# 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 23:34

**Series Name:** M30

## Table of Contents



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

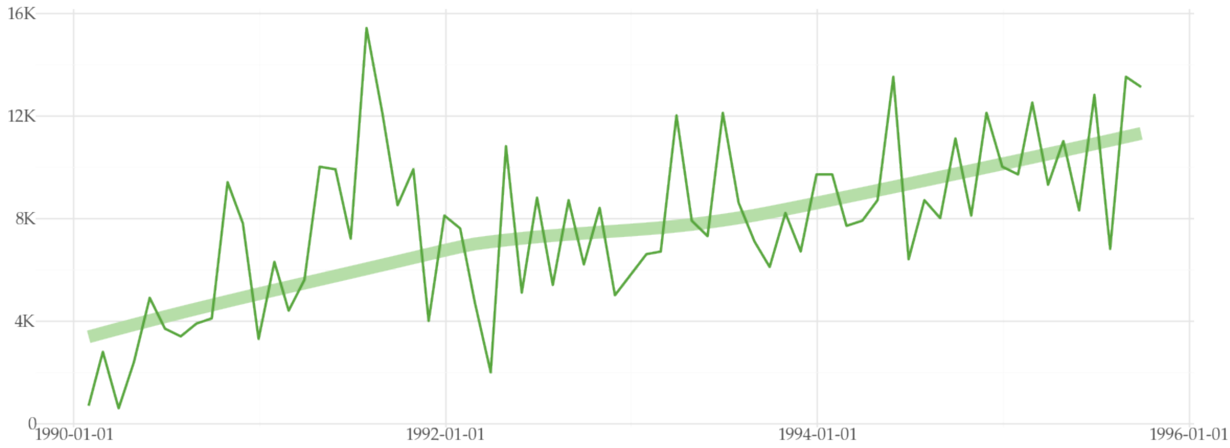


Figure 1: Time series line plot.

- A total of 69 observations spanning from January 1990 to September 1995. These are collected with a monthly sampling frequency.
- The data ranges from a minimum of 600 to a maximum of 15400, starting in 700 and ending in 13100 during the observed period. The average growth percentage per observation is 23.13% (median equal to -0.5%), with an average value of 7750.72. There are no missing values in the time series.

## Trend, Seasonality, and Residuals

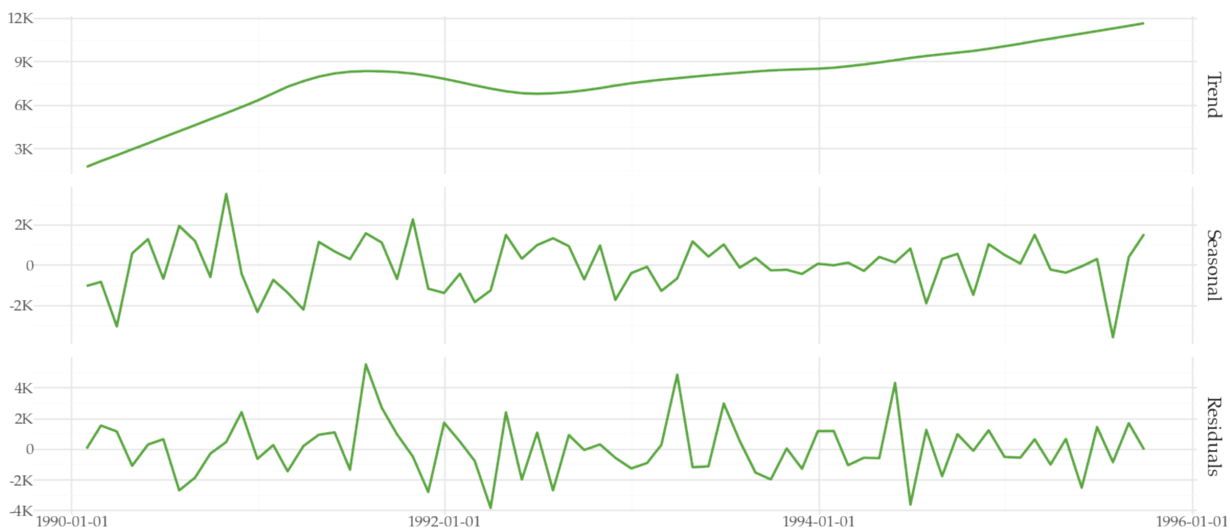


Figure 2: Seasonal, Trend, and Residuals components after decomposition on a monthly frequency using the STL (Season-Trend decomposition using LOESS) method.

