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 23:12

Series Name: M900

Table of Contents

Other aspects were explored but omitted from the final report:

Variance

Hypothesis testing suggests that the time series has constant variance (homoskedasticity). In preliminary tests, common transformations for variance stabilization did not improve forecasting accuracy

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

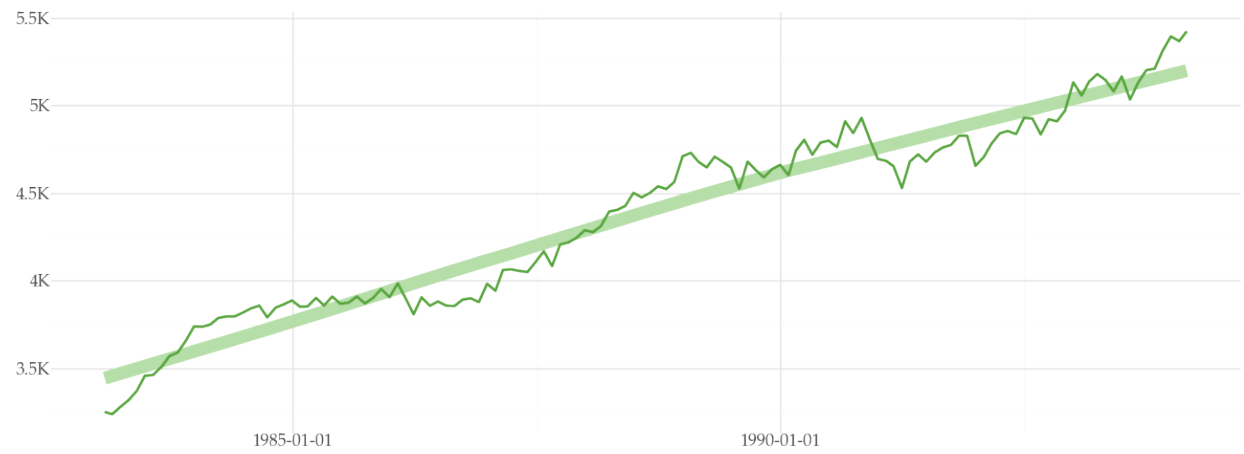

Figure 1: Time series line plot.

- A total of 134 observations spanning from January 1983 to February 1994. These are collected with a monthly sampling frequency.
- The data ranges from a minimum of 3236.96 to a maximum of 5423.5, starting in 3249.6 and ending in 5423.5 during the observed period. The average growth percentage per observation is 0.4% (median equal to 0.53%), with an average value of 4364.34. There are no missing values in the time series.

