Time Series on the United States Quarterly Gross Domestic Product Growth

Daniel Ramos

2025-03-21

This project analyzes the quarterly percentage growth of U.S. real GDP from 1950 to 2024 using time series modeling. Seasonal differencing was applied to stabilize the series, and a SARIMA model was initially fitted to account for autocorrelation and seasonality. However, diagnostic tests revealed volatility clustering in the residuals, particularly around the early 1980s, 2008, and 2020, prompting the use of a GARCH model to better capture time-varying variance. A unified ARMA-GARCH model was then fitted to the differenced GDP series, effectively capturing both short-term dynamics and conditional volatility, with all key parameters statistically significant. The results demonstrate that incorporating volatility modeling provides a more complete understanding of the behavior and uncertainty in U.S. GDP growth over time.

Table of Contents

I. Introduction

One of the most important indicators used globally to assess economic performance is the Gross Domestic Product (GDP). The monetary value of all goods and services produced within a country serves as a general measure of economic health. Viewed as a time series, GDP data reveals cycles of expansion and contraction, as well as the impact of sudden economic shocks. This project uses the quarterly GDP growth rate of the United States from 1950 to 2024 to analyze seasonal trends and volatility patterns.

The purpose of the project is to compare several models' ability to capture predictable patterns in GDP growth, while simultaneously accounting for periods with higher uncertainty. The simple model is a SARIMA, seasonally adjusted for quarterly data. A second model, ARMA-GARCH, is then introduced to account for the irregular volatility caused by economic shocks such as financial crises and global pandemics. Additionally, a forecast of the most probable range for the next eight quarters is provided.

GDP growth was chosen as the subject of this project because of its relevance to economic policy, forecasting, and cross-country comparison. Having studied macroeconomic indicators such as unemployment and interest rates previously, this project offered an opportunity to examine GDP dynamics through a statistical framework rather than a purely economic one.

II. Data

The dataset consists of quarterly U.S. GDP values, measured in millions of U.S. dollars, sourced from the U.S. Bureau of Economic Analysis (BEA). The original data spans from Q1 1942 to Q3 2024; however, the series was trimmed to begin in Q1 1950, yielding 299 quarterly observations after transformation. The early years were excluded due to the extreme volatility associated with the final years of World War II and its immediate aftermath, which made reliable model fitting impractical. The data is publicly available through the Federal Reserve Bank of St. Louis: https://fred.stlouisfed.org/series/NA000334Q

To prepare the data for time series modeling, GDP levels were transformed into quarter-over-quarter percentage changes. This transformation stabilizes the variance, enables meaningful comparison across time and with other countries, and makes the series more suitable for modeling. Logarithmic transformation was not feasible given the presence of negative values, and standardization was avoided to preserve the interpretability of positive and negative growth. Seasonal differencing with lag 4 was then applied to remove recurring quarterly patterns and ensure stationarity, a required condition for the models used in this analysis.

The BEA compiles GDP estimates using data from the Census Bureau, the Bureau of Labor Statistics, and the Internal Revenue Service, making it a highly reliable and widely cited source in economic research and policy analysis. The dataset's long time span and quarterly frequency make it particularly well suited for time series modeling, as it captures multiple distinct economic regimes and a range of major historical events including the 2008 financial crisis and the COVID-19 pandemic.

III. Methodology

This project applies two time series modeling approaches to analyze and forecast quarterly U.S. GDP growth: a Seasonal ARIMA (SARIMA) model and an ARMA-GARCH model. The SARIMA model captures predictable patterns in the mean of the series, while the ARMA-GARCH model additionally accounts for volatility clustering in the residuals, providing a more complete picture of both the predictable behavior and the uncertainty in GDP growth over time.

a. SARIMA (p, d, q) x (P, D, Q) model

- p, d, q are the non-seasonal autoregressive order, differencing degree, and moving average order, respectively
- P, D, Q are their seasonal counterparts, and
- S is the seasonal period (set to 4 for quarterly data).

