMA(1, 0, 1)(1, 1, 1)52 No Intercept	201	0.0590595	18.798831	32.090871	0.423	10.798831	0.008603		3.715058	0.187242
		Seasonal ARIMA(0, 1, 1)(0, 1, 1)52	201	0.0609195	19.013396	28.967756	0.488	13.013396	0.007727		3.566232	0.180332
		Seasonal ARIMA(0, 1, 1)(0, 1, 1)52 No Intercept	202	0.0614775	19.104467	25.740707	0.484	15.104467	0.007383		3.603856	0.182060
		Seasonal ARIMA(1, 0, 1)(1, 0, 1)52	252	0.0551339	19.276959	37.022340	0.790	9.2769593	0.006773		4.062772	0.184231
		Seasonal ARIMA(1, 0, 1)(1, 1, 0)52	201	0.0596765	19.363295	32.655334	0.459	11.363295	0.006487		3.655692	0.184364

Appendix F: Published Time Series: Best Seasonal ARIMA Values from JMP

The following JMP output includes the best seasonal ARIMA values computed for the Published status sorted by lowest AIC followed by lowest SBC. The researcher explored all options in accordance with the Box-Jenkins Methodology in order to produce a model with significant parameters that was the least complex possible. This required at least 625 permutations for an equation with an intercept and another 625 without an intercept. Additional permutations were calculated to verify the best result.

Time Series Log[Published]

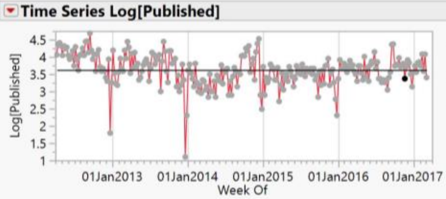

Mean	3.6215343
Std	0.4511277
N	257
Zero Mean ADF	-1.104706
Single Mean ADF	-9.524027
Trend ADF	-9.764479

Time Series Basic Diagnostics

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
	✓	Seasonal ARIMA(1, 1, 2)(1, 1, 0)52	199	0.1435461	208.09099	224.68159	0.093	198.09099	0.026484		9.246431	0.316599
		Seasonal ARIMA(1, 1, 2)(3, 1, 0)52	197	0.1330542	208.90652	232.13336	0.099	194.90652	0.017615		9.130876	0.312147
		Seasonal ARIMA(1, 1, 2)(1, 1, 2)52	197	0.1275094	209.38064	232.60748	0.101	195.38064	0.013898		9.148346	0.312746
		Seasonal ARIMA(2, 1, 1)(1, 1, 0)52	199	0.144668	209.38066	225.97126	0.087	199.38066	0.013897		9.244203	0.316767
		Seasonal ARIMA(1, 1, 2)(0, 1, 3)52	197	0.1151825	209.56592	232.79276	0.101	195.56592	0.012668		9.156361	0.313073
		Seasonal ARIMA(1, 1, 2)(1, 1, 1)52	198	0.1434427	209.64837	229.55709	0.094	197.64837	0.012156		9.244105	0.316268
✓		Seasonal ARIMA(1, 1, 1)(1, 1, 0)52 No Intercept	201	0.1488412	209.69939	219.65375	0.073	203.69939	0.011850		9.285754	0.318511
		Seasonal ARIMA(2, 1, 2)(1, 1, 0)52	198	0.1441446	209.71891	229.62763	0.096	197.71891	0.011735		9.219734	0.315748
		Seasonal ARIMA(1, 1, 2)(2, 1, 0)52	198	0.1438471	209.81466	229.72338	0.094	197.81466	0.011186		9.248561	0.316521
		Seasonal ARIMA(1, 1, 3)(1, 1, 0)52	198	0.1441542	209.85022	229.75894	0.095	197.85022	0.010989		9.230443	0.316086
		Seasonal ARIMA(2, 1, 1)(3, 1, 0)52	197	0.1330029	210.00455	233.23139	0.093	196.00455	0.010173		9.130010	0.312301
		Seasonal ARIMA(1, 1, 2)(1, 1, 0)52 No Intercept	200	0.1482357	210.12460	223.39708	0.078	202.1246	0.009580		9.218173	0.316846
✓		Seasonal ARIMA(2, 1, 1)(1, 1, 0)52 No Intercept	200	0.1486082	210.34407	223.61655	0.077	202.34407	0.008585		9.226350	0.316780
		Seasonal ARIMA(1, 1, 2)(0, 1, 1)52	199	0.1410218	210.60250	227.19310	0.078	200.6025	0.007544		9.259274	0.316468
		Seasonal ARIMA(1, 1, 2)(2, 1, 1)52	197	0.1303883	210.62257	233.84941	0.097	196.62257	0.007469		9.200971	0.314583
		Seasonal ARIMA(2, 1, 1)(1, 1, 2)52	197	0.1274447	210.65521	233.88205	0.095	196.65521	0.007348		9.149128	0.313007
		Seasonal ARIMA(1, 1, 2)(4, 1, 0)52	196	0.0014278	210.70304	237.24800	0.100	194.70304	0.007174		9.120955	0.311852
		Seasonal ARIMA(1, 1, 2)(0, 1, 2)52	198	0.1439475	210.71358	230.62230	0.089	198.71358	0.007137		9.263495	0.316855
		Seasonal ARIMA(2, 1, 1)(0, 1, 3)52	197	0.1158291	210.72269	233.94953	0.096	196.72269	0.007104		9.153552	0.313180
		Seasonal ARIMA(3, 1, 1)(1, 1, 0)52	198	0.1448093	210.75165	230.66037	0.090	198.75165	0.007002		9.278803	0.317817
✓		Seasonal ARIMA(0, 1, 3)(1, 1, 0)52 No Intercept	200	0.1489018	210.85022	224.12270	0.075	202.85022	0.006665		9.213806	0.316004