- The trend strength is 0.61 (ranges from 0 to 1). The following tests indicate that the time series is non-stationary in trend or level: Augmented Dickey-Fuller. On the other hand, other tests (KPSS and Philips-Perron) fail to reject the hypothesis that the data is stationary.

- The seasonal strength is 0.37 (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: 50.72% of residuals are positive and 49.28% negative. The average magnitude of positive residuals is 1360.107 compared to -1346.867 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

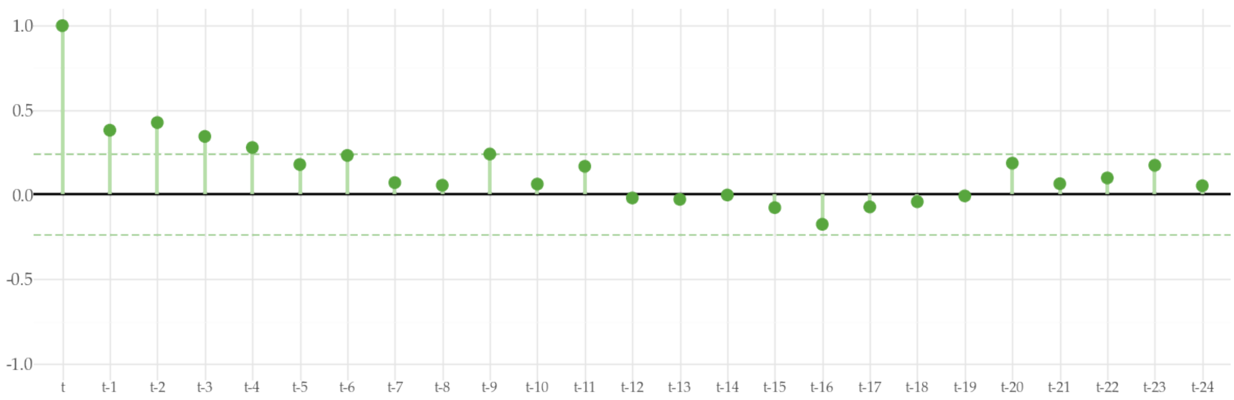


Figure 3: Auto-correlation plot up to 24 lags.

- The following lags show significant autocorrelation: t-1, t-2, t-3, and t-4. The autocorrelation is positive for all lags with a significant value.
- All lags relative to the seasonal period (t-12 and t-24) show a significant autocorrelation.

## Partial Auto-Correlation

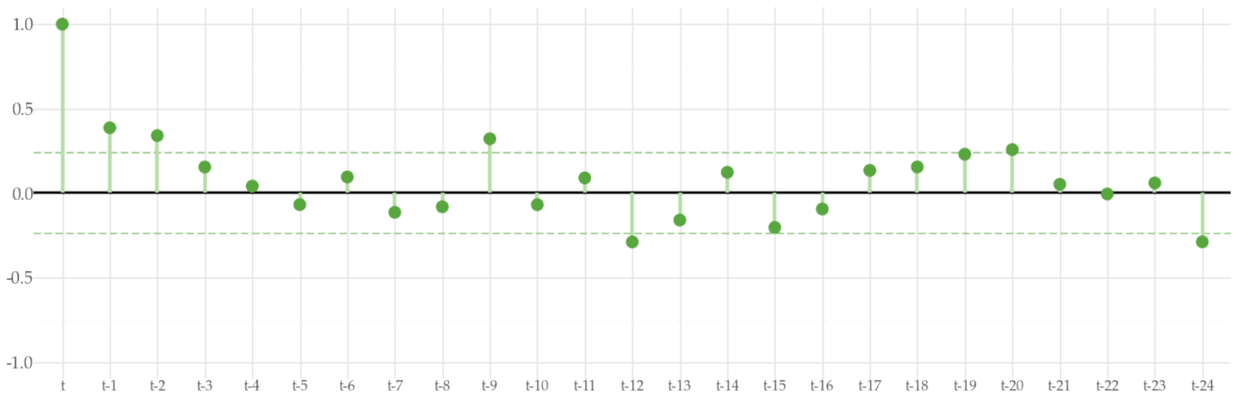

Figure 4: Partial Auto-correlation plot up to 24 lags. At each lag, the partial auto-correlation takes into account the previous correlations.