  Prior to fitting the model, two conditions must be satisfied: a constant mean and constant residual variance. The first condition was achieved through seasonal differencing at lag 4. Based on inspection of the ACF and PACF plots and iterative diagnostic testing, a SARIMA (0,0,2) × (0,0,2) [4] specification was selected. This translates to a model in which current GDP growth depends on shocks from the previous two quarters, as well as shocks from the same quarter in each of the previous two years.

b. ARMA-GARCH Model

To capture the volatility patterns that the SARIMA model could not explain, an ARMA-GARCH model was fitted to the same transformed GDP growth series. This approach allows for simultaneous modeling of both the mean and the variance of the series. The GARCH component models time-varying variance, specifically how the current level of volatility depends on past forecast errors and past variance.

Based on parameter significance, model diagnostics, and AIC comparison across multiple specifications, an ARMA(2,2)-GARCH(1,1) model was selected. This specification assumes that current GDP growth is influenced by its own past values and past shocks up to two lags, and that expected volatility in any given quarter depends on the volatility and shock size from the previous quarter.

IV. Results

a. Results from SARIMA (0, 0, 2) x (0, 0, 2)

Figure 1

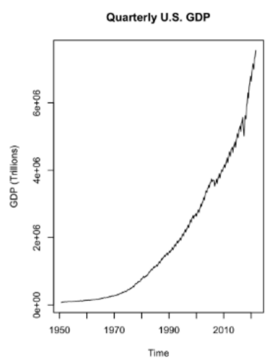

Figure 2

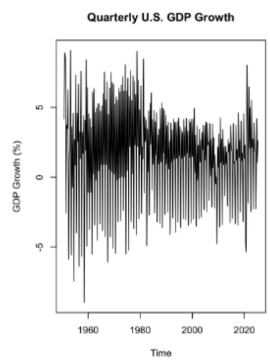

Figure 3

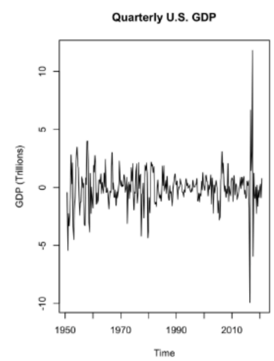

Figure 4

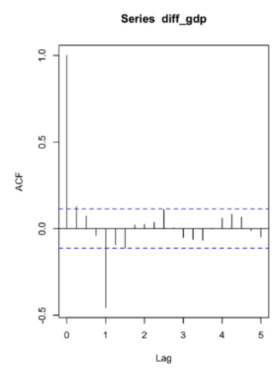

Figure 5

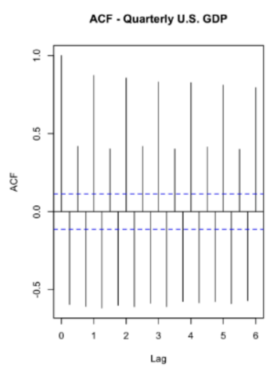

Figure 6

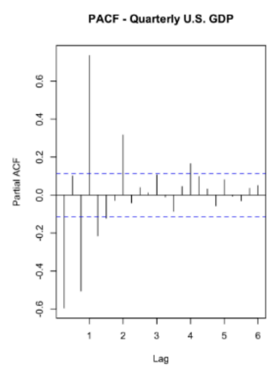

Figure 7

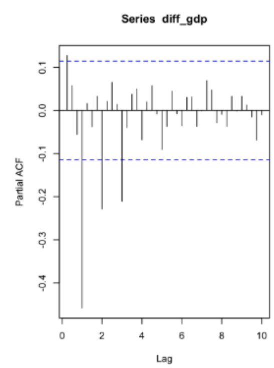

Figures 1 through 7 present the exploratory plots used to identify an appropriate model specification. Figures 1–4 establish the need for differencing, while Figures 5–7 guide parameter selection.

- **Figure 1 - Quarterly U.S. GDP (Levels):** The series shows clear exponential growth over time, confirming that the raw dollar values are non-stationary and unsuitable for direct modeling.
- **Figure 2 - Quarterly U.S. GDP Growth (%):** Converting to percentage changes substantially reduces the trend, though visible volatility clustering remains.
- **Figure 3 - GDP Growth After Seasonal Differencing:** Each bar represents the quarterly change in GDP growth from the same quarter one year prior. After differencing, the series fluctuates more uniformly around zero with no persistent trend, satisfying the stationarity requirement for SARIMA modeling.
- **Figure 4 - ACF of Differenced Series:** The ACF drops sharply after lag 1, with all subsequent lags falling within the confidence bounds. This clean cutoff suggests a moving average order of q = 1 as the appropriate starting point for model selection.
- **Figure 5 - ACF of Original GDP Growth:** Autocorrelation remains strongly significant across all displayed lags and shows no sign of decay, confirming that the original series is non-stationary and requires differencing before modeling.