Appendix G: Cancelled Time Series: Best Seasonal ARIMA Values from JMP

The following JMP output includes the best seasonal ARIMA values computed for the Cancelled status sorted by lowest AIC followed by lowest SBC. The researcher explored all options in accordance with the Box-Jenkins Methodology in order to produce a model with significant parameters that was the least complex possible. This required at least 625 permutations for an equation with an intercept and another 625 without an intercept. Additional permutations were calculated to verify the best result.

Time Series Log[Cancelled]

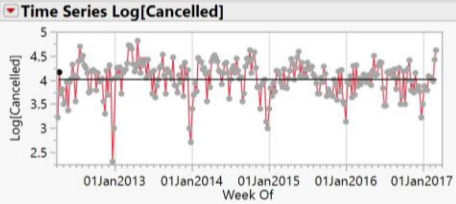

Mean	4.0208277
Std	0.3533596
N	257
Zero Mean ADF	-0.484238
Single Mean ADF	-9.67892
Trend ADF	-9.704816

Time Series Basic Diagnostics

Model Comparison

Report	Graph	Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	Weights	.2 .4 .6 .8	MAPE	MAE
		Seasonal ARIMA(2, 0, 0)(2, 0, 0)52	252	0.0819538	108.75940	126.50478	0.284	98.759397	0.074267		5.799628	0.222581
		Seasonal ARIMA(2, 0, 0)(1, 0, 2)52	251	0.0807933	109.18050	130.47496	0.288	97.1805	0.060166		5.789562	0.222144
		Seasonal ARIMA(2, 0, 0)(2, 0, 1)52	251	0.0812113	109.36294	130.65740	0.287	97.36294	0.054921		5.791022	0.222108
		Seasonal ARIMA(1, 0, 1)(2, 0, 0)52	252	0.0822501	109.37359	127.11897	0.283	99.373594	0.054629		5.807545	0.222986
		Seasonal ARIMA(1, 0, 1)(1, 0, 2)52	251	0.0808219	109.50633	130.80079	0.287	97.506334	0.051121		5.791665	0.222294
		Seasonal ARIMA(2, 0, 1)(2, 0, 0)52	251	0.0822092	109.59551	130.88997	0.288	97.595514	0.048892		5.766330	0.221058
		Seasonal ARIMA(1, 0, 1)(2, 0, 1)52	251	0.0812701	109.74210	131.03655	0.287	97.742096	0.045437		5.792694	0.222247
		Seasonal ARIMA(2, 0, 1)(1, 0, 2)52	250	0.0813123	110.40670	135.25023	0.291	96.406702	0.032590		5.773953	0.221382
		Seasonal ARIMA(2, 0, 1)(2, 0, 1)52	250	0.0817298	110.53942	135.38296	0.291	96.539425	0.030498		5.775167	0.221315
		Seasonal ARIMA(1, 0, 2)(2, 0, 0)52	251	0.082407	111.02716	132.32162	0.284	99.027163	0.023898		5.809244	0.222984
		Seasonal ARIMA(2, 0, 0)(2, 0, 2)52	250	0.0812718	111.16123	136.00476	0.288	97.161231	0.022348		5.786897	0.222017
		Seasonal ARIMA(1, 0, 0)(1, 0, 2)52	252	0.0822247	111.22217	128.96755	0.280	101.22217	0.021678		5.865789	0.225409
		Seasonal ARIMA(2, 0, 2)(2, 0, 0)52	250	0.0823391	111.39227	136.23580	0.288	97.392266	0.019910		5.762874	0.220937
		Seasonal ARIMA(1, 0, 2)(1, 0, 2)52	250	0.081179	111.41268	136.25621	0.287	97.412675	0.019708		5.794207	0.222355
		Seasonal ARIMA(1, 0, 0)(2, 0, 0)52	253	0.0838771	111.44032	125.63662	0.276	103.44032	0.019438		5.884539	0.226169
		Seasonal ARIMA(1, 0, 1)(2, 0, 2)52	250	0.0812556	111.50001	136.34354	0.287	97.500007	0.018866		5.790113	0.222221
		Seasonal ARIMA(1, 0, 0)(2, 0, 1)52	252	0.0825921	111.52407	129.26945	0.279	101.52407	0.018640		5.874133	0.225643
		Seasonal ARIMA(1, 0, 2)(2, 0, 1)52	250	0.0816043	111.60162	136.44515	0.287	97.601617	0.017931		5.796567	0.222353
		Seasonal ARIMA(2, 0, 0)(1, 0, 1)52	252	0.0812798	111.82098	129.56636	0.281	101.82098	0.016069		5.922705	0.227397
		Seasonal ARIMA(0, 0, 2)(2, 0, 0)52	252	0.0837525	112.11797	129.86335	0.279	102.11797	0.013851		5.847380	0.224440
		Seasonal ARIMA(2, 0, 2)(1, 0, 2)52	249	0.0813631	112.14440	140.53701	0.291	96.144399	0.013669		5.771347	0.221290

