Time-Series Statistical Model for Forecasting Revenue and Risk Management

Mihir Kelkar[1], Cosmin Borsa[2], and Lina Kim[2#]

[1]Lynbrook High School, San Jose, CA, USA
[2]University of California, Santa Barbara, CA, USA
[#]Advisor

ABSTRACT

Following a Low-Cost Carrier (LCC) model, Southwest Airlines has consistently demonstrated growing annual revenues up until the start of the COVID-19 pandemic. Southwest's quarterly revenue shows that there exists a strong seasonal component with the revenue in the first quarter of the fiscal year (September) significantly higher than other quarters. Using the quarterly revenue data we constructed a time-series model: a Seasonal Autoregressive Integrated Moving Average (Seasonal ARIMA or SARIMA) model to forecast Southwest's revenue over 2020. We then performed a cost and solvency risk analysis using the company's financial results from its annual reports to analyze Southwest's financial performance due to COVID-19, and proposed business strategies to keep Southwest financially stable.

Introduction

An income statement reflects the effect of management's operating decisions on business performance, and accounts for the resulting profit or loss for a company [1, p. 47].

Analyzing financial statements is a complex activity that uses a system of models, methods, and techniques to assess financial position, performances, and potential of an enterprise. In this paper, we employ time-series analysis techniques for forecasting revenue to assess the financial situation of Southwest.

Time series datasets are datasets in which the order of observation is essential [2]. Time series analyses are about examining patterns in data that emerge over time, of a length pertinent to the problem at hand. Time series forecasting is performed in nearly every organization that works with quantifiable data. Retail and transportation sales time series tend to portray seasonal trends in that there is a repeating pattern over a set interval of time. Time series data which have a trend and exhibit cyclicity can be trained into a Seasonal Autoregressive Integrated Moving Average (Seasonal ARIMA or SARIMA) model [3, p. 157].

The time series investigated in this paper refers to the quarterly revenue data published by Southwest Airlines (trading ticker: LUV) over the course of 14 years from 2005 to 2019 [4] (FIG. 1, revenue in millions USD). Prior to 2005, Southwest was dealing with the economic losses of the September 11th, 2001, terrorist attack on the World Trade Center. This reduced flight travel significantly, and airline revenues were lower than their predicted outcome following the fiscal year of 2000. In order to appropriately forecast with the SARIMA model, we acquired data following 2005, when the company had recovered from the 9/11 tragedy and started displaying seasonal growth in revenue. We looked to analyze the calm period following September 11 and up to the start of the COVID-19 pandemic.

Critical to selecting data for the SARIMA model is consistency in the company's performance: Southwest has seen constant growth over the past decade. We collected the most recent quarterly revenue reports released by Southwest from 2005 to 2019, to forecast revenue in 2020, had COVID-19 not happened. Points from 2005-2018 were used to build the model, and points from 2019 were used to test the accuracy of the model.
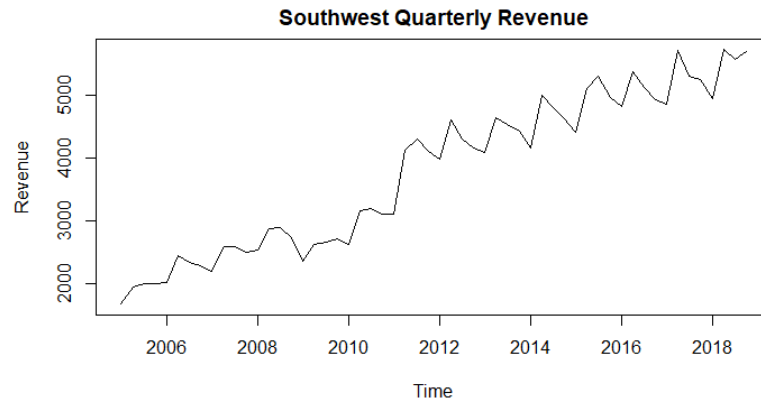
Figure 1. Graph of Southwest's quarterly revenue

The goal of this project is to build and test the SARIMA model on Southwest. Combining the model's predictions with Southwest's balance sheets, we examine potential risks and simulate possible real-world decisions Southwest could take. With the programming language R [5], we analyze and interpret the data into a SARIMA model and explain the impact of the pandemic.

Revenue Data Analysis

Southwest Airlines' sales have been highly seasonal with upside and downside swings, since revenue is dependent on factors such as leisure travel, which usually occurs during seasons of travel. However, the revenue has also followed a long-term positive linear trend due to an overall increase in commerce and population. Gradually, increasing volumes of people have been traveling.