- The following lags show significant partial autocorrelation: t-1, t-2, t-9, t-12, t-20, and t-24.
- All lags relative to the seasonal period (t-12 and t-24) show a 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

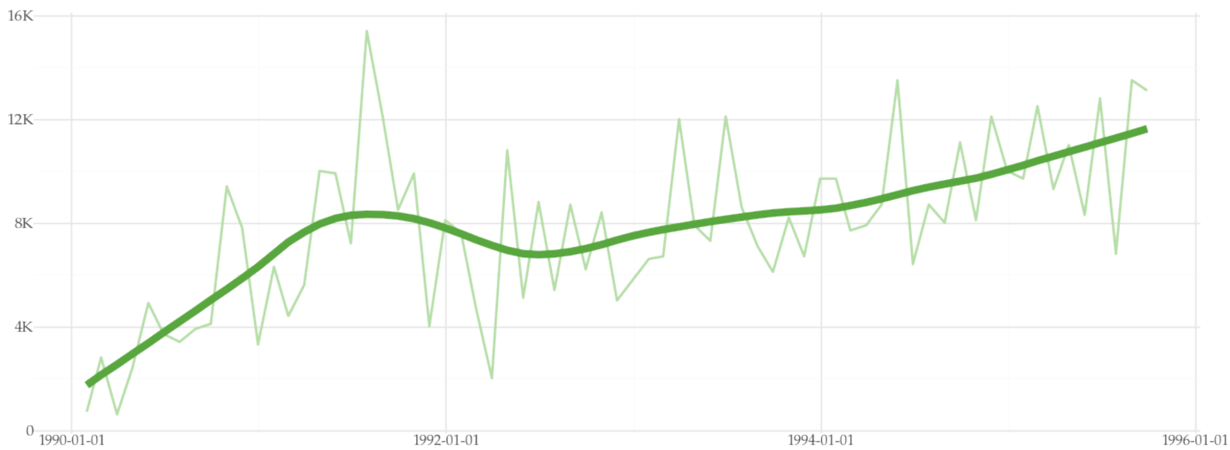


Figure 5: Time series trend plot.

- There is a moderate upward trend. The following test suggest that the trend is non-stationary (i.e. not deterministic): Augmented Dickey-Fuller. But, the tests KPSS 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. The following tests reject this hypothesis: KPSS and Augmented Dickey-Fuller. But, the test Philips-Perron suggest stationarity.
- **Preliminary experiments:** Including a trend explanatory variable which denotes the position (row id) of each observation improves 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 32.64% on the test set. Including the trend variable improved the SMAPE to 26.8%.

## Long-term Growth

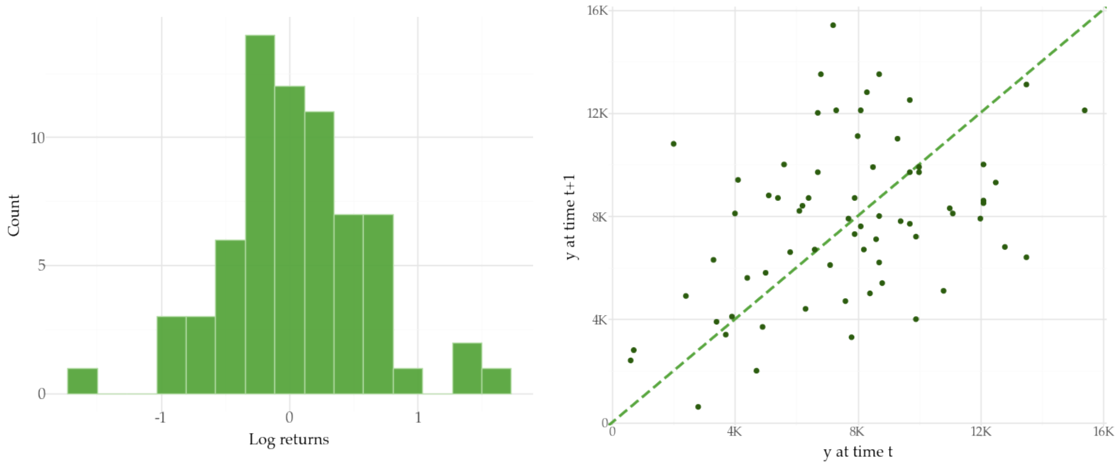


Figure 6: Distribution of log differences (left), and a Lag-plot (right). These plots help to understand how the data changes over consecutive observations. The histogram show the distribution of these changes using log returns. The lag-plot depicts the randomness in the data. The time series shows greater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