Appendix H: 1 Year of Air Refueling Hourly Duration Request Distribution

Distributions

Confirmed

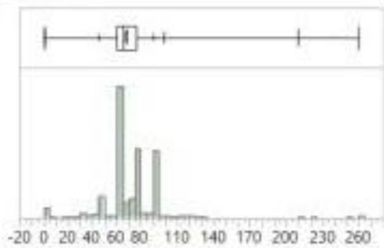

Quantiles

100.0%	maximum	260
99.5%		210
97.5%		98.925
90.0%		90
75.0%	quartile	76
50.0%	median	66
25.0%	quartile	60
10.0%		45
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	68.782372
Std Dev	23.013394
Std Err Mean	0.4120058
Upper 95% Mean	69.590202
Lower 95% Mean	67.974542
N	3120

Bought

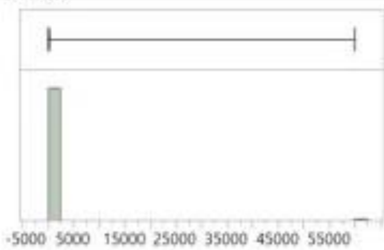

Quantiles

100.0%	maximum	60000
99.5%		186
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	50
10.0%		40
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	74.299081
Std Dev	876.38461
Std Err Mean	12.809301
Upper 95% Mean	99.411344
Lower 95% Mean	49.186819
N	4681

Filled

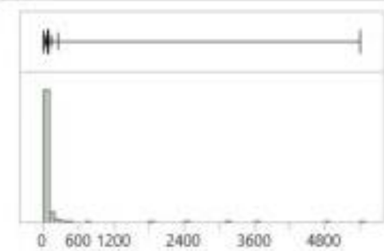

Quantiles

100.0%	maximum	5400
99.5%		240
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	54
10.0%		36
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	71.775663
Std Dev	156.63165
Std Err Mean	1.7471594
Upper 95% Mean	75.200548
Lower 95% Mean	68.350777
N	8037

Published

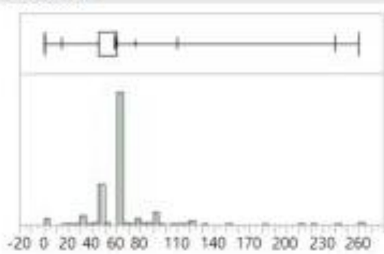

Quantiles

100.0%	maximum	260
99.5%		240
97.5%		110
90.0%		75
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		45
2.5%		15
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	59.515975
Std Dev	24.828205
Std Err Mean	0.5421833
Upper 95% Mean	60.579249
Lower 95% Mean	58.452702
N	2097

Cancelled

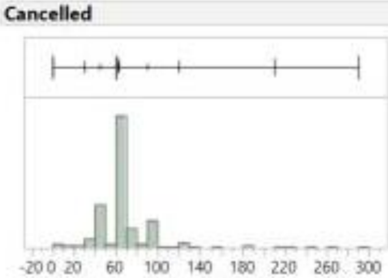

Quantiles

100.0%	maximum	290
99.5%		210
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	62.92604
Std Dev	23.482973
Std Err Mean	0.4335339
Upper 95% Mean	63.776101
Lower 95% Mean	62.075978
N	2934