Trend, Seasonality, and Residuals

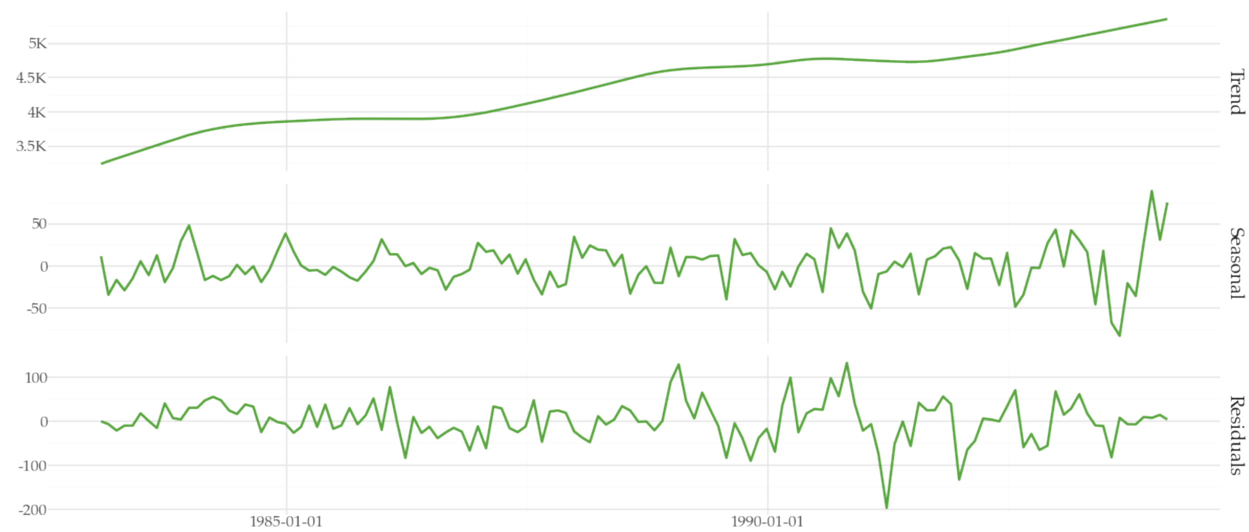

Figure 2: Seasonal, Trend, and Residuals components after decomposition on a monthly frequency using the STL (Season-Trend decomposition using LOESS) method.

- The trend strength is 0.99 (ranges from 0 to 1). The following tests indicate that the time series is non-stationary in trend or level: Augmented Dickey-Fuller and

Philips-Perron. On the other hand, other tests (KPSS) fail to reject the hypothesis that the data is stationary.

- The seasonal strength is 0.32 (ranges from 0 to 1). All hypothesis tests carried out (Wang-Smith-Hyndman and OCSB) indicate that the time series is stationary in seasonality.

- The STL decomposition residuals show balanced behavior: 50.0% of residuals are positive and 50.0% negative. The average magnitude of positive residuals is 34.653 compared to -32.518 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

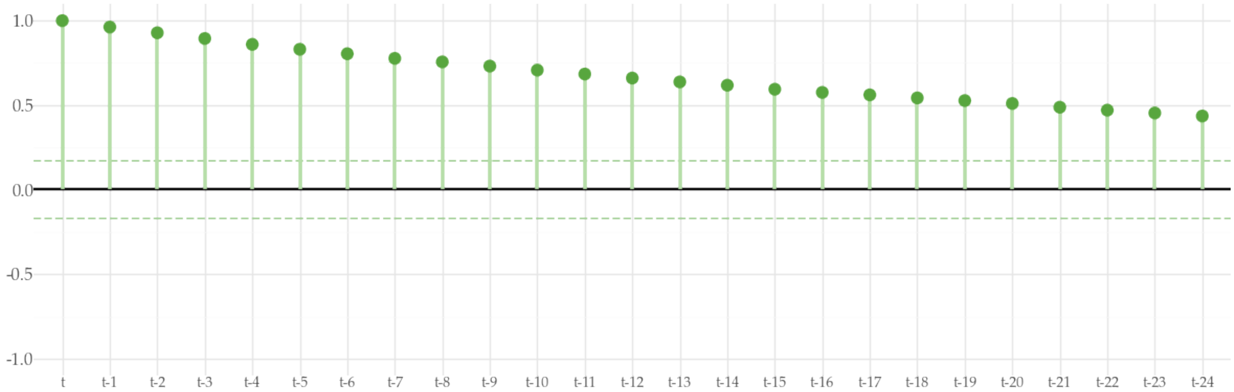

Figure 3: Auto-correlation plot up to 24 lags.

- The following lags show significant autocorrelation: t-1, t-2, t-3, t-4, t-5, t-6, t-7, t-8, t-9, t-10, t-11, t-12, t-13, t-14, t-15, t-16, t-17, t-18, t-19, t-20, t-21, t-22, t-23, and t-24. The autocorrelation is positive for all lags with a significant value.

- None of the lags relative to the seasonal period (t-12 and t-24) show any significant autocorrelation.