To see these features of the data, we plot the sample autocorrelation function (ACF) as well as the histogram (FIG. 2) of the Revenue. In simple terms, ACF shows relation between present values and past values over different lag times. As variance is for single variable statistics, ACF is for time series. The sample autocorrelation function allows us to assess the degree of dependence (correlation) in the data at various times [3, p. 16]. A series is addressed as having low correlation if its plot demonstrates randomness, and the sample ACF is within the confidence interval (Figure 2, marked with blue dashed lines) for each lag except for lag 0, which should have an ACF of 1.
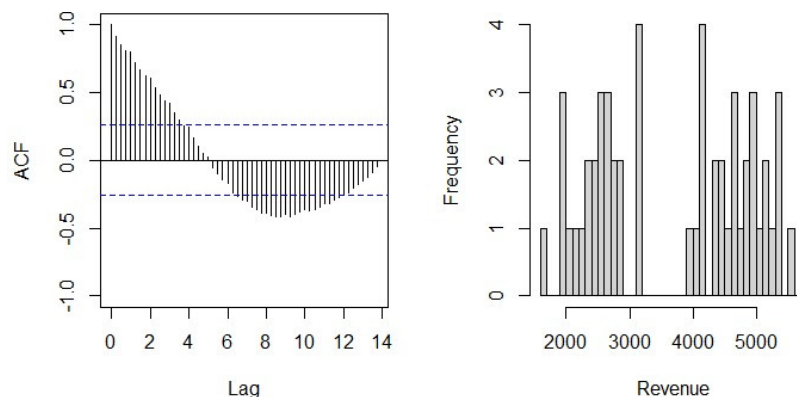

Figure 2. ACF and Histogram of Revenue Data

Here, the ACF for revenue data shows a strong linear trend in the spikes, and the histogram plot does not assume a normal distribution. Consequently, we have to difference the data to make it stationary. To determine

necessary transformations, we decompose the time series to observe its individual components (FIG 3): a trend component, a seasonal component, and a random noise component that is stationary [3, p. 20].
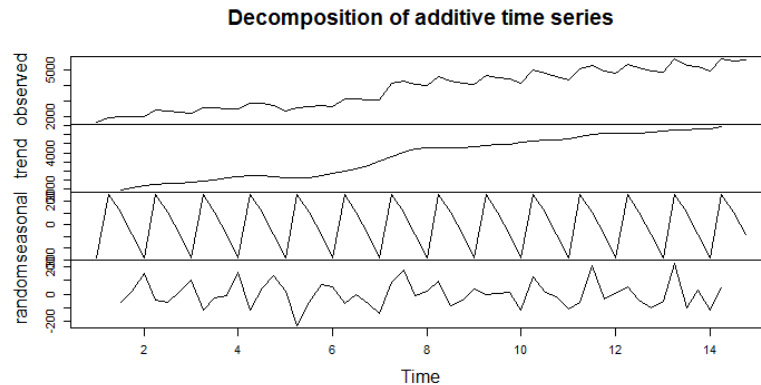

Figure 3. Decomposition Plot of Revenue Data

The second and third rows of the decomposition graph show the positive linear trend and the seasonal up-downs respectively. The variance of the random component seems to resemble that of Gaussian white noise, however, to see if it can be stabilized further, we will perform a Box-Cox transformation [6]. There are also large spikes in the random component, which indicate the white noise might follow a heavy-tailed distribution.

Data Transformation

To build a model for the data, we must transform the data to acquire a stationary time series. The resulting plot should resemble Gaussian White Noise. We will test if a Box-Cox transformation stabilizes the series. Then we difference the dataset at specific lags for seasonality and the linear trend to obtain a stationary time series.

Testing a Box-Cox Transformation

In order to properly assess the output of the ACFs, the data in Figure 1 needs to be sized down such that it resembles a plot of Gaussian White Noise: a plot of seemingly random spikes. Box-Cox transformations are useful in removing the increasing or decreasing variability of the time series with each level i.e., it will remove the trend where the difference between the dip and peak is increasing, such that the difference remains constant [3, p. 390]. We tentatively perform a Box-Cox transformation (FIG. 4), shown in equation 1, to smoothen the variance.

Equation 1: box-cox transformation of the time series from Figure 1:

$$Y_t = (X_t{}^{\lambda} - 1)/\lambda$$

Where,
$Y_t = Box - Cox\ transformed\ series$
$X_t = Original\ time\ series$
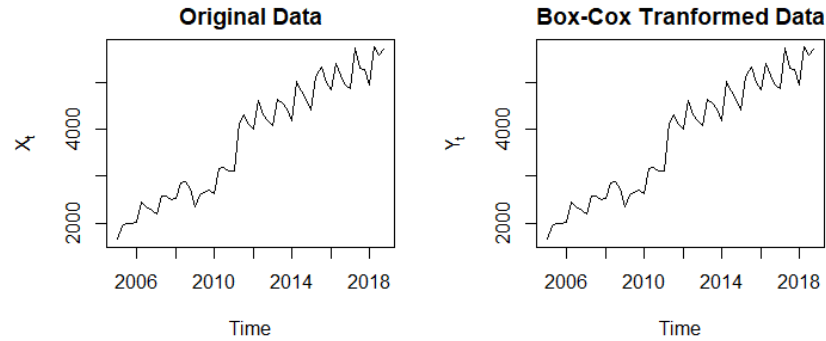$\lambda = Box - Cox\ Transformation\ parameter$

Figure 4. Graph of Box-Cox transformed revenue data compared to original revenue data

Comparing the two plots, we notice close to no change in the plots, and opposite to our intention, the variance of the Box-Cox transformed data (Y_t) increased. Therefore, we proceeded with the original data (X_t)for the next set of transformations.