Appendix I: 2 Years of Air Refueling Hourly Duration Request Distribution

Distributions

Confirmed

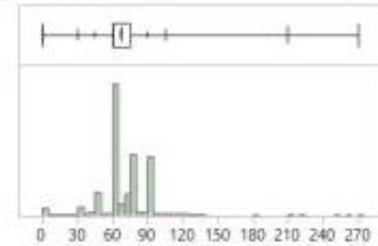

Quantiles

100.0%	maximum	270
99.5%		210
97.5%		105
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	68.085583
Std Dev	22.487502
Std Err Mean	0.2876637
Upper 95% Mean	68.649506
Lower 95% Mean	67.521661
N	6111

Bought

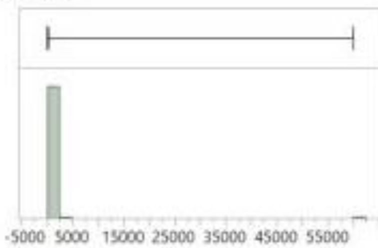

Quantiles

100.0%	maximum	60000
99.5%		169.04
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		40
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	68.028166
Std Dev	667.80276
Std Err Mean	7.4223218
Upper 95% Mean	82.577825
Lower 95% Mean	53.478506
N	8095

Filled

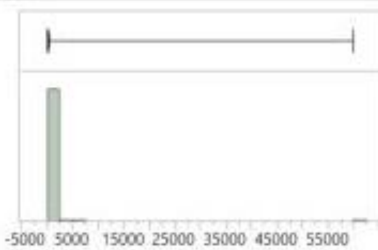

Quantiles

100.0%	maximum	60000
99.5%		300
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		40
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	79.708179
Std Dev	493.89281
Std Err Mean	3.7016752
Upper 95% Mean	86.963822
Lower 95% Mean	72.452535
N	17802

Published

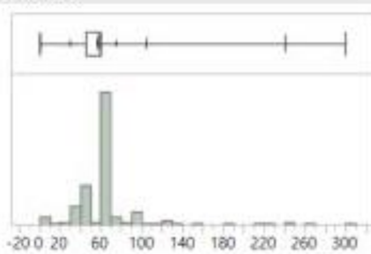

Quantiles

100.0%	maximum	300
99.5%		240
97.5%		105
90.0%		75
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	57.800468
Std Dev	24.38635
Std Err Mean	0.3930727
Upper 95% Mean	58.571118
Lower 95% Mean	57.029817
N	3849

Cancelled

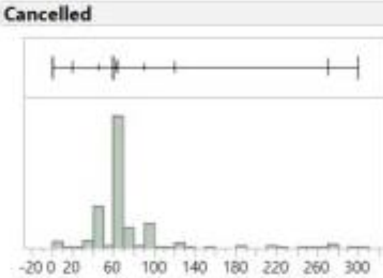

Quantiles

100.0%	maximum	300
99.5%		270
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		20
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	64.232061
Std Dev	32.067838
Std Err Mean	0.4176647
Upper 95% Mean	65.050837
Lower 95% Mean	63.413285
N	5895

Appendix J: 3 Years of Air Refueling Hourly Duration Request Distribution

Distributions

Confirmed

Quantiles

100.0%	maximum	6090
99.5%		194.25
97.5%		90.375
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	66.357126
Std Dev	68.971825
Std Err Mean	0.7479285
Upper 95% Mean	67.823248
Lower 95% Mean	64.891004
N	8504

Bought

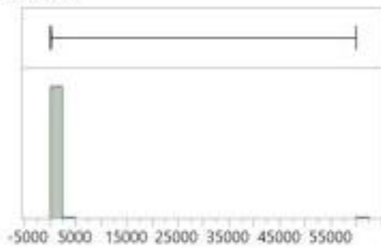

Quantiles

100.0%	maximum	60000
99.5%		150
97.5%		120
90.0%		90
75.0%	quartile	64
50.0%	median	60
25.0%	quartile	45
10.0%		35
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	65.651857
Std Dev	581.26058
Std Err Mean	5.6216229
Upper 95% Mean	76.671283
Lower 95% Mean	54.632431
N	10691

Filled

Quantiles

100.0%	maximum	60000
99.5%		240
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		40
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	74.296916
Std Dev	382.77961
Std Err Mean	2.2138938
Upper 95% Mean	78.636244
Lower 95% Mean	69.957588
N	29894

Published

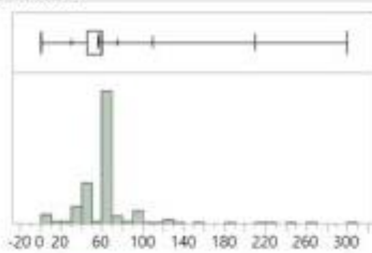

Quantiles

100.0%	maximum	300
99.5%		210
97.5%		110
90.0%		75
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	57.327986
Std Dev	24.580385
Std Err Mean	0.3245802
Upper 95% Mean	57.964286
Lower 95% Mean	56.691686
N	5735