- **Figure 6 - PACF of Original GDP Growth:** Large spikes appear at lags 1, 2, and 3, with the spike at lag 1 being particularly prominent. These spikes correspond to the end-of-year seasonal pattern in quarterly data and are attributable to seasonality rather than genuine autoregressive structure, supporting a choice of p = 0.
- **Figure 7 - PACF of Differenced Series:** After differencing, the partial autocorrelation structure is considerably cleaner, with most lags falling within the confidence bounds. The remaining mild spikes are within acceptable range and do not suggest additional autoregressive terms are needed.

Figure 8

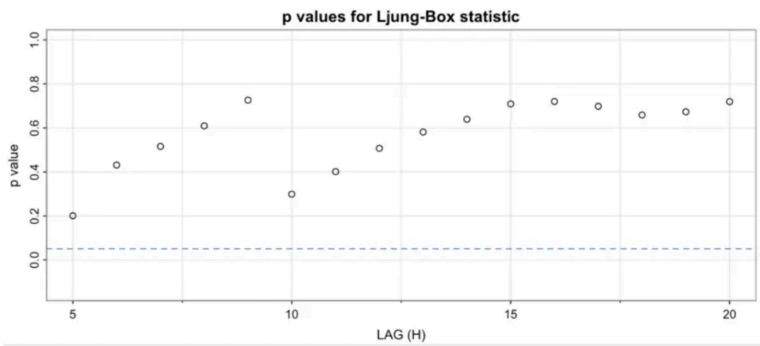

Figure 9

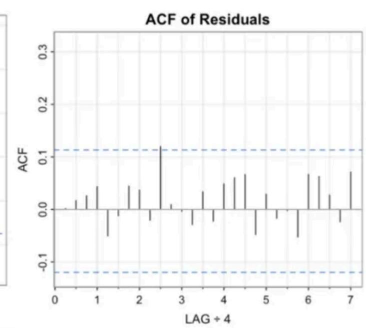

- **Figure 8 - Ljung-Box p-values for SARIMA Residuals:** All p-values across lags 5 through 20 fall well above the 0.05 significance threshold, indicated by the dashed line. This confirms that the residuals are consistent with white noise and that the SARIMA model has adequately captured the autocorrelation structure of the series.
- **Figure 9 - ACF of SARIMA Residuals:** The residual autocorrelation is statistically insignificant at nearly all lags, with the exception of a mild spike at lag 10 quarters (displayed as 2.5 on the LAG ÷ 4 axis). While this spike marginally crosses the confidence bound, it does not constitute a systematic pattern and does not materially compromise model adequacy.

Figure 10

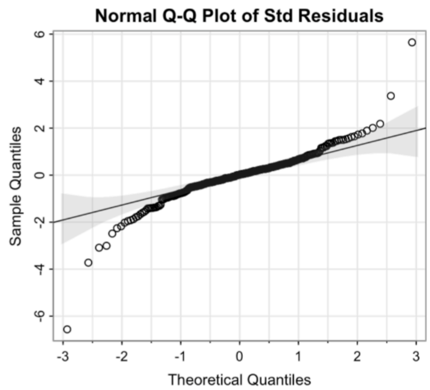

Figure 11

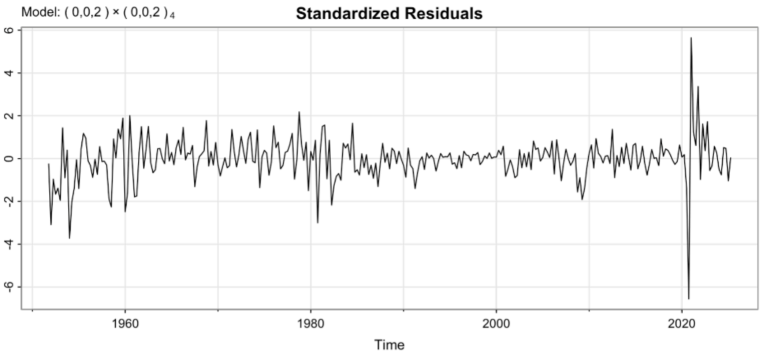

- **Figure 10 — Normal Q-Q Plot of Standardized Residuals:** Residuals follow the normal line closely across most of the range, supporting the assumption of approximate normality. The outliers in both tails are attributable to the 2020 pandemic, which produced swings far outside the range of typical quarterly variation.
- **Figure 11 — Standardized Residuals Over Time:** Residuals fluctuate consistently around zero throughout the series with no systematic pattern, indicating a well-fitted mean equation. The sharp spike near 2020 reflect the pandemic contraction and recovery. The rebound appears dramatic largely because GDP was recovering from an abnormally low baseline rather than experiencing genuine accelerated growth.

The model was kept relatively simple to avoid overfitting. A SARIMA (0,0,1) × (0,0,1) [4] was considered as a simpler alternative, but retaining q = Q = 2 was justified by the Ljung-Box p-values, which approached the significance threshold under the simpler specification.

b. ARMA(2,2)-GARCH(1,1)

To better capture both the mean dynamics and volatility of U.S. GDP growth, an ARMA(2,2)-GARCH(1,1) model was fitted to the differenced series. The ARMA(2,2) specification was selected after comparing multiple (p,q) combinations, with (2,2) producing the lowest AIC of 3.36 while keeping the parameter count manageable. The GARCH(1,1) component was added to capture volatility clustering, where large movements in GDP growth tend to be followed by similarly large movements. All parameters were statistically significant, and diagnostic tests confirmed that the model adequately captured both the autocorrelation structure and time-varying variance in the series.

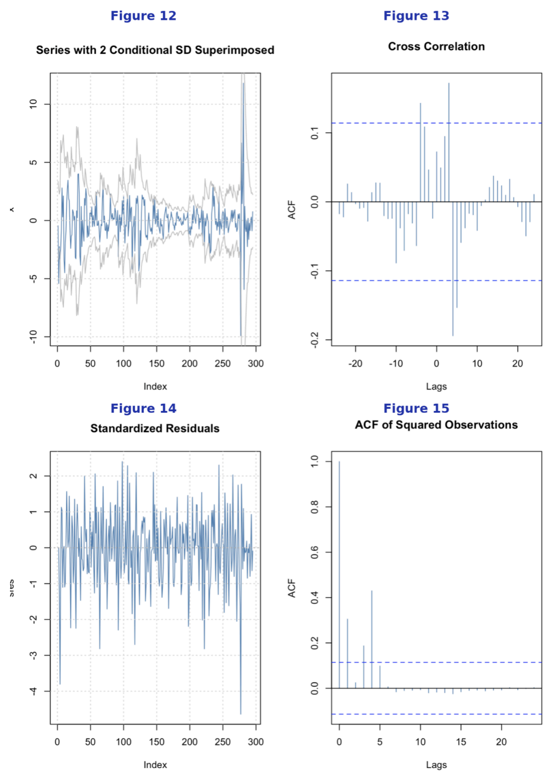

- **Figure 12 - Series with 2 Conditional SD Superimposed:** The x-axis represents each of the 299 quarterly observations in the dataset rather than calendar years. The grey bands show two conditional standard deviations above and below the GDP growth series, widening during periods of elevated volatility and narrowing during stable ones. This is the key visual takeaway from the GARCH model: rather than assuming constant variance, the model adjusts its uncertainty estimates based on recent behavior. When the preceding quarters show large swings, the model anticipates continued volatility and widens the band accordingly. The bands are notably wide in the early portion of the series and compress through the more stable period of the 1990s and 2000s, before widening dramatically near observation 290 during the 2020 pandemic.

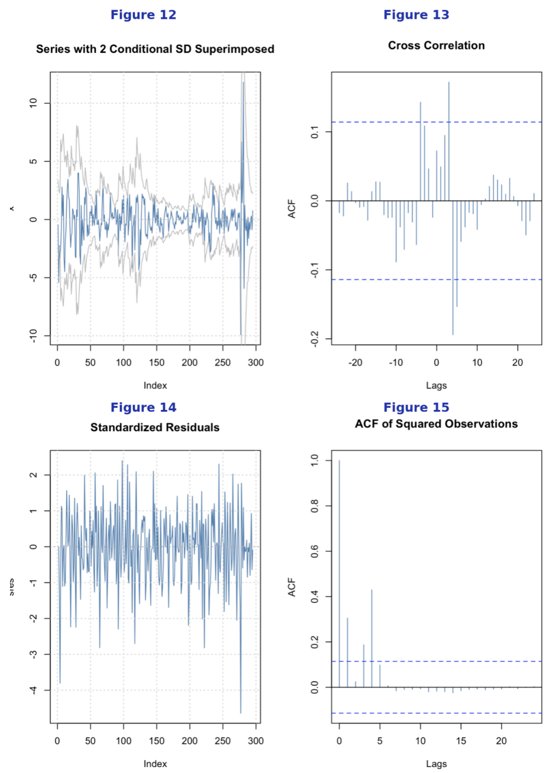

- **Figure 13 — Cross Correlation:** The cross correlation between standardized residuals and squared residuals shows no systematic pattern, with most values falling within the confidence bounds. This indicates that the mean and variance equations of the model are not misspecified.
- **Figure 14 — Standardized Residuals:** The standardized residuals fluctuate around zero with relatively consistent spread throughout the series, suggesting the GARCH component has successfully accounted for the time-varying variance. No persistent clustering of large values is visible outside of the pandemic period near observation 280
- **Figure 15 — ACF of Squared Observations:** The ACF of the squared observations drops sharply after lag 1 and remains within the confidence bounds at most subsequent lags. This confirms that the GARCH(1,1) specification has effectively captured the volatility clustering present in the original series.

c. Forecast

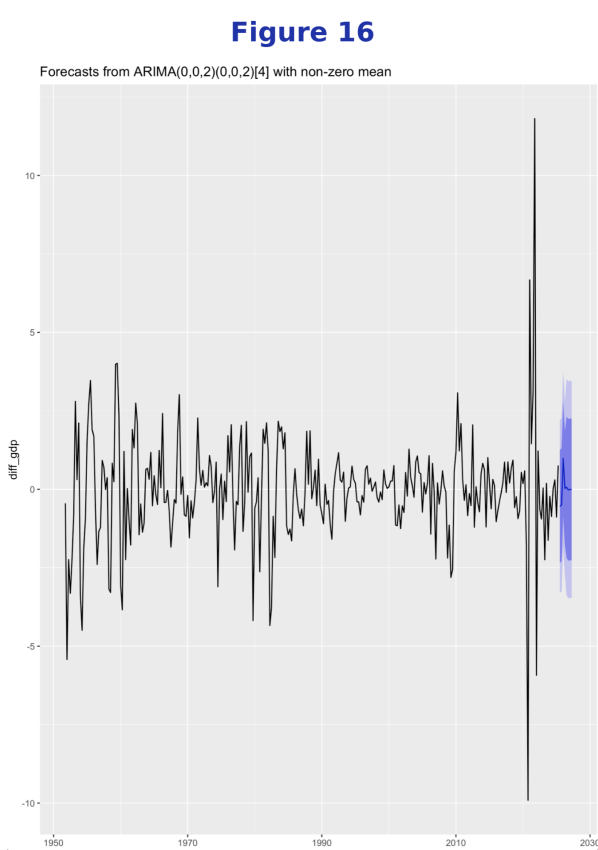

- **Figure 16 - 8-Quarter Forecast from SARIMA (0,0,2) × (0,0,2) [4]:** The shaded blue region at the right of the plot represents the forecast interval for the next 8 quarters, extending from late 2024 through 2026. The darker inner band shows the 80% confidence interval and the lighter outer band shows the 95% interval. The point forecast converges to approximately 3% quarterly growth but it is worth noting that this forecast assumes no major economic shocks. Events like a financial crisis or recession would push growth outside these bounds, as the historical series clearly demonstrates.

V. Conclusion and Future Study

This project analyzed the quarterly percentage growth of U.S. real GDP from 1950 to 2024 using two time series modeling approaches. The goal was to identify patterns in GDP growth, capture shifts in economic uncertainty, and produce a forecast for the near-term future.

Following a Box-Jenkins approach, a SARIMA model was first applied to capture the regular seasonal and short-term patterns in the series. The model identified a consistent seasonal structure in quarterly GDP growth and confirmed that short-term changes are influenced by recent past values. Its key limitation, however, is the assumption of constant variance, which does not hold during periods of economic turbulence.

The ARMA-GARCH model addressed this limitation directly. One of the most meaningful findings of this project is that uncertainty surrounding U.S. GDP growth is not constant — it rises sharply during periods of crisis and compresses during stable periods. The model detected notably elevated volatility during the early 1980s, the 2008 financial crisis, and the COVID-19 pandemic, as illustrated in Figure 12. This highlights the importance of modeling not just average growth, but the risk and uncertainty surrounding it.

The ARMA-GARCH model performed well overall, capturing both the predictable dynamics of GDP growth and the clustering of volatility that follows economic shocks. These insights would not be visible from a simple inspection of the plotted data alone.

This analysis represents a starting point rather than a complete picture. GDP growth is closely interconnected with unemployment and inflation, and incorporating these variables into a multivariate model would likely improve forecast accuracy and provide a richer understanding of the economic dynamics at play. A future study expanding in this direction would be a natural and worthwhile extension of this work.

VI. References

Federal Reserve Bank of St. Louis. (n.d.). *Gross Domestic Product: Final consumption expenditure of households (NA000334Q)* [Data set]. FRED, Federal Reserve Economic Data. https://fred.stlouisfed.org/series/NA000334Q

Santander Group. (n.d.). *What is GDP and why is it important in economics?* https://www.santander.com/en/stori
is-gdp-and-why-is-it-important-in-economics?

Shumway, R. H., & Stoffer, D. S. (2017). *Time series analysis and its applications: With R examples* (4th ed.). Springer.

U.S. Bureau of Economic Analysis. (n.d.). *U.S. Bureau of Economic Analysis (BEA)*. https://www.bea.gov

VII. Appendix (R Code)

```
library(astsa)
library(ggplot2)
library(tidyverse)
library(tseries)
library(xts)
library(forecast)
library(dplyr)
library(fracdiff)

# getwd()  # Get the current working directory
gdp <- read.csv("US_GDP4.csv", header = TRUE) # load CSV File
colnames(gdp) <- c("Date", "GDP") # rename columns
gdp$Date <- as.Date(gdp$Date)
gdp <- gdp %>%
  arrange(Date) %>% # organize by date just in case
  mutate(Growth = (GDP - lag(GDP)) / lag(GDP) *100) # change to percentage

gdp <- na.omit(gdp) # removes first value since it will be zero

start_year <- year(gdp$Date[1]) # 1950
start_month <- month(gdp$Date[1]) # April
end_year <- year(gdp$Date[nrow(gdp)]) # 2024
end_month <- month(gdp$Date[nrow(gdp)]) # October
head(gdp)
```

```
        Date   GDP    Growth
2 1950-04-01 70872  4.174506
3 1950-07-01 77178  8.897731
4 1950-10-01 83746  8.510197
5 1951-01-01 81603 -2.558928
6 1951-04-01 84645  3.727804
7 1951-07-01 87586  3.474511
```

```
gdp_ts <- ts(gdp$Growth, # turn into time series
             start = c(start_year, start_month),
             frequency = 4) # quarterly data => frequency=4
ts.plot(gdp_ts, # plot time series
        main="Quarterly U.S. GDP Growth",
        ylab="GDP Growth in Percent")
```

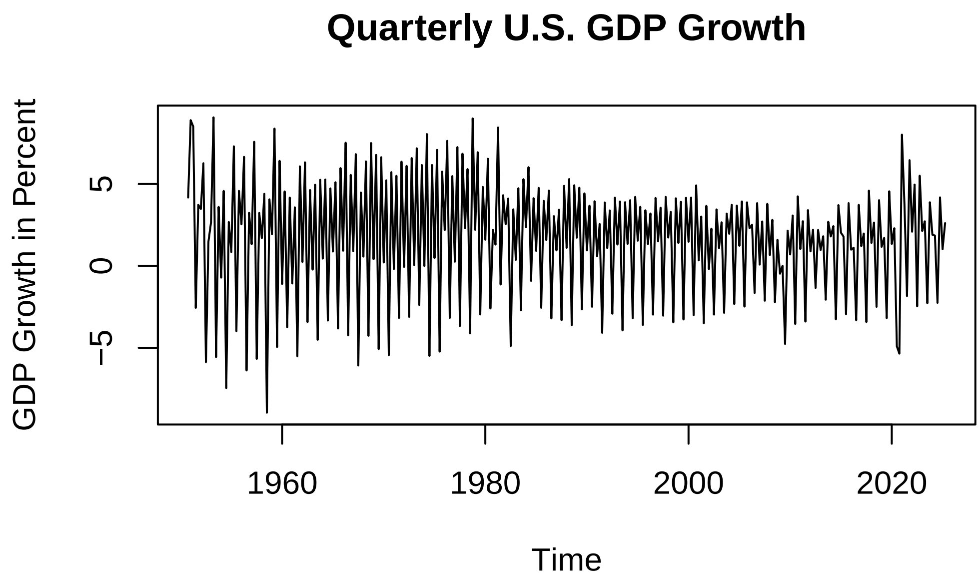


```
# testing stationarity
adf.test(gdp_ts)  # If p-value > 0.05, the series is non-stationary
diff_gdp <- diff(gdp_ts, lag=4, differences=1)  # Seasonal differencing
par(mfrow=c(2,2))
ts.plot(ts(diff_gdp, start = c(start_year, start_month), frequency = 4.2),
        main="Quarterly U.S. GDP", ylab="GDP (Trillions)")
acf(diff_gdp, lag.max = 20) # high at 1, dies
pacf(diff_gdp, lag.max = 40) # high at 1,2,3, dies
adf.test(diff_gdp)
ts.plot(diff_gdp)
```

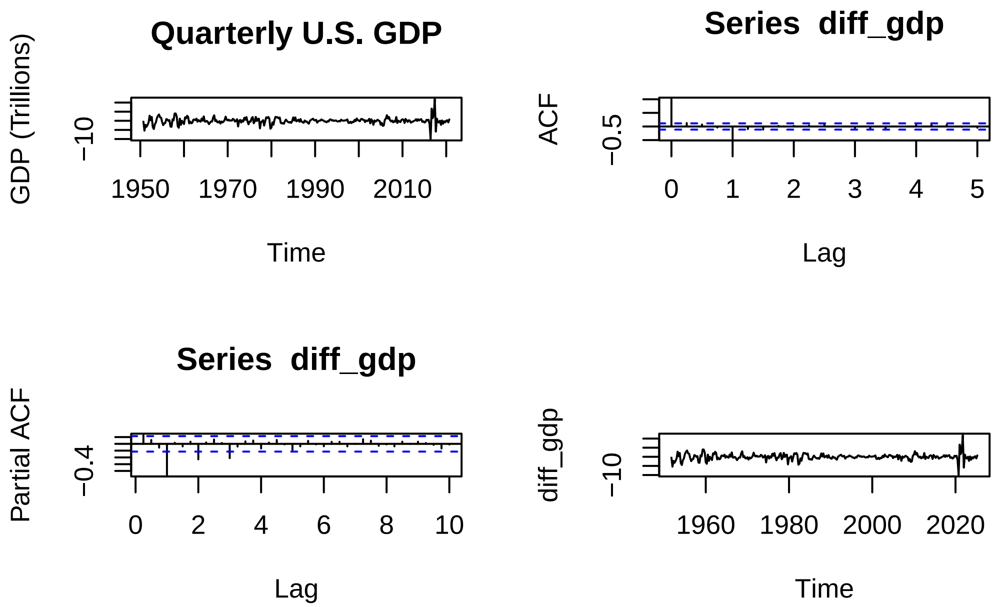


```r
real <- sarima(diff_gdp, p=0, d=0, q=2, P=0, D=0, Q=2, S=4)
```

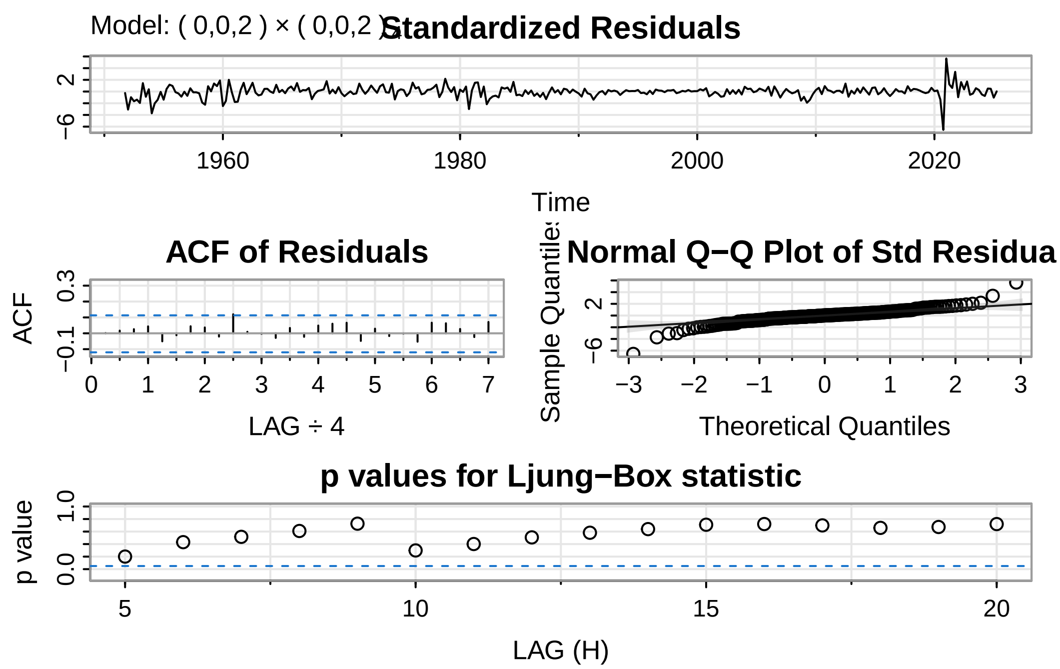


```r
library(fGarch)
garch_model <- garchFit(~ arma(2,2) + garch(1,1),
                                data = diff_gdp,
                                trace = FALSE)
summary(garch_model)
```

```
Title:
 GARCH Modelling

Call:
 garchFit(formula = ~arma(2, 2) + garch(1, 1), data = diff_gdp,
    trace = FALSE)

Mean and Variance Equation:
 data ~ arma(2, 2) + garch(1, 1)
<environment: 0x133aa1378>
 [data = diff_gdp]

Conditional Distribution:
 norm

Coefficient(s):
       mu        ar1        ar2        ma1        ma2      omega     alpha1
 0.016428   0.182971  -0.395467  -0.026886   0.998922   0.033677   0.412424
    beta1
 0.668891

Std. Errors:
 based on Hessian

Error Analysis:
        Estimate  Std. Error  t value Pr(>|t|)
mu      0.016428    0.089587    0.183    0.855
ar1     0.182971    0.038846    4.710 2.47e-06 ***
ar2    -0.395467    0.033223  -11.903  < 2e-16 ***
ma1    -0.026886    0.004177   -6.436 1.22e-10 ***
ma2     0.998922    0.004834  206.647  < 2e-16 ***
omega   0.033677    0.026944    1.250    0.211
alpha1  0.412424    0.100696    4.096 4.21e-05 ***
beta1   0.668891    0.059417   11.258  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Log Likelihood:
 -488.4089    normalized:  -1.655624

Description:
 Thu Mar 26 17:09:41 2026 by user:


Standardised Residuals Tests:
                                 Statistic      p-Value
 Jarque-Bera Test   R    Chi^2  51.9662834 5.195955e-12
```

```
 Shapiro-Wilk Test  R    W       0.9766275 9.564103e-05
 Ljung-Box Test     R    Q(10)  50.9515018 1.782785e-07
 Ljung-Box Test     R    Q(15)  58.5706565 4.433435e-07
 Ljung-Box Test     R    Q(20)  70.9585620 1.269291e-07
 Ljung-Box Test     R^2  Q(10)   7.4997892 6.775681e-01
 Ljung-Box Test     R^2  Q(15)   8.6836624 8.934684e-01
 Ljung-Box Test     R^2  Q(20)  10.8027033 9.511793e-01
 LM Arch Test       R    TR^2    8.1710681 7.716231e-01

Information Criterion Statistics:
     AIC      BIC      SIC     HQIC
3.365484 3.465470 3.364065 3.405521
```

```
par(mfrow = c(2,2))
std_resid <- residuals(garch_model, standardize = TRUE)
acf(std_resid^2, main = "ACF of Standardized Squared Residuals")
Box.test(std_resid^2, lag = 12, type = "Ljung-Box")
```

```
    Box-Ljung test

data:  std_resid^2
X-squared = 7.5059, df = 12, p-value = 0.8225
```

```
# p value is high, tells us the garch aspect was able to capture volatility
```