Differencing

From Figure 1, we see two clear trends in the data: an upward, linear business trend, and a seasonal trend of periodic peaks and dips. We can remove these trends through differencing, where the seasonal and linear components are removed from the data by repeated differencing at specific lags to create a stationary white noise time series [3, p. 25].

To remove the positive linear trend, we difference at lag-4 in Equation 2. (FIG. 5). This is because the seasonality is based on quarterly reported revenue for the fiscal years.

Equation 2: Differencing for seasonality to arrive at differenced time series

$$Y_t = X_t - X_{t-4},$$

Where,
$Y_t \ is \ lag - 4 \ differenced \ time \ series$
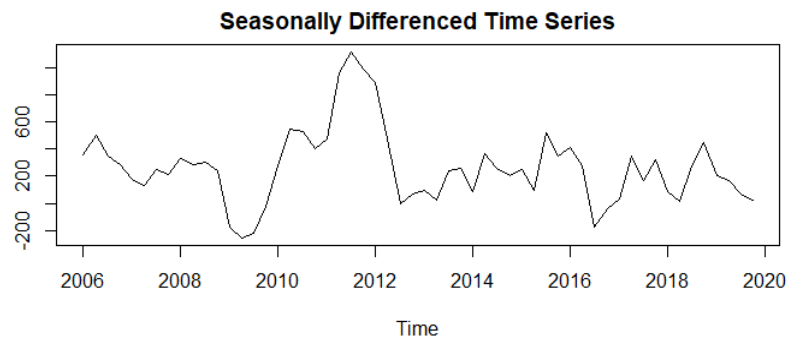$X_t \ is \ original \ time \ series$

Figure 5. Seasonally Differenced Time Series

This equation can be rearranged (equation 3) to incorporate a backward shift operator (B) which is a variable that equates a time series to an instance immediately preceding it [3, p. 25].

Equation 3: Introducing the backward shift operator (B) into Equation 2,

$$Y_t = (1 - B^4)X_t$$

Where,

$B^4 = backward\ shift\ operator\ for\ lag-4$

The resulting time series in Figure 5 is still not entirely random, as at certain points throughout the series, it exhibits some linearity. Using Equation 4, we remove the spontaneous linearity by differencing at lag-1 on Y_t from Equation 3 for quarterly revenue.

Equation 4: we difference the resulting equation from equation 3 for its linear trend

$$Y_t = (1 - B^4)(1 - B)X_t$$

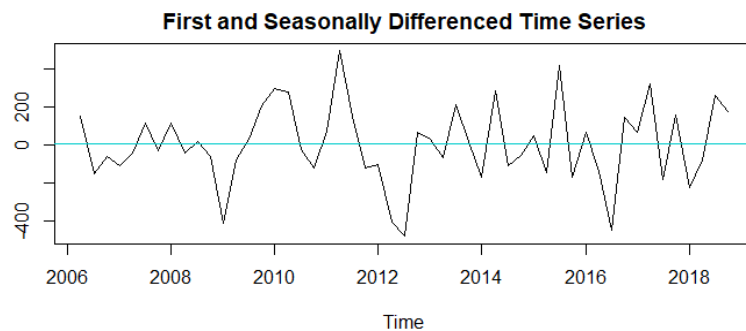Where,

$B = backward\ shift\ operator\ for\ lag-1$

Figure 6. First and Seasonally Differenced Time Series

Through a sequence of differencing for seasonality and linear trend, the time series looks stationary, and there is no visible trend (FIG. 6). The ACF shows no specific trend, and while the histogram still demonstrates heavy tails, it looks considerably close to a normal distribution than from Figure 1 (FIG. 7).
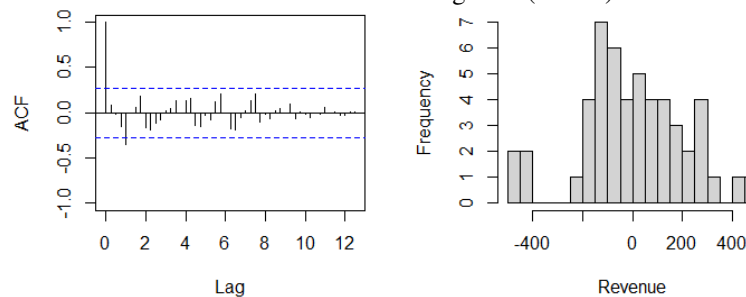

Figure 7. ACF and Histogram of First and Seasonally Differenced Data

Building The Model

The Seasonal ARIMA model (SARIMA), is a non-stationary seasonal model derived from the class of ARMA models. ARMA models are the combination of Autoregressive and Moving Average models used to describe a particular time series as a function of a Gaussian white noise [7, p. 77]. The "I" letter of the ARIMA model acronym corresponds to the initial differencing steps applied to eliminate the non-stationarity part of the time series. Our observed time series required differencing of order 1 in order to remove the linear trend component. The seasonal "S" component of the seasonal ARIMA model comes from differencing the time series at lag 4 to remove the cyclicity of

the data. [3, p. 177] To summarize, each component of the Seasonal ARIMA model manipulates Gaussian white noises to produce the time series being observed.