Cancelled

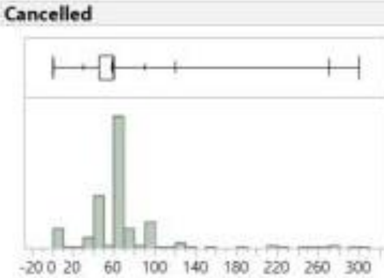

Quantiles

100.0%	maximum	300
99.5%		270
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	59.144738
Std Dev	31.550115
Std Err Mean	0.3318859
Upper 95% Mean	59.79531
Lower 95% Mean	58.494167
N	9037

Appendix K: 4 Years of Air Refueling Hourly Duration Request Distribution

Distributions

Confirmed

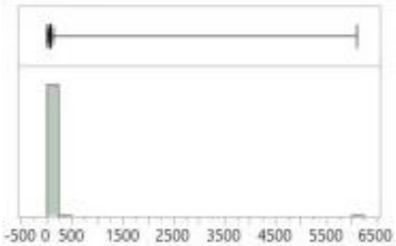

Quantiles

100.0%	maximum	6090
99.5%		137.52
97.5%		90
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	64.122963
Std Dev	61.034698
Std Err Mean	0.5759272
Upper 95% Mean	65.251881
Lower 95% Mean	62.994045
N	11231

Bought

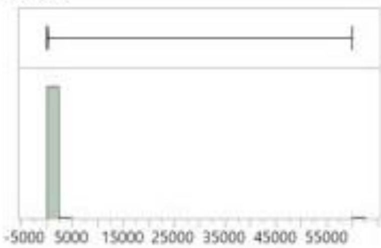

Quantiles

100.0%	maximum	60000
99.5%		150
97.5%		120
90.0%		90
75.0%	quartile	65
50.0%	median	60
25.0%	quartile	45
10.0%		35
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	64.307934
Std Dev	533.08563
Std Err Mean	4.7272072
Upper 95% Mean	73.573972
Lower 95% Mean	55.041896
N	12717

Filled

Quantiles

100.0%	maximum	60000
99.5%		240
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		40
2.5%		30
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	72.223733
Std Dev	338.06327
Std Err Mean	1.7236016
Upper 95% Mean	75.602036
Lower 95% Mean	68.845429
N	38470

Published

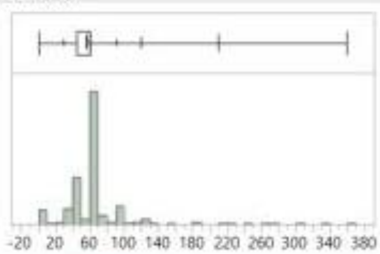

Quantiles

100.0%	maximum	360
99.5%		210
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	57.512616
Std Dev	27.201623
Std Err Mean	0.3062943
Upper 95% Mean	58.113034
Lower 95% Mean	56.912198
N	7887

Cancelled

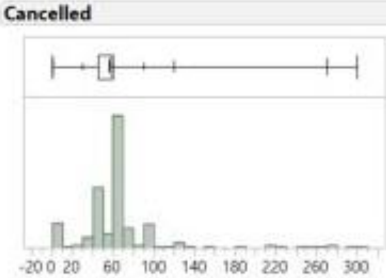

Quantiles

100.0%	maximum	300
99.5%		270
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	57.053636
Std Dev	30.649764
Std Err Mean	0.2756769
Upper 95% Mean	57.594006
Lower 95% Mean	56.513267
N	12361

Appendix L: All Test Data of Air Refueling Hourly Duration Request Distribution

Distributions

Confirmed

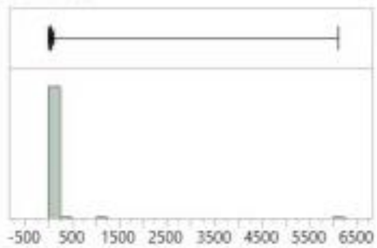

Quantiles

100.0%	maximum	6090
99.5%		125
97.5%		90
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	45
10.0%		1
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	56.052238
Std Dev	57.531116
Std Err Mean	0.4657665
Upper 95% Mean	56.965196
Lower 95% Mean	55.13928
N	15257

Bought

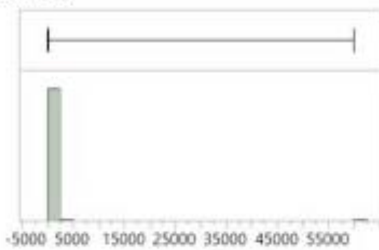

Quantiles

100.0%	maximum	60000
99.5%		150
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	63.287917
Std Dev	515.15402
Std Err Mean	4.4138407
Upper 95% Mean	71.939654
Lower 95% Mean	54.636179
N	13622