- The time series has an average growth (log returns) of 0.04 (median equal to -0.01). The volatility of the returns in terms of standard deviation is 0.56. The skewness of the log differenced series is equal to 0.3, which is close to zero. This indicates a symmetric distribution, though there is a slight right skewness. The excess kurtosis of the log differenced series is equal to 0.99. This value is similar to that found from data following a Gaussian distribution.

- Concerning the symmetry of returns, 48.53% of the log differences are positive. The average of positive returns is 0.48, while the average of negative returns (50.0% of all returns) is -0.38. Overall, there are 47 return direction changes (70.15% of the data points)

- In the tails, 5.88% of returns fall beyond 2 standard deviations from the mean. The largest positive return is 1.69 on March 1992. Conversely, the largest decline is -1.54 (on February 1990).

- **Preliminary experiments:** Modeling the time series of first differences does not seem to improve forecasting accuracy. Experiments were conducted using a LightGBM algorithm and evaluated using SMAPE loss function. Using the original time series led to a 32.64% SMAPE. The scores using the differenced and log differenced time series are 42.46% and 47.2%, 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)

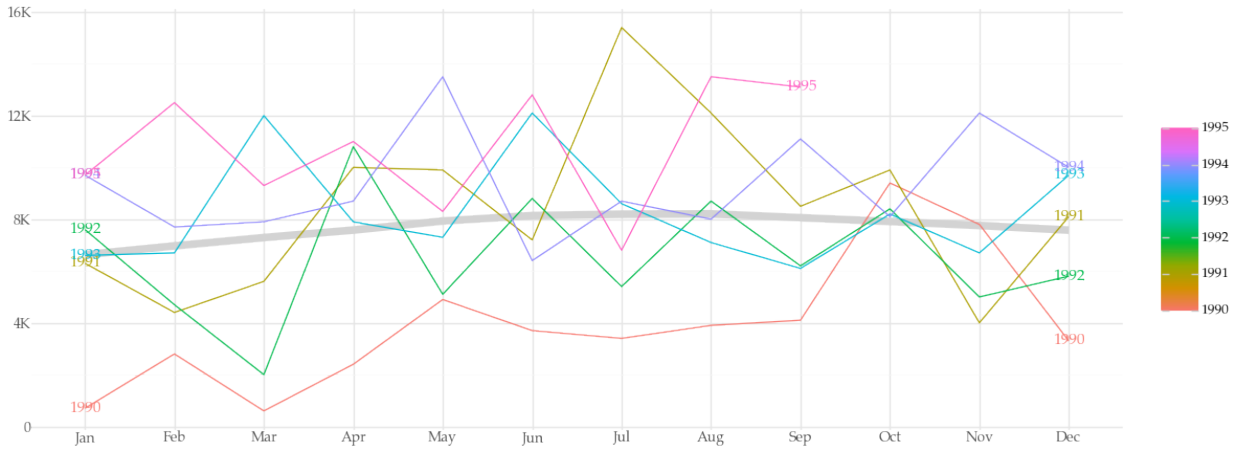


Figure 7: Seasonal plot of monthly values grouped by year.

- The seasonal strength is 0.37. 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 (32.64% SMAPE):
  - Fourier terms: 33.74% SMAPE
  - Seasonal differencing: 28.6% SMAPE
  - Monthly time features: 32.64% SMAPE