Partial Auto-Correlation

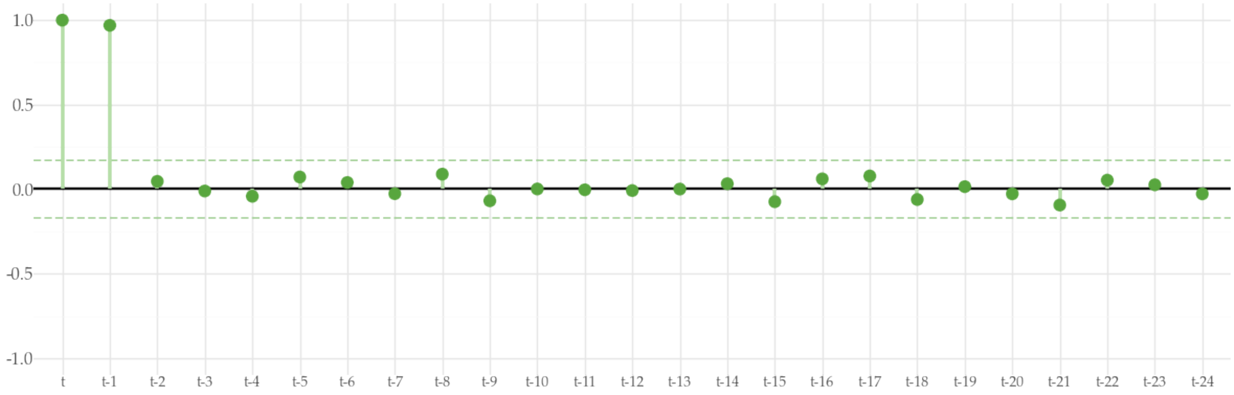

Figure 4: Partial Auto-correlation plot up to 24 lags. At each lag, the partial auto-correlation takes into account the previous correlations.

- The following lags show significant partial autocorrelation: t-1.
- 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

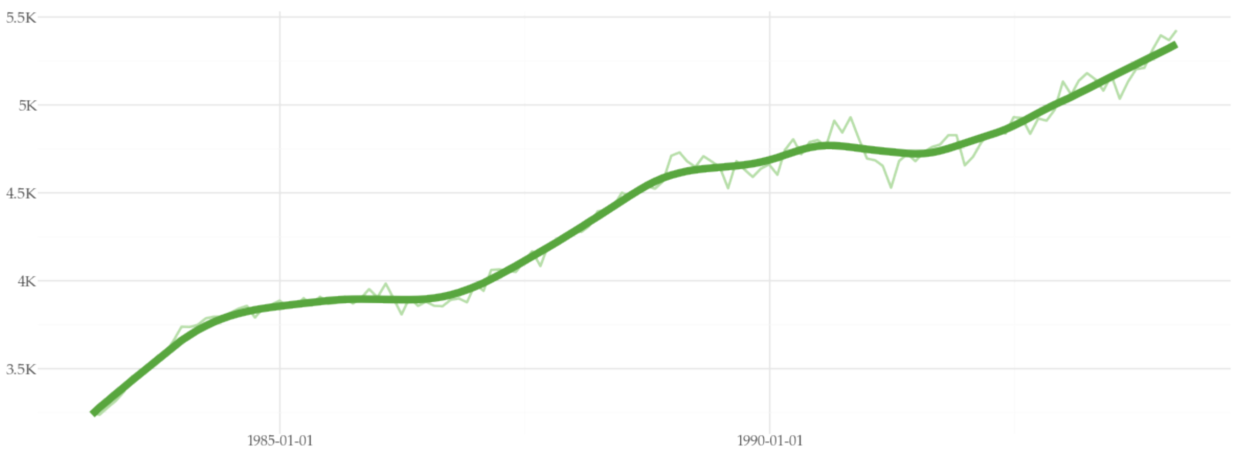

Figure 5: Time series trend plot.

- There is a strong upward trend. The following tests suggest that the trend is non-stationary (i.e. not deterministic): Augmented Dickey-Fuller and Philips-Perron.

But, the test KPSS 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 reject this hypothesis.

- **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 4.3% on the test set. Including the trend variable improved the SMAPE to 3.4%.