Filled

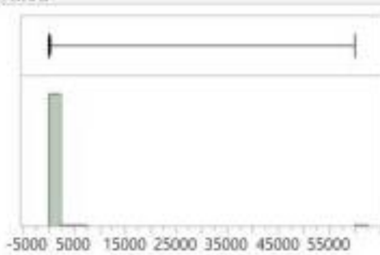

Quantiles

100.0%	maximum	60000
99.5%		240
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	50
10.0%		36
2.5%		18
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	70.959344
Std Dev	326.24662
Std Err Mean	1.604442
Upper 95% Mean	74.104085
Lower 95% Mean	67.814603
N	41347

Published

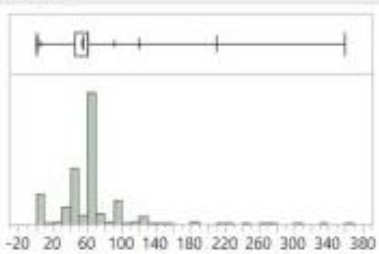

Quantiles

100.0%	maximum	360
99.5%		210
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		5.6
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	55.046262
Std Dev	30.098515
Std Err Mean	0.2933545
Upper 95% Mean	55.621292
Lower 95% Mean	54.471232
N	10527

Cancelled

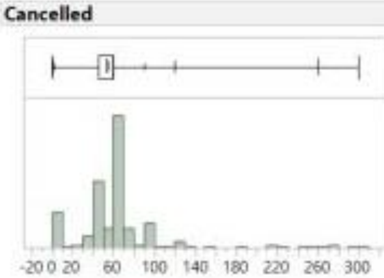

Quantiles

100.0%	maximum	300
99.5%		260
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	45
10.0%		1
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	54.805979
Std Dev	30.939228
Std Err Mean	0.2513392
Upper 95% Mean	55.298634
Lower 95% Mean	54.313324
N	15153

Appendix M: Validation Data of Air Refueling Hourly Duration Distribution

Distributions

Confirmed

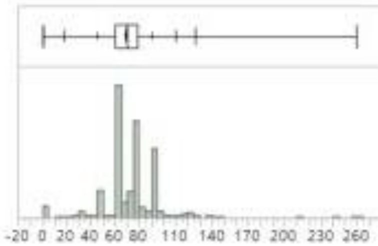

Quantiles

100.0%	maximum	260
99.5%		126
97.5%		110
90.0%		90
75.0%	quartile	78
50.0%	median	70
25.0%	quartile	60
10.0%		45
2.5%		17.375
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	69.456835
Std Dev	20.814105
Std Err Mean	0.3763908
Upper 95% Mean	70.194839
Lower 95% Mean	68.71883
N	3058

Bought

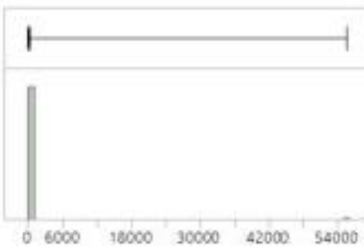

Quantiles

100.0%	maximum	55555
99.5%		180
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	48
10.0%		30
2.5%		15
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	74.608449
Std Dev	891.01066
Std Err Mean	14.300629
Upper 95% Mean	102.64591
Lower 95% Mean	46.570988
N	3882

Filled

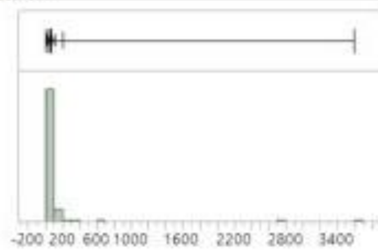

Quantiles

100.0%	maximum	3600
99.5%		210
97.5%		120
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		36
2.5%		24
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	65.989177
Std Dev	63.524771
Std Err Mean	0.700532
Upper 95% Mean	67.362396
Lower 95% Mean	64.615957
N	8223

Published

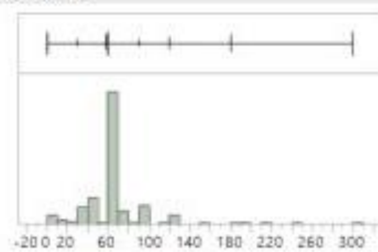

Quantiles

100.0%	maximum	300
99.5%		180
97.5%		120
90.0%		90
75.0%	quartile	60
50.0%	median	60
25.0%	quartile	60
10.0%		30
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	59.57746
Std Dev	25.771553
Std Err Mean	0.4591825
Upper 95% Mean	60.477788
Lower 95% Mean	58.677133
N	3150

Cancelled

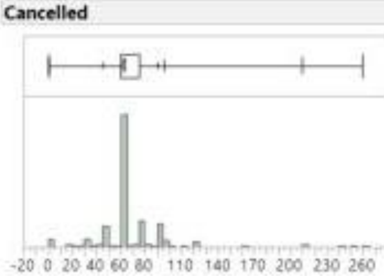

Quantiles

100.0%	maximum	260
99.5%		210
97.5%		96
90.0%		90
75.0%	quartile	75
50.0%	median	60
25.0%	quartile	60
10.0%		45
2.5%		0
0.5%		0
0.0%	minimum	0

Summary Statistics

Mean	63.411838
Std Dev	22.178904
Std Err Mean	0.3904277
Upper 95% Mean	64.177349
Lower 95% Mean	62.646326
N	3227

Bibliography

AFCAA/FMCY. (2018). *Many Faces of Operations and Support*.

Boeing. (2018). KC-135 Stratotanker. Retrieved December 19, 2018, from http://www.boeing.com/history/products/kc-135-stratotanker.page

Box, G. E. P., & Jenkins, G. M. (1970). *Time Series Analysis: Forecasting and Control* (1st ed.). San Fransico: Holden-Day.

Commandeur, J. J. F., & Koopman, S. J. (2007). *An Introduction to State Space Time Series Analysis*. (J. Doornik & B. Hall, Eds.). New York: Oxford University Press.

Defense Science Board. (2004). *Defense Science Board Task Force Report on Aerial Refueling Requirements*.

Gertler, J. (2013). *Air Force KC-46A Tanker Aircraft Program*. Retrieved from https://www.everycrsreport.com/files/20130514_RL34398_d25149b537c171bba7f9c2714a468344aba246b7.pdf

Grismer, M. W. (2011). Fiscally sound options for a flawed tanker recapitalization strategy. *Air and Space Power Journal*, *26*(3), 63–73.

HQ AMC/A3O. (2016). *AIR FORCE INSTRUCTION 11-221: AIR REFUELING MANAGEMENT*. Retrieved from https://static.e-publishing.af.mil/production/1/af_a3/publication/afi11-221/afi11-221.pdf

HQ AMC Office of History. (2018). *Command Aircraft 1990-2015*. Retrieved from http://www.amc.af.mil/Units/AMC-History

Hyndman, R. J. (2000). *Box-Jenkins modelling*. Retrieved from https://pdfs.semanticscholar.org/5b99/a115799e519176b55fae2bde44c778c87e21.pdf?_ga=2.92465237.958563023.1555083215-97551347.1554668301

Hyndman, R. J. (2001). ARIMA processes. *The Informed Student Guide to Management Science*, (May), 2. https://doi.org/10.2307/1391375

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. OTEXTS. Retrieved from https://otexts.org/fpp2/index.html

Insinna, V. (2019). Boeing delivers first KC-46, but fixes to technical problems still years away. Retrieved March 20, 2019, from https://www.defensenews.com/breaking-news/2019/01/10/boeing-delivers-first-kc-46-but-fixes-to-technical-problems-still-years-away/

JMP Support. (2018). Modeling Reports | JMP 12. Retrieved February 10, 2019, from http://www.jmp.com/support/help/Modeling_Reports.shtml

Jones, A., & Arnold, T. (2019). Equations for the Smoothing Models. Retrieved April 15, 2019, from http://support.sas.com/documentation/cdl/en/etsug/63348/HTML/default/viewer.htm#etsug_tffordet_sect014.htm

Kennedy, M., Baldwin, L. H., Boito, M., Calef, K. M., Chow, J. S., Cornuet, J., … Tonkinson, J. (2006). *Analysis of Alternatives (AoA) for KC-135 Recapitalization. Executive Summary*. https://doi.org/RAND

Klein, J. L. (1997). *Statistical Visions in Time: A History of Time Series Analysis, 1662-1938*. Cambridge, United Kingdom: Cambridge University Press.

Knight, W., & Bolkcom, C. (2008). *Air Force Air Refueling: The KC-X Aircraft Acquisition Program*. Retrieved from https://www.everycrsreport.com/files/20080304_RL34398_ae83c68d866d75211f6084cf7b040ea934fbb635.pdf

Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and Applications* (Third Edit). Hoboken, NJ: Wiley and Sons, Inc.

McAndrew, A. J. (2018). *Fiscal Year (FY) 2019 Department of Defense (DoD) Fixed Wing and Helicopter Reimbursement Rates*.