## Seasonal Sub-series Plot (Quarterly)

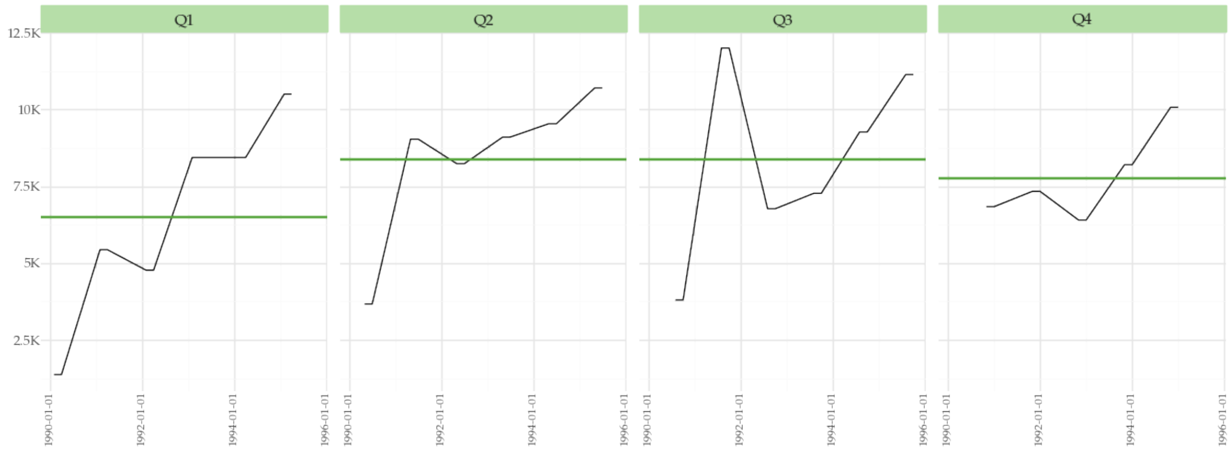


Figure 8: Quarterly seasonal sub-series. This plot helps to understand how the data varies within and across quarterly groups.

- 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.

- All hypothesis tests carried out (Wang-Smith-Hyndman and OCSB) indicate that the time series is stationary in seasonality.

- **Preliminary experiments:** Statistical tests show no evidence for a quarterly seasonal pattern. Yet, modeling quarterly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (32.64% SMAPE):
    - Fourier terms: 32.45% SMAPE
    - Quarterly seasonal differencing: 34.79% SMAPE
    - Quarterly time features: 32.64% 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

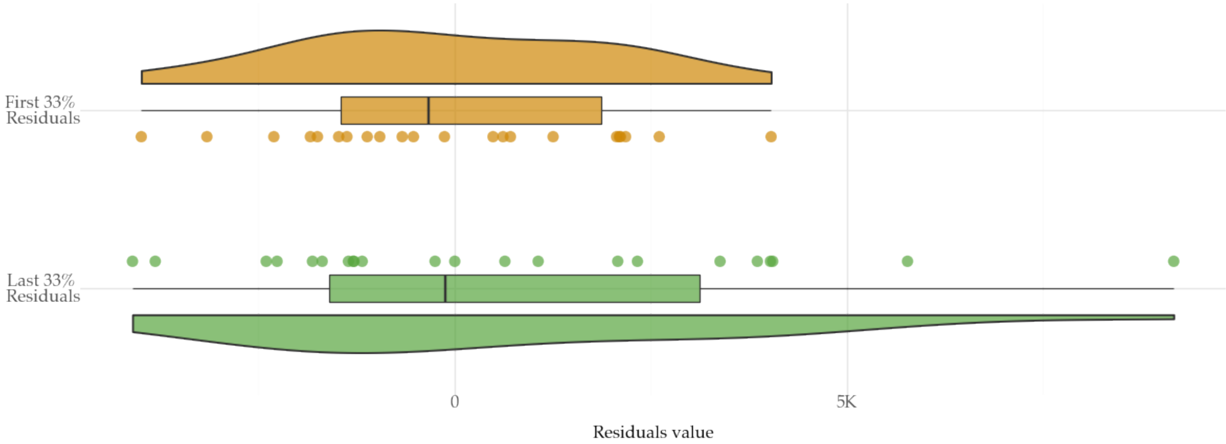


Figure 9: Time series residuals analysis based on a linear trend model. Difference in the distribution of the residuals in the first and last thirds of the series, following a Goldfeld-Quand partition.

- The following tests suggest that the time series is heteroskedastic: White and Breusch-Pagan. But, other tests (Goldfeld-Quandt) fail to reject the hypothesis that the time series has a constant variance. 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 32.64% SMAPE using lag-based features):
  - Log returns: 47.2% SMAPE
  - Log transformation: 40.12% SMAPE
  - Box-Cox transformation: 34.7% SMAPE