

Univariate Time Series Data and Model Card

Generated by [Cardtale](#) - Automated Model and Data Card Generator

This report provides an automated, comprehensive analysis of univariate time series data. Generated by Cardtale, it explores basic aspects and potential challenges in your data to support informed decision-making and modeling choices.

Generated: 2025-02-04 23:33

Series Name: M17

Table of Contents

- 1 Data Overview**
Time series fundamental characteristics and statistical properties
- 2 Trend**
Long-term time series growth and dynamics. Analysis of level stabilization methods.
- 3 Seasonality**
Analysing recurring patterns in the time series. Assessing the impact of different seasonality modeling strategies
- 4 Variance**
Exploring the variability of values over time. Assessing the impact of variance stabilization methods

Other aspects were explored but omitted from the final report:

Change Detection

No change point was found according to offline change detection methods

Data Overview

This section examines the core characteristics and statistical properties of the time series. Understanding these attributes is important for assessing data quality and

gaining a preliminary context. We explore the temporal structure, summary statistics, and distribution patterns to create a baseline understanding of your data.

Time Series Plot




Figure 1: Time series line plot.

- A total of 68 observations spanning from January 1990 to August 1995. These are collected with a monthly sampling frequency.
- The data ranges from a minimum of 120 to a maximum of 4600, starting in 2160 and ending in 1560 during the observed period. The average growth percentage per observation is 33.5% (median equal to 0%), with an average value of 1977.94. There are no missing values in the time series.

Trend, Seasonality, and Residuals




Figure 2: Seasonal, Trend, and Residuals components after decomposition on a monthly frequency using the STL (Season-Trend decomposition using LOESS) method.

- The trend strength is 0.23 (ranges from 0 to 1). All hypothesis tests carried out (KPSS, Augmented Dickey-Fuller, and Philips-Perron) indicate that the time series is stationary in trend or level.
- The seasonal strength is 0.64 (ranges from 0 to 1). Tests for yearly seasonal stationarity show mixed results: the OCSB test indicates presence of a seasonal unit root, while the Wang-Smith-Hyndman test suggests stationarity.
- The STL decomposition residuals show balanced behavior: 50.0% of residuals are positive and 50.0% negative. The average magnitude of positive residuals is 359.576 compared to -355.154 for negative residuals. In terms of auto-correlation structure, the residuals show significant temporal dependency in some of the first 12 lags according to the Ljung-Box test. This suggests that the decomposition method is missing some systematic patterns.

Auto-Correlation




Figure 3: Auto-correlation plot up to 24 lags.

- There are no lags with significant correlation, which suggest that the series may be white noise.
- None of the lags relative to the seasonal period (t-12 and t-24) show any significant autocorrelation.

Partial Auto-Correlation




Figure 4: Partial Auto-correlation plot up to 24 lags. At each lag, the partial auto-correlation takes into account the previous correlations.

- There are no lags with significant correlation, which suggest that the series may be white noise.
- None of the lags relative to the seasonal period ($t-12$ and $t-24$) show any significant partial autocorrelation.

Trend

Trend refers to the long-term change in the mean level of a time series. It reflects systematic and gradual changes in the data over time. Understanding the trend is important for identifying long-term growth or decline, structural changes, and making informed modeling decisions. This section examines the characteristics of the trend of the time series.

Trend Line Plot




Figure 5: Time series trend plot.

- There is a slight upward trend. The tests KPSS, Augmented Dickey-Fuller, and Philips-Perron did not find evidence for non-stationarity around a deterministic trend.
- The same tests were applied to analyse stationarity around a constant level. All tests failed to reject this hypothesis.
- **Preliminary experiments:** Including a trend explanatory variable which denotes the position (row id) of each observation does not improve forecasting accuracy. These experiments were conducted using a LightGBM algorithm and evaluated using SMAPE loss function. Using only lag-based features the model achieved a SMAPE of 45.5% on the test set. Including the trend variable leads to a score equal to 49.05%.

Long-term Growth




Figure 6: Distribution of log differences (left), and a Lag-plot (right). These plots help to understand how the data changes over consecutive observations. The histogram show the distribution of these changes using log returns. The lag-plot depicts the randomness in the data. The time series shows greater randomness as the points deviate from the dotted line.

- The time series has an average growth (log returns) of -0.005 (median equal to 0.0). The volatility of the returns in terms of standard deviation is 0.76. The skewness of the log differenced series is equal to -0.28, which is close to zero. This indicates a symmetric distribution, though there is a slight left skewness. The excess kurtosis of the log differenced series is equal to 2.8. This indicates a heavy tailed distribution.
- Concerning the symmetry of returns, 46.27% of the log differences are positive. The average of positive returns is 0.55, while the average of negative returns (46.27% of all returns) is -0.56. Overall, there are 38 return direction changes (57.58% of the data points)

- In the tails, 7.46% of returns fall beyond 2 standard deviations from the mean. The largest positive return is 2.19 on May 1990. Conversely, the largest decline is -2.76 (on April 1990).
- **Preliminary experiments:** Modeling the time series of first differences may improve forecasting accuracy. Experiments were conducted using a LightGBM algorithm and evaluated using SMAPE loss function. Using the original time series led to a 45.5% SMAPE. The scores using the differenced and log differenced time series are 40.42% and 40.42%, respectively.

Seasonality

Seasonality represents recurring patterns or cycles that appear at regular intervals in time series data. These are predictable fluctuations that reflect periodic influences such as monthly, quarterly, or yearly cycles. Understanding seasonal patterns is crucial for forecasting, trend analysis, and identifying anomalies. This section examines the presence, strength, and characteristics of seasonal components in the input time series.

Seasonal Line Plot (Monthly)




Figure 7: Seasonal plot of monthly values grouped by year.

- The seasonal strength is 0.64. This score ranges from 0 to 1 and values above 0.64 are considered significant. The following tests indicate that the time series is non-stationary in seasonality for a yearly period: OCSB. On the other hand, other tests (Wang-Smith-Hyndman) fail to reject the stationary null hypothesis.
- **Preliminary experiments:** Modeling yearly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (45.5% SMAPE):
 - Fourier terms: 47.12% SMAPE
 - Seasonal differencing: 44.43% SMAPE
 - Monthly time features: 49.95% SMAPE

Seasonal Sub-series Plot (Quarterly)




Figure 8: Quarterly seasonal sub-series. This plot helps to understand how the data varies within and across quarterly groups.

- Statistical analysis of quarterly data shows no evidence of systematic differences across quarters, with both means (Kruskal-Wallis test) and variances (Levene's test) being statistically similar.
- Tests for quarterly seasonal stationarity show mixed results: the OCSB test indicates presence of a seasonal unit root, while the Wang-Smith-Hyndman test suggests stationarity.
- **Preliminary experiments:** There is evidence for a quarterly seasonal pattern based on statistical tests. Besides, modeling quarterly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (45.5% SMAPE):
 - Fourier terms: 38.68% SMAPE
 - Quarterly seasonal differencing: 50.22% SMAPE
 - Quarterly time features: 45.62% SMAPE

Variance

Variance measures how data points spread around the average value in your time series. This section examines whether the variability remains stable (homoskedastic) or changes (heteroskedastic) over time. Understanding variance patterns is crucial for selecting appropriate modeling techniques, which can have a significant impact on forecasting accuracy.

Heteroskedasticity Testing




Figure 9: Time series residuals analysis based on a linear trend model. Difference in the distribution of the residuals in the first and last thirds of the series, following a Goldfeld-Quandt partition.

- No statistical evidence was found for the hypothesis that the time series is heteroskedastic, according to the White, Breusch-Pagan, and Goldfeld-Quandt tests. The residuals are based on a linear trend model.
- Variance in seasonal periods according to Levene's test
 - Quarterly groups: no differences in variance
 - Monthly groups: no differences in variance
- **Preliminary experiments:** Three variance stabilization preprocessing techniques were tested to improve the forecast accuracy of an auto-regressive LightGBM (with 45.5% SMAPE using lag-based features):
 - Log returns: 40.42% SMAPE
 - Log transformation: 43.13% SMAPE
 - Box-Cox transformation: 45.46% SMAPE