Meredith, W. C., Stone, T. F., Dey, M. J., Newell, K. W., & Ragsdale, N. L. (1996). *U.S. COMBAT AIR POWER: Aging Refueling Aircraft Are Costly to Maintain and Operate*. Retrieved from https://www.gao.gov/assets/160/155520.pdf

Myung, I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology*, *47*, 90–100. https://doi.org/10.1016/S0022-2496(02)00028-7

Nau, R. (2014a). Forecasting with moving averages. Retrieved January 19, 2019, from http://people.duke.edu/~rnau/Notes_on_forecasting_with_moving_averages--Robert_Nau.pdf

Nau, R. (2014b). Lecture notes on forecasting Introduction to ARIMA models. Retrieved January 4, 2019, from http://people.duke.edu/~rnau/forecasting.htm

Nau, R. (2018). Stationarity and differencing of time series data. Retrieved January 7, 2019, from http://people.duke.edu/~rnau/411diff.htm

Ngo, T. H. D. (2013). The Box-Jenkins Methodology for Time Series Models. *SAS Global Forum 2013*, *6*, 1–11. Retrieved from http://support.sas.com/resources/papers/proceedings13/454-2013.pdf

Omega Air. (2018). Omega Air Refueling. Retrieved December 21, 2018, from http://www.omegaairrefueling.com/

Pawlyk, O. (2018). Will Ongoing KC-46 Tanker Problems Affect KC-10 Retirement? Retrieved March 30, 2019, from https://www.military.com/defensetech/2018/03/23/will-ongoing-kc-46-tanker-problems-affect-kc-10-retirement.html

Pitfield, D. E. (1993). Predicting Air-Transport Demand. *Environment and Planning A: Economy and Space*, *25*(4), 459–466. https://doi.org/10.1068/a250459

Singh, G. (2018). 7 methods to perform Time Series forecasting. Retrieved January 7, 2019, from https://www.analyticsvidhya.com/blog/2018/02/time-series-forecasting-methods/

U.S. Air Force. (2014a). F-35A Lightning II. Retrieved December 21, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/478441/f-35a-lightning-ii-conventional-takeoff-and-landing-variant/

U.S. Air Force. (2014b). KC-10 Extender. Retrieved December 19, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104520/kc-10-extender/

U.S. Air Force. (2015a). F-16 Fighting Falcon. Retrieved December 21, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104505/f-16-fighting-falcon/

U.S. Air Force. (2015b). F-22 Raptor. Retrieved December 21, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104506/f-22-raptor/

U.S. Air Force. (2016). KC-46A Pegasus. Retrieved December 19, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104537/kc-46a-pegasus/

U.S. Air Force. (2018). KC-135 Stratotanker. Retrieved December 19, 2018, from https://www.af.mil/About-Us/Fact-Sheets/Display/Article/1529736/kc-135-stratotanker/

USTRANSCOM TCJ5. (1999). *Concept Development Report on Contracted Aerial Refueling*.

Wensveen, J. G. (2015). *Air Transportation: A Management Perspective* (Eighth). Burlington, VT: Ashgate Publishing Company.

Wickham, R. R. (1995). *Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation*. Massachusetts Institute of Technology. Retrieved from https://dspace.mit.edu/handle/1721.1/68101

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, *50*, 159–175. Retrieved from www.elsevier.com/locate/neucom

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From – To)*
14-06-2019	Graduate Research Paper	May 2018 – June 2019

4. TITLE AND SUBTITLE
Time Series Forecasting of Tanker Training Demand

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Hawkins, Steven R., Major, USAF

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S)
Air Force Institute of Technology
Graduate School of Engineering and Management(AFIT/ENS)
2950 Hobson Way
WPAFB, OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT-ENS-MS-19-J-033

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Major General Ricky Rupp, Director of Operations, J-3,
United States Transportation Command
Scott AFB, IL 62225-5357
312-770-7000
Ricky.Rupp@us.af.mil

10. SPONSOR/MONITOR'S ACRONYM(S)
USTC/J-3

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution Statement A. Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Since 1998 the USAF tanker fleet has decreased by 155 aircraft, with the additional anticipated retirement of the equivalent of 201 KC-135 tankers by 2029. Meanwhile, the new KC-46A is scheduled to replace only 179 boom-equipped tankers by 2029, resulting in a total decrease of 131 aircraft which corresponds to a reduction of 24.5% in total air refueling fuel capacity in just over three decades. This research examines historical tanker training requests, drawn from the 618th AOC's Air Refueling Scheduling Tool (ARST), and uses multiple forecasting techniques, including autoregressive integrated moving average (ARIMA) models, in order to create a model for predicting future air refueling training demand and communicating that demand in terms of aircraft flight hours. Air refueling remains a supply and demand problem in which there will always be more demand than the ability to supply with any realistically sized tanker fleet. The ability to understand, predict, and prepare for increased air refueling demand holds real value to planners, tanker units, and receiver units. This research is a first step in more clearly understanding unsupported air refueling training demand in terms of tanker aircraft flight hours.

15. SUBJECT TERMS
Time Series Forecasting, Air Refueling, Tankers, Demand, ARIMA

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Cunningham, William A. III, Ph.D., AFIT/ENS
					19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	UU	135	(937) 255-6565, x 4283 (William.Cunningham@afit.edu)