The SARIMA model is represented by SARIMA (p, d, q) x (P, D, Q)$_s$ [3, p. 177]. Parameters p and q represent the orders of the non-seasonal autoregressive (AR) and moving average (MA) terms respectively. P and Q represent the orders of the seasonal AR and MA terms, respectively. The parameter s represents the seasonal frequency of the data. The d and D correspond to the orders of differencing associated with the "I" component of non-stationarity component. Since quarterly revenue data is used in this case study, the s value is set at 4. The d and D parameters represent the times we difference for trend and seasonality, respectively. In this case, they are both 1.

While the d, D, and s parameters have already been determined, due to the randomness exhibited in the decomposition graph (FIG. 3), the p, P, q, and Q parameters will need to be chosen according to the ACF, and a slightly modified version of the ACF, the partial ACF (PACF). The ACF of a time series is used to examine how it is related to the Moving Average model. The PACF operates similarly, however it examines how a time series is related to the Autoregressive model. To make sure the right parameters are chosen, we suggest potential candidates based on our analysis of the ACF and PACF graphs (FIG. 7).

To select an appropriate SARIMA model for the revenue, we take a look at the ACF and also the partial autocorrelation function (PACF) of the seasonally differenced time series, and then at the first and seasonally differenced time series. Specifically, we look at spikes in both plots that fall outside the confidence interval indicated by the dashed blue lines, except for ACF and PACF spikes at lag 0, which are always present.

There are quite a few lags in both the ACF and PACF plots which meander close to the confidence interval, either just above it or just below. The lag values correspond to that many periods, where one period is a quarter in the fiscal year. Hence the spikes at lag 2 corresponds to quarter 8 in the quarterly revenue data. When determining parameters, we consider whether there is a spike at lag 1 for p and q, and then whether there are spikes occurring at lags greater than 1 to determine P and Q. Depending on how close a spike is to its confidence interval, we will need to check whether the corresponding parameter is truly its designated value, or whether the spike being outside the confidence interval is simply due to random noise.
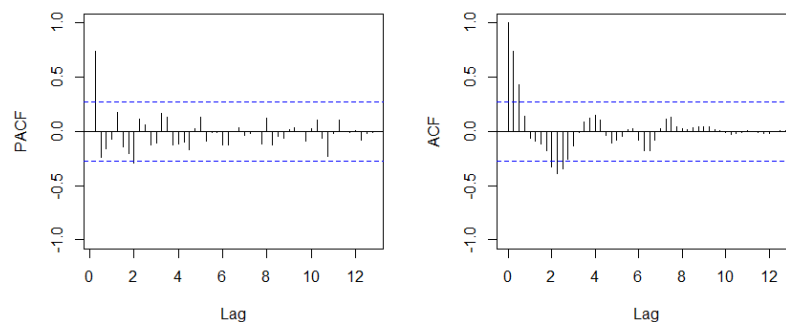

Figure 7. PACF and ACF of Seasonally Differenced Data

Looking at the ACF and PACF from Figure 7, in both plots, there is no spike outside the confidence interval at lag 1, so p and q are likely 0 and there is a spike at lag 2 in the PACF and ACF, so it is reasonable to suggest that $P, Q \geq 1$. However, since the PACF lag 1 spike approaches the confidence interval, we will test with various values of p, q.

Having rough estimates of SARIMA model parameters, we formulated four possible model candidates.

Equation 5. SARIMA model candidates

$$SARIMA(1,0,1) \times (1,1,1)_4$$
$$SARIMA(1,1,1) \times (0,1,2)_4$$
$$SARIMA(1,0,0) \times (0,1,1)_4$$

$$SARIMA(0,1,0) \times (0,1,1)_4$$

We evaluate the model candidates according to their Akaike's Information Criterion, corrected for bias (AICc) [8]. The AIC is a calculation on the residuals of the model to estimate its accuracy to the given data. It measures the goodness of fit by balancing the error of the fit against the amount of model parameters. When comparing model candidates, we should pick the model with the smallest AICc value (TAB. 1) [7, p. 51].

Table 1. Akaike's Information Criterion for the four model candidates

Model Candidate	AICc
$SARIMA(1,0,1) \times (1,1,1)_4$	750.3461
$SARIMA(1,1,1) \times (0,1,2)_4$	736.9157
$SARIMA(1,0,0) \times (0,1,1)_4$	746.0329
$SARIMA(0,1,0) \times (0,1,1)_4$	730.7539

From the AICc values, model (4) has the smallest AICc and is therefore the most appropriate seasonal ARIMA model for the revenue data. Model (4) in equation 5 shows the representation of the SARIMA model with its parameters. The mathematical function can be seen in equation 6, which expresses our observed time (X_t) series as a function of white noise (Z_t).

Equation 6. Time series equation for model (4).

$$(1 - B)(1 - B^4)X_t = (1 - 0.597B^4)Z_t$$

Unit Roots are used to tell us whether our model needs to be differenced further, or it has been over-differenced. A root near 1 for the autoregressive indicates that the data should be differenced before modeling. A root near 1 for the moving average suggests the data may have been over-differenced [3, p. 169]. The polynomial $(1 - 0.597B^4)$ has its roots outside the unit circle (FIG. 8) indicating model invertibility, and no further transformations need to be executed.
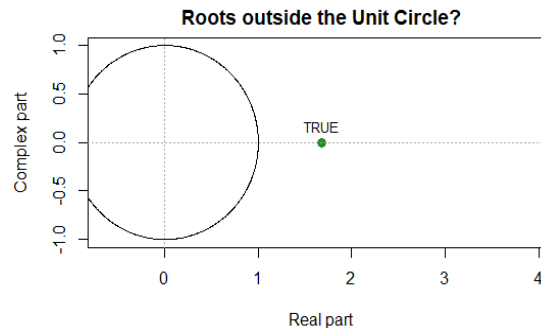

Figure 8. Green roots are outside the unit circle

Diagnostic Checking

The goodness of fit of statistical model (4) from equation 5 can be tested based on the behavior of its residuals. Since our chosen SARIMA model was a pure MA model, and it is causal, we may proceed with diagnostic checking [7, p. 140]. We start by examining whether the residuals resemble a Gaussian white noise.
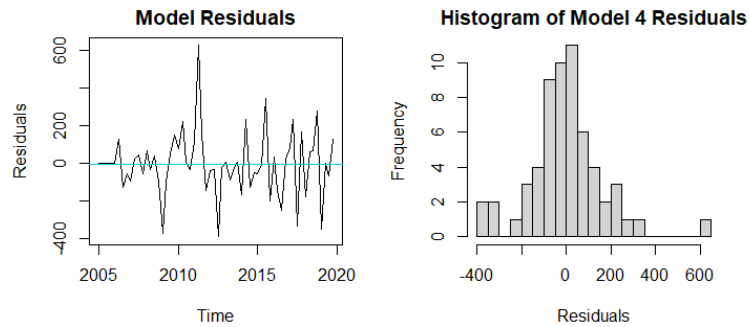

Figure 9. Graph and histogram of the residuals of $SARIMA(0,1,0)\times(0,1,1)_4$

If the fitted model is appropriate, the graph of the residuals should resemble that of a white noise [3, p. 145]. The graph gives no indication of a non-zero mean or non-constant variance. So, it is reasonable to assume that the residuals resemble a Gaussian white noise. The histogram also indicates a mean at 0, however there seem to be heavy tails on either ends (FIG. 9).

We reach the same conclusion from plotting a Quantile-Quantile (Q-Q) plot. Q-Q plots are used to compare shapes of distributions and are commonly used to compare a data set to a theoretical model [9]. We compare the residuals to a sample quantities model seen below (FIG. 10).
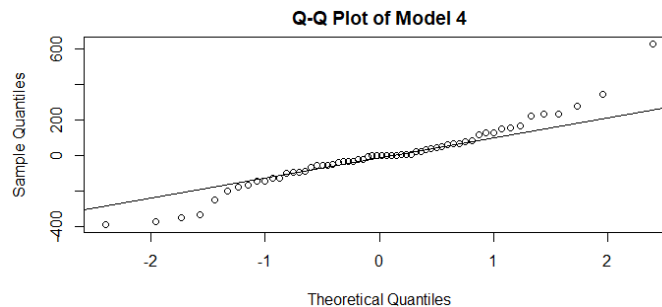

Figure 10. Normal Q-Q Plot of model Residuals

We see that at deviations of 1 on either side of 0, the residuals align with the plotted line, while veering off near the edges.

We can test whether or not the observed residuals are consistent with white noise by examining their sample autocorrelation functions [3, p. 146] Therefore we pass the residuals into the ACF and PACF. The spikes all fall well within the confidence interval indicating that the residuals resemble Gaussian White noise (FIG. 11). The selected model seems to appropriately fit the selected data as excluding the ACF lag 0, none of the lags have spikes outside the confidence interval (dotted lines).
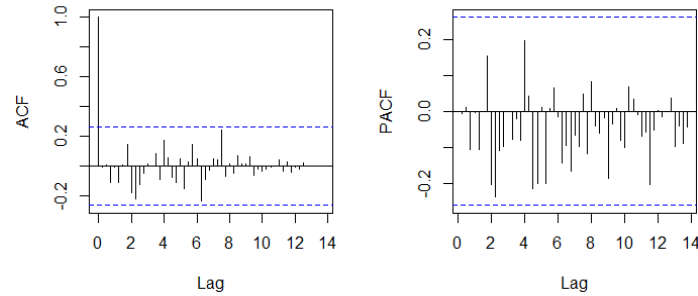
Figure 11. ACF and PACF of Model (4) Residuals

Next, we perform Portmanteau tests [10] to assess the normality and independence of the residuals and whether they resemble white noise (TAB 2). The portmanteau tests are chi squared statistics that measure the independence of the residuals to the model. A p-value greater than 0.05 indicates that we fail to reject the null hypothesis that the residuals are Gaussian white noise. This in tune shows that the model is appropriate for forecasting.

Table 2. Table of the portmanteau statistical tests and their p-values

Statistical Test	p-value
Shapiro-Wilk Test	0.008
Box-Pierce Test	0.730
Ljung-Box Test	0.598
McLeod-Li Test	0.783

Shapiro-Wilk Test: The Shapiro-Wilk test checks for the normality of the residuals (EQ. 7) where the $x_{(i)}$ are the ith order statistic, and the a_i's are constants generated from means, variances, and covariances of the order statistics of a sample size n [11].

Equation 7. Equation that calculates the test-statistic of the Shapiro-Wilk Test

$$W = \frac{\left(\sum_{i=1}^n a_i x_{(i)}\right)^2}{\sum_{i=1}^n (x_i - \bar{x})^2}$$

The computed p-value of 0.008 (Table 2) is less than 0.05. Therefore, we reject the null hypothesis that the distribution is normal. However, this is no surprise as the residuals follow a heavy-tailed distribution, not Gaussian.

Box-Pierce Test: The Box-Pierce test assesses the independence of residuals (EQ. 8) where $h = \sqrt{n}$.

Equation 8. Equation to calculate the test statistic of the Box-Pierce Test

$$Q_W = n \sum_{j=1}^h \hat{p}^2 w(j) \sim x_{1-\alpha}^2 (h - p - q)$$

Since we have 60 observations, we may approximate h by rounding down: $h = 7$. For the selected model we have p and q both equal to 0. The test statistic is going to be a χ^2 (chi-squared) distribution with 7 degrees of freedom [12]. As seen by a p-value of 0.730 (Table 2) which is greater than 0.05, we fail to reject the null hypothesis about the data being independent and conclude that the residuals resemble a white noise.

Ljung-Box Test: The Ljung-Box test also checks for independence of residuals. It considers the magnitude of the autocorrelation of each h grouping in the data [10]. The test statistic is going to be a χ^2 (chi-squared) distribution with 7 degrees of freedom (EQ. 9).

Equation 9. Equation for the Ljung-Box test

$$\hat{Q}_W = n(n + 2) \sum_{j=1}^{h} \frac{\hat{p}^2 w(j)}{n - j} \sim x_{1-\alpha}^2(h - p - q)$$

Since the computed p-value of 0.598 (Table 2) is also greater than 0.05, we again fail to reject the data independence hypothesis; therefore, stating the residuals resemble a white noise.

McLeod-Li Test: Another portmanteau test, the McLeod-Li test can be used to further analyze the white noise hypothesis, stating that if the data resemble white noise from the Ljung-Box test (EQ. 9), then the squared data are also white noise [13]. We use the McLeod-Li test (EQ. 10) where $h = \sqrt{n}$.

Equation 10. Equation of the McLeod-Li test calculation

$$\hat{Q}_{WW} = n(n + 2) \sum_{j=1}^{h} \frac{\hat{p}^2 ww(j)}{n - j} \sim x_{1-\alpha}^2(h)$$

The p-value far exceeds 0.05 (Table 2), so we do not reject the hypothesis that the square residuals are white noise, and state the residuals resemble white noise. Since the results of the portmanteau tests indicate that the residuals demonstrate strong model selection, we proceed by comparing a forecast of Southwest's revenue to the actual reported revenue.

Following the statistical tests, we check whether the model forecasts revenue for Southwest accurately. We train the model on a dataset of quarterly reported revenues from 2005 to 2018 and test it on the 4 quarters of 2019 (FIG. 12, revenue in millions USD). We see that the forecast (marked with red orbs) was approximately 5% higher than the true observation (marked by the blue line). In order to resolve this issue, we reduced the predictions of the model by 5% and arrived at Figure 13 (revenue listed in millions USD). Following this change, the percent error was reduced to a maximum of 2.11%.
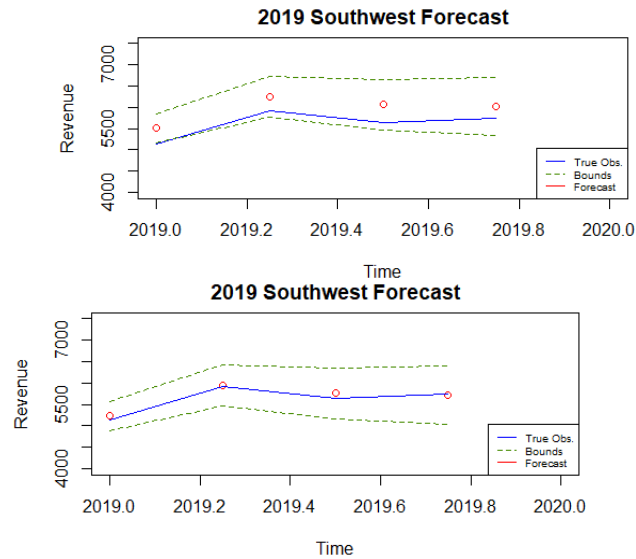
2019 Southwest Forecast

2019 Southwest Forecast

Figure 12. Forecast of 2019 Southwest Revenue　　　　**Figure 13.** Forecast with 5% reductions

The reason for this 5% reduction was due to the grounding of the BOEING 737-MAX. On March 10, 2019, the plane was grounded due to the two 737-MAX tragedies with Lion Air and Ethiopian Airlines. Due to Southwest's airline business model, its fleet of planes reduced [14]. The loss in flight capacity amounted to approximately 8% due to the groundings [15]. As a result, the company reported significantly lower revenue for the second quarter of the 2019 fiscal year, and the following quarters were lower than the forecast, likely due to the recovery period following the 737-MAX grounding.

Forecasting and Risk

Having determined that $SARIMA(0,1,0) \times (0,1,1)_4$ can accurately predict Southwest's revenue, we look at two cases. We examine Southwest's financial performance, had the pandemic not taken place, to understand Southwest's cash flow and the workable earnings for the year of 2020. We then take a look at the impact COVID-19 had on the revenue for Southwest, and the implications of this impact based on the unchanging financial costs that would still need to be paid for the fiscal year of 2020.

In order to understand the financial position of Southwest, we had a look at Southwest's financial statement as well, specifically the Income Statement and Balance Sheet [16]. We assumed that the costs in the income statement and assets plus liabilities in the balance sheet would assume a growth as they have been for the past decade.

For the income statement, to calculate the total operating expenses, we multiplied its components: the selling, general, administrative (SGA), salaries, wages, benefits, and fuel, oil expenses by their mean growth percentage in each respective category from their 2019 values. The total operating income was computed by subtracting the Total Operating Expenses from the Operating revenue for the fiscal year.

For the Balance sheet, we did the same, computing assets and liabilities by multiplying 2019s value by a mean percent growth of the past decade.

Base Case Forecasting

A look at Figure 14 shows that the revenue of Southwest would have been close to approximately 5 billion each quarter. The total sum for the fiscal year of 2020 was forecasted to be approximately $23.3 billion in revenue. Through the process described, we computed the total operating expenses to be approximately $20.3 billion. Hence, the operating income, which is money the company must use for its programs, was approximately $3 billion. Through calculations of the balance sheet, the total current liabilities subtracted from the assets computed a difference of

approximately $3.8 billion, which is the company's deficit. This deficit is resolved by the operating income. Subtracting the two shows that in a normal base case scenario, Southwest airlines would be running close to $800 million in debt, which can be absolved with bank loans or reducing fuel costs.
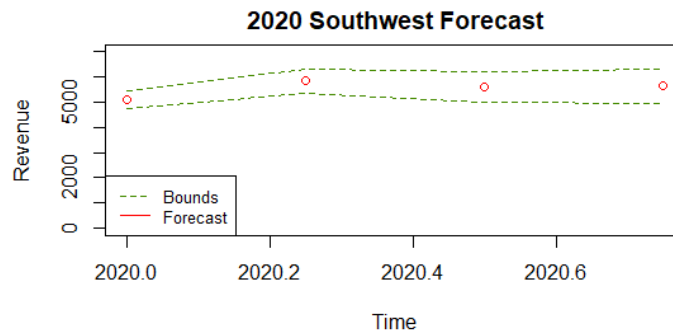

Figure 14. Forecast of 2020 Southwest Revenue

Worst Case Scenario Analysis

Had the pandemic not taken place, we analyze the worst-case scenario for Southwest airlines, indicated by the brown lower bounds in Figure 15. Summing these values, we arrive at an expected $21.2 billion in annual revenue. Assuming the same operating expenses, the operating income is calculated to be $1.2 billion only. Assuming the same deficit as the base case, subtracting the deficit from the operating income results in $2.5 billion negative, which Southwest would incur as debt. That is an additional $1.7 billion in debt which is troublesome for Southwest and a reason for concern. Southwest incurred huge losses due to the pandemic.
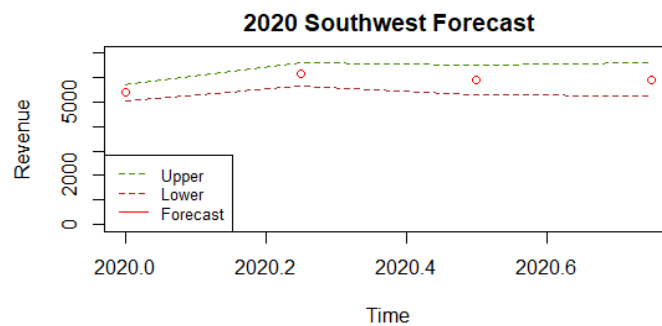

Figure 15. Lower bounds highlighted in Brown Color

Impact of COVID-19 Pandemic

The impact of COVID-19 was drastic on the financial stability of Southwest in 2020 (FIG. 16). Southwest reported an annual revenue of $9 billion. $14.4 billion less than the SARIMA model prediction. As a result, the loss of operating income due to COVID-19 was approximately $6.8 billion dollars.
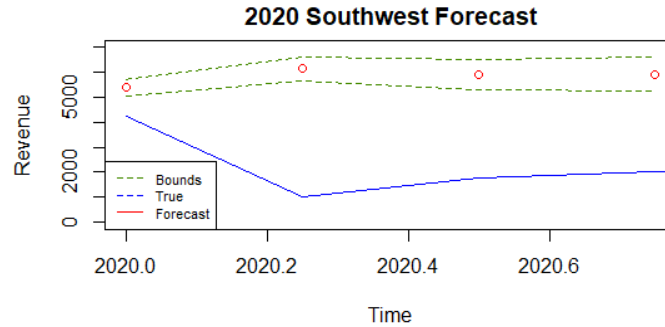
Figure 16. 2020 Forecast with Real Values Included

Southwest experienced $6.8 billion dollars in economic losses. Their annual revenue was not enough to supplement their Total Operating Expenses. The only change in the Operating Expenses was that due to the lack of travel, Southwest did not have to spend as much as it normally would on fuel and oil costs. However regardless of the fuel cuts, Southwest was negative $6.8 in operating income. They were incurring debts. Added to that, neither the assets nor the liabilities changed much from what their expected values were. Thus, Southwest had the previously discussed $3.8 billion dollars in deficit. The $6.8 billion is added to this deficit to arrive at an estimated debt of $10 billion. This came in the form of the $25 billion government bailout [17]. As a result, the reported Total Current Assets for Southwest in 2020, which should have been approximately $6.3 billion, were $15.1 billion. The Long-Term Debt also grew from merely $2 billion in 2019 to an astounding $10.4 billion in 2020. The goal is to gradually pay off this increasing debt as the pandemic progresses and travel starts to reopen.

Conclusion

In this paper we used Southwest Airlines as a case study for business forecasting and risk management. We took the quarterly revenue data for Southwest and differenced it on seasonality and trend, arriving at a stationary time series resembling white noise. Following that, we analyzed the ACFs and PACFs of the seasonally adjusted time series to determine parameter values for p, q, P, Q. Seeing that the spikes were relatively close to the confidence interval, we proposed several values for each parameter, and came up with 4 SARIMA model candidates. We finally arrived at Model (4), $SARIMA (0,1,0) \times (0,1,1)_4$ which had the lowest AICc. We ran diagnostic tests on the model and checked its accuracy with pre-existing data. Since it was accurate, we ran a solvency risk analysis on Southwest Airlines had the COVID-19 pandemic not happened, and the impact of the pandemic on Southwest's financial situation. Seeing the debt and data for southwest, we propose that Southwest should reduce its SG&A and fuel costs in order to make up for the COVID-19 recovery. We also look to develop a model that could estimate the amount of time it would take for Southwest to reapproach the revenue numbers it had prior to the pandemic. While cutting costs, Southwest can look to strike deals with investors and look to resolve the debt as soon as passengers start flying again.

Acknowledgments

References

[1]   F. Fridson M. and Alvarez, Financial Statement Analysis: A Practitioner's Guide, John Willey, 2011.

[2]   G. Box, G. Jenkins, G. Reinsel and G. Ljung, Time Series Analysis: Forecasting and Control, John Wiley, 2015.

[3]   P. Brockwell and R. Davis, Introduction to time series and forecasting, Springer, 2016.

[4]   Macrotrends LLC, "Southwest Airlines Revenue 2005-2020 Chart," [Online]. Available: https://www.macrotrends.net/stocks/charts/LUV/southwest/revenue. [Accessed 28 6 2021].

[5]   R. Shumway and D. Stoffer, Time series: a data analysis approach using R, CRC Press, 2019.

[6]   J. Spitzer, "A Primer on Box-Cox Estimation," The Review of Economics and Statistics, vol. 64, no. 2, pp. 307-313, 1982.

[7]   R. Shumway and D. Stoffer, Time Series Analysis and Its Applications, Springer, 2016.

[8]   C. Hurvich and C. Tsai, "A corrected Akaike information criterion for vector autoregressive model selection," Journal of time series analysis, vol. 14, no. 3, pp. 271-279, Journal of time series analysis.

[9]   R. Johnson and D. Wichern, Applied multivariate statistical analysis (Vol. 6), London, UK: Pearson, 2014.

[10]   G. Ljung and G. Box, "On a measure of lack of fit in time series models," Biometrika, vol. 65, no. 2, pp. 297-303, 1978.

[11]   S. Shapiro and M. Wilk, "An analysis of variance test for normality (complete samples).," Biometrika, vol. 52, no. 3/4, pp. 591-611, 1965.

[12]   G. Box and D. Pierce, "Distribution of residual autocorrelations in autoregressive-integrated moving average time series models," Journal of the American statistical Association, vol. 65, pp. 1509-1526, 1970.

[13]   A. McLeod and W. Li, "Diagnostic checking ARMA time series models using squared-residual autocorrelations," Journal of time series analysis, vol. 4, no. 4, pp. 269-273, 1983.

[14]   A. Sider, "American, southwest under pressure from 737 MAX grounding," 24 10 2019. [Online]. Available: https://www.wsj.com/articles/american-airlines-lowers-earnings-guidance-11571918373. [Accessed 3 11 2021].

[15]   B. Sumers, "Southwest may bear biggest brunt after FAA grounding of Boeing Max Jets," 14 3 2019. [Online]. Available: https://skift.com/2019/03/13/southwest-may-bear-biggest-brunt-after-faa-grounding-of-max-jets/. [Accessed 3 11 2021].

[16]   marketbeat, [Online]. Available: https://www.marketbeat.com/stocks/NYSE/LUV/financials/.

[17]   A. Rappeport and N. Chokshi, "Crippled airline industry to get $25 Billion Bailout, part of it as loans," 14 4 2020. [Online]. Available: https://www.nytimes.com/2020/04/14/business/. [Accessed 16 8 2021].