Long-term Growth

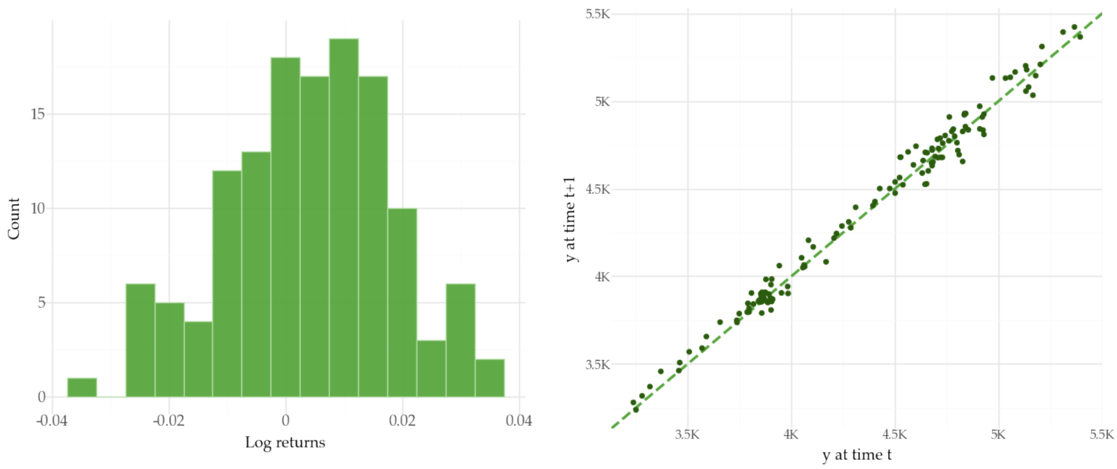

Figure 6: Distribution of log differences (left), and a Lag-plot (right). These plots help to understand how the data changes over consecutive observations. The histogram show the distribution of these changes using log returns. The lag-plot depicts the randomness in the data. The time series shows greater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

- The time series has an average growth (log returns) of 0.004 (median equal to 0.01). The volatility of the returns in terms of standard deviation is 0.01. The skewness of the log differenced series is equal to -0.22, 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 -0.16. This value is similar to that found from data following a Gaussian distribution.

- Concerning the symmetry of returns, 63.16% of the log differences are positive. The average of positive returns is 0.01, while the average of negative returns (36.84% of all returns) is -0.01. Overall, there are 69 return direction changes (52.27% of the data points)

- In the tails, 4.51% of returns fall beyond 2 standard deviations from the mean. The largest positive return is 0.03 on July 1989. Conversely, the largest decline is -0.04 (on November 1991).

- **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 4.3% SMAPE. The scores using the differenced and log differenced time series are 1.14% and 1.44%, 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)

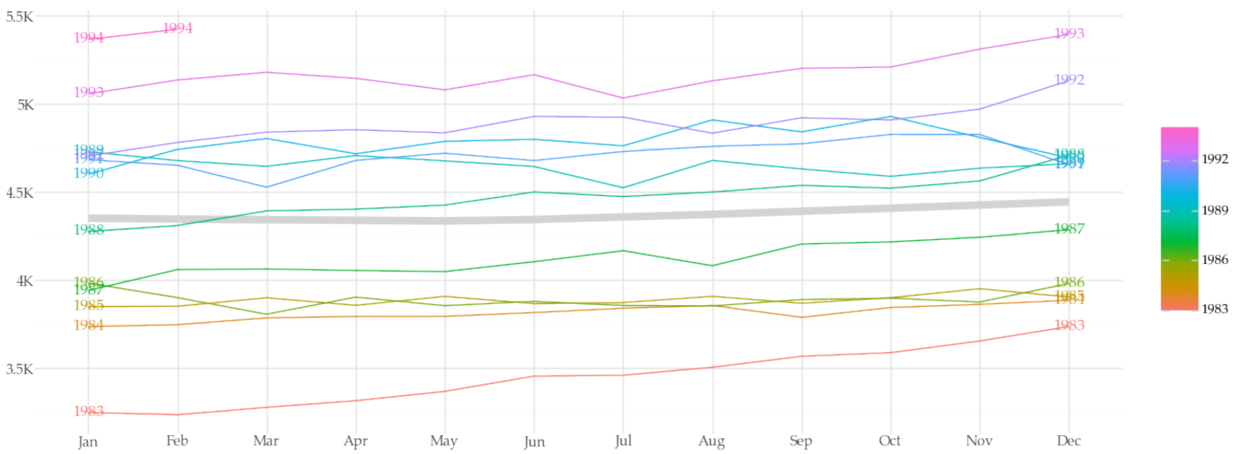

Figure 7: Seasonal plot of monthly values grouped by year.

- The seasonal strength is 0.32. This score ranges from 0 to 1 and values above 0.64 are considered significant. All hypothesis tests carried out (Wang-Smith-Hyndman and OCSB) indicate that the time series is stationary in yearly seasonality.

- **Preliminary experiments:** Modeling yearly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (4.3% SMAPE):
    - