Time Series Analysis of San Francisco International Airport Quarterly Air Traffic Passenger Counts

TRENZY ODERO

Purpose of the Analysis

The aim of this time series analysis is to model and forecast the total number of air traffic passengers in the United States using quarterly data from 1999 to 2025. Particular attention is given to identifying seasonality, long-term trends, and structural shifts such as those caused by the COVID-19 pandemic. The forecasts produced may inform transportation authorities and policymakers regarding infrastructure, scheduling, and capacity planning decisions.

Studies such as Song and Li (2008) have shown that SARIMA models perform well in tourism and travel demand forecasting [2]. Additionally, Gudmundsson et al. (2021) emphasized the role of recovery forecasting in the air transport market post-COVID [3], while Wu et al. (2021) explored separating supply and demand effects in forecasting frameworks [4]

Description of the Data

The data for this project is sourced from the San Francisco International Airports monthly Air Traffic Passenger Statistics, published through the City and County of San Francisco OpenData portal.

Originally reported monthly by airlines, the data was aggregated to a quarterly level to support the time series analysis. The dataset reflects total passenger counts, including both domestic and international flights, and is inherently seasonal, with peaks typically observed during summer quarters. Documentation recommends year-over-year comparisons to account for this seasonality.

For the analysis, four variables were selected: Year, Quarter, Year_Quarter, and Total_Passengers. The Year and Quarter columns structure the time index chronologically and capture seasonal effects. Year_Quarter combines the two for labeling and visualization. Total_Passengers serves as the response variable for modeling and forecasting. Other attributes, such as airline-specific or monthly breakdowns, were excluded to maintain focus on aggregate passenger trends.

This selection provided a clean quarterly dataset appropriate for modeling seasonal patterns and long-term trends in U.S. air traffic passenger volumes.

Proposed Analyses

This section outlines the statistical methods selected to address the modeling objective. The analysis follows the BoxJenkins methodology, beginning with time series visualization to examine trend, seasonality, and outliers. Stationarity is formally tested using the Augmented DickeyFuller (ADF) test. Due to evidence of non-stationarity, seasonal differencing (lag 4) and first-order differencing are applied.

The autocorrelation function (ACF) and partial autocorrelation function (PACF) plots of the differenced series guide model identification. Based on these diagnostics, a seasonal ARIMA model of the form SARIMA$(p, d, q)(P, D, Q)_4$ is proposed. Model estimation is performed via maximum likelihood estimation.

Residual diagnosticsincluding the LjungBox test, histograms, and QQ plotsare employed to assess model adequacy. If the SARIMA model fails to adequately capture the underlying structure, the analysis allows for alternative approaches, such as exponential smoothing methods or outlier adjustments.

Finally, forecasting is performed for the next three quarters, and 95% prediction intervals are included to quantify forecast uncertainty.

Analysis Result

Time Plot

The time series plot in (Figure 0.1) in the appendix of quarterly U.S. airline passengers from 1999 to 2025 shows a clear seasonal pattern and a steady upward trend through 2019. In 2020, passenger volumes dropped sharply due to the onset of the COVID-19 pandemic, marking a major structural disruption. A recovery begins in 2021, but volumes remain volatile and below pre-pandemic levels. These features support the use of SARIMA modeling to capture seasonality, trend, and the pandemic shock.

Stationarity Check

The Augmented Dickey-Fuller test provides weak and inconsistent evidence of stationarity across model types, with most p-values exceeding 0.05. The ACF shows a slow decay, and the PACF displays strong early spikes, both indicating non-stationarity. Although the time series plot suggests seasonality, the ACF does not show clear seasonal spikes due to the dominating trend and the structural break caused by the COVID-19 pandemic. Together, these results confirm the need for both first-order and seasonal differencing to stabilize the series and reveal its underlying seasonal structure.

Differencing

After applying both first-order and seasonal differencing at lag 4, the series appears stabilized. The ACF shows a significant spike at lag 1 and a weaker seasonal signature, while the PACF displays no strong remaining structure. The autocorrelations decay quickly and fall mostly within the confidence bounds, suggesting that the series is now approximately stationary. These results confirm that differencing was effective and that a SARIMA(0,1,1)(0,1,1)$_4$ model is appropriate for capturing both short-term and seasonal dependencies.

Model Estimation

Among the models evaluated, the SARIMA$(0,1,1)(0,1,1)_4$ model achieved the lowest Akaike Information Criterion (AIC) value of 3056, indicating the best fit with minimal complexity. Although SARIMA$(1,1,1)(0,1,1)_4$ had a comparable AIC of 3058, the simpler SARIMA$(0,1,1)(0,1,1)_4$ model was preferred due to its parsimony and nearly equivalent performance.

The fitted model can be expressed as:

$$(1-B)(1-B^4)Y_t = (1+0.2542B)(1-0.9274B^4)\varepsilon_t - 345.0$$

where:

- Y_t is the air passenger series at time t,
- B is the backshift operator,
- $(1-B)$ represents regular first differencing,
- $(1-B^4)$ represents seasonal differencing with a period of 4 (quarterly),
- ε_t is a white noise error term,

- Model parameters are:

- $\theta_1 = +0.2542$ (non-seasonal MA, $p = 0.0139$),
- $\Theta_1 = 0.9274$ (seasonal MA, $p < 0.001$),
- $\mu = -345.0$ (mean term, $p = 0.8294$, not statistically significant).

This model effectively captures the trend component through regular differencing and seasonal patterns through quarterly differencing. The significance of the MA parameters at both regular and seasonal levels suggests that the model appropriately accounts for short-term shocks and seasonal persistence in the data.

Residual Diagnostics

The residual correlation diagnostics indicate that the SARIMA$(0,1,1)(0,1,1)_4$ model adequately captures the autocorrelation structure in the data. The ACF, PACF, and IACF plots of the residuals show that most spikes lie within the 95% confidence bounds, suggesting no significant remaining autocorrelation. This is further supported by the LjungBox test, whose p-values exceed 0.05 across multiple lags, affirming that the residuals resemble white noise.

Normality Test on Residuals

Normality tests conducted on the model residuals show clear evidence against the assumption of normality. The ShapiroWilk, KolmogorovSmirnov, Cramrvon Mises, and AndersonDarling tests all yield p-values less than 0.01, leading to rejection of the null hypothesis. The histogram of residuals is sharply peaked and skewed, while the QQ plot exhibits strong deviations from the diagonal line, particularly in both tails. Although this non-normality does not compromise the accuracy of point forecasts, it cautions against relying on normal-theory-based inference or prediction intervals without adjustment.

Forecasting Results

The SARIMA$(0, 1, 1)(0, 1, 1)_4$ model was used to forecast U.S. airline passenger totals for the next eight quarters (2025Q22027Q1). The forecasts reflect the seasonal structure of the data, with expected peaks in Q3 of each year and slight dips in Q1 and Q4.

The forecast plot and table show that predicted passenger volumes remain below pre-pandemic highs but suggest steady recovery. However, the 95% confidence intervals widen substantially over time, indicating increased uncertainty. Notably, lower bounds fall below zero starting in 2026a nonphysical result due to residual non-normality and the limits of Gaussian-based interval estimates.

Standard errors grow from approximately 1.3 million to over 9.8 million by the eighth quarter, reinforcing this uncertainty. Despite these limitations, the model captures key seasonal dynamics and provides reliable short-term forecasts, consistent with recovery projections found in [3].

Conclusion

The SARIMA$(0, 1, 1)(0, 1, 1)_4$ model effectively captured the seasonal structure, long-term trends, and the disruption caused by the COVID-19 pandemic in quarterly U.S. airline passenger data. The use of first-order and seasonal differencing led to a well-specified model, with the lowest AIC value indicating optimal balance between fit and complexity.

Residual diagnostics confirmed the removal of autocorrelation, with ACF, PACF, and Ljung-Box results suggesting white-noise residuals. However, strong deviations from normality were observed in the QQ plot and histogram, including skewness and heavy tails. While this does not compromise the accuracy of point forecasts, it undermines the reliability of normal-theory-based confidence intervals.

Forecasts for 2025 to 2027 suggest a gradual recovery in air travel demand, with recurring seasonal peaks in Q3. Yet, the widening prediction intervals and negative lower bounds in later quarters highlight the limitations of the model in long-range forecasting.

Recommendations. Future work could explore SARIMA extensions that handle non-normality, such as bootstrap-based intervals or SARIMA-GARCH models for heteroskedastic errors. Incorporating exogenous variablessuch as fuel prices, economic indicators, or policy changesmay improve explanatory power. Additionally, continued model validation using post-2027 data is recommended to ensure robustness in dynamic recovery periods.

APPENDICES

References

[1] Zahed, M. (2025). *STAT 5037 Time Series Analysis Course Notes*. East Tennessee State University, Spring 2025.

[2] Song, H., & Li, G. (2008). Forecasting tourism demand using ARIMA: The case of China. *Annals of Tourism Research*, 35(4), 11151127. https://doi.org/10.1016/j.annals.2008.05.006

[3] Gudmundsson, S., Cattaneo, L., & Redondi, C. (2021). Forecasting recovery time in air transport markets in the presence of large economic shocks: COVID-19. *Journal of Air Transport Management*, 91, 102007. https://doi.org/10.1016/j.jairtraman.2020.102007

[4] Wu, C., Zhang, H., & Rico, M. I. (2021). Separation of supply and demand effects in the passenger air transport market during the COVID-19 pandemic. *Transportation Research Part E: Logistics and Transportation Review*, 147, 102228. https://doi.org/10.1016/j.tre.2021.102228

SAS Output

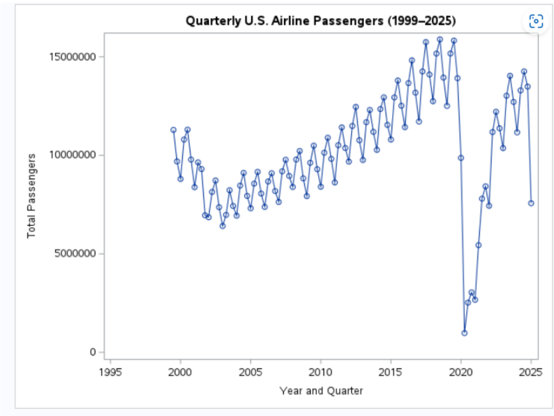

Autocorrelation Check for White Noise									
To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	177.66	6	<.0001	0.805	0.584	0.537	0.508	0.318	0.135
12	183.40	12	<.0001	0.127	0.138	0.018	-0.101	-0.068	0.006
18	186.25	18	<.0001	-0.056	-0.118	-0.038	0.063	0.026	-0.012
24	193.96	24	<.0001	0.095	0.162	0.087	-0.010	0.036	0.119

Augmented Dickey-Fuller Unit Root Tests							
Type	Lags	Rho	Pr < Rho	Tau	Pr < Tau	F	Pr > F
Zero Mean	0	-1.7546	0.3593	-1.03	0.2697		
	1	-2.1308	0.3142	-1.07	0.2540		
	2	-0.6393	0.5388	-0.57	0.4684		
Single Mean	0	-19.2504	0.0106	-3.20	0.0231	5.13	0.0347
	1	-31.5154	0.0009	-3.84	0.0036	7.37	0.0010
	2	-11.1739	0.0927	-2.25	0.1917	2.52	0.4358
Trend	0	-21.9088	0.0362	-3.38	0.0592	5.74	0.0817
	1	-36.1272	0.0010	-4.01	0.0113	8.11	0.0075
	2	-13.0354	0.2409	-2.36	0.3979	2.80	0.6213

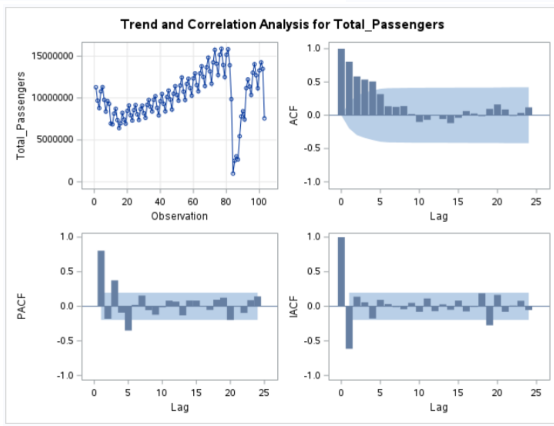

Figure 1: Time plot and stationarity check

Tests for Normality				
Test	Statistic		p Value	
Shapiro-Wilk	W	0.571786	Pr < W	<0.0001
Kolmogorov-Smirnov	D	0.263895	Pr > D	<0.0100
Cramer-von Mises	W-Sq	1.919826	Pr > W-Sq	<0.0050
Anderson-Darling	A-Sq	10.46389	Pr > A-Sq	<0.0050

Goodness-of-Fit Tests for Normal Distribution				
Test	Statistic		p Value	
Kolmogorov-Smirnov	D	0.2638948	Pr > D	<0.010
Cramer-von Mises	W-Sq	1.9198264	Pr > W-Sq	<0.005
Anderson-Darling	A-Sq	10.4638896	Pr > A-Sq	<0.005

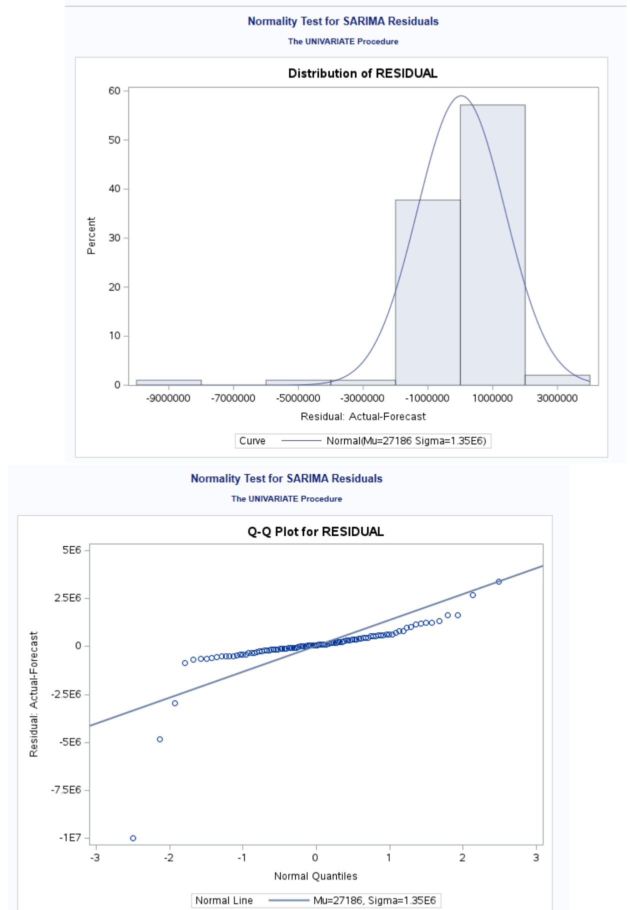

Figure 2: nomality for residuals check

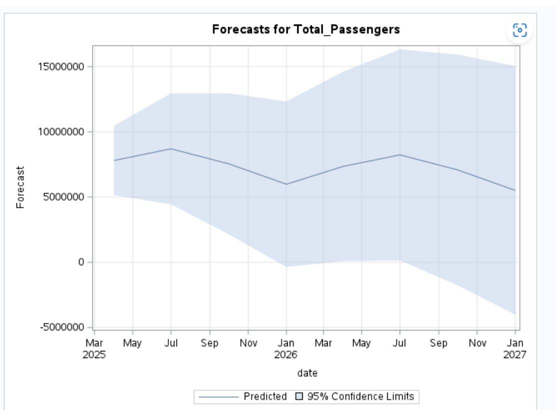

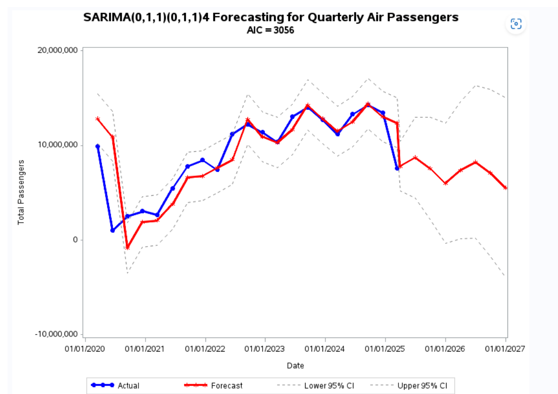

Forecast Values for Air Passengers

date	Total_Passengers	FORECAST	L95	U95
2025Q1	7,570,655	12,395,662	9,736,988	15,054,336
2025Q2	.	7,805,008	5,150,723	10,459,293
2025Q3	.	8,700,377	4,442,720	12,958,035
2025Q4	.	7,550,788	2,146,155	12,955,421
2026Q1	.	5,984,349	-363,286	12,331,985
2026Q2	.	7,347,121	87,995	14,606,247
2026Q3	.	8,237,145	147,352	16,326,938
2026Q4	.	7,082,210	-1,760,560	15,924,980
2027Q1	.	5,510,426	-4,026,052	15,046,904

Figure 3: Forecasting

SAS CODES

Listing 1: Filtering Data in SAS

```
/* data import */

FILENAME REFFILE '/home/u63992502/final_quarterly_data.csv';

PROC IMPORT DATAFILE=REFFILE
        DBMS=CSV
        OUT=passenger;
        GETNAMES=YES;
RUN;

/* Step 1: Explore the Cleaned Data */
PROC CONTENTS DATA=passenger;
RUN;

PROC PRINT DATA=passenger (OBS=5);
RUN;

/* Step 2: Visualize the Original Time Series */
PROC SGPLOT DATA=passenger;
    SERIES X=Year_Quarter Y=Total_Passengers / markers;
    XAXIS LABEL="Year and Quarter";
    YAXIS LABEL="Total Passengers";
    TITLE "Quarterly U.S. Airline Passengers (1999  2025)";
RUN;
```

```
/* Step 3: Stationarity Check (ADF Test) */
PROC ARIMA DATA=passenger;
    IDENTIFY VAR=Total_Passengers;
    TITLE "ADF Test for Stationarity";
RUN;

/* Step 4: ACF and PACF for the Differenced Series */
/* Apply first differencing and seasonal differencing (lag 4) */
PROC ARIMA DATA=passenger;
    IDENTIFY VAR=Total_Passengers(1,4);
    TITLE "Differenced Series - First and Seasonal Differencing (Lag 4)";
RUN;

/* Step 6: Model Estimation - Fit SARIMA(0,1,1)(0,1,1)_4 */
TITLE "SARIMA(0,1,1)(0,1,1)_4 Model Estimation";
PROC ARIMA DATA=passenger;
    IDENTIFY VAR=Total_Passengers(1,4);
    ESTIMATE q=(1)(4) METHOD=ML;
    FORECAST OUT=forecast_resid ID=Year_Quarter INTERVAL=QUARTER;

RUN;

/* Step 7: Forecast the Next 4 Quarters */
PROC UNIVARIATE DATA=forecast_resid NORMAL;
    VAR residual;
    HISTOGRAM / NORMAL;
```

```
    QQPLOT / NORMAL(MU=EST SIGMA=EST);
    TITLE "Normality Test for SARIMA Residuals";
RUN;

/* Step 8: Forecast the Next 4 Quarters */
PROC ARIMA DATA=passenger;
    IDENTIFY VAR=Total_Passengers(1,4);
    ESTIMATE q=(1)(4) METHOD=ML;
    FORECAST LEAD=24 OUT=forecast_passengers;
    TITLE "Forecasting U.S. Airline Passengers (Next 4 Quarters)";
RUN;



/******************/
/* Forecasting */
/* Step 1: Use the air passenger dataset */
data data1;
    set passenger;
    /* Ensure date variable exists for time series - create if needed */
    date = mdy(Quarter*3, 15, Year);
    format date mmddyy10.;
run;

/* Step 2: Fit SARIMA(0,1,1)(0,1,1)4 model to Total_Passengers */
proc arima data=data1;
    identify var=Total_Passengers(1,4);  /* First difference and seasonal dif
    estimate q=(1)(4) method=ml;  /* SARIMA(0,1,1)(0,1,1)4 specification */
```

```
    forecast printall lead=8 interval=qtr id=date out=forecast;
run; quit;

/* Step 3: Plot the forecast with enhanced visuals */
goptions reset=all;
title "SARIMA(0,1,1)(0,1,1)4 Forecasting for Quarterly Air Passengers";
title2 "AIC = 3056";
axis1 label=(angle=90 'Total Passengers');
axis2 label=('Date') order=('1JAN2020'd to '1JAN2027'd by year);
symbol1 i=join l=1 w=2 c=blue v=dot;       /* Actual values */
symbol2 i=join l=1 w=2 c=red v=star;       /* Forecast values */
symbol3 i=join l=2 w=1 c=gray v=none;      /* Lower 95% CI */
symbol4 i=join l=2 w=1 c=gray v=none;      /* Upper 95% CI */
legend1 label=none frame
    value=('Actual' 'Forecast' 'Lower 95% CI' 'Upper 95% CI')
    position=(bottom) across=4;

proc gplot data=forecast;
    plot (Total_Passengers forecast l95 u95) * date / overlay
        legend=legend1 vaxis=axis1 haxis=axis2 vminor=0 hminor=0;
    format Total_Passengers forecast l95 u95 comma12.0;
run; quit;

/* Step 4: Create a table of forecast values */
title "Forecast Values for Air Passengers";
proc print data=forecast(where=(date > '01JAN2025'd)) noobs;
    var date Total_Passengers forecast l95 u95;
```

```
    format date yyq. Total_Passengers forecast l95 u95 comma12.0;
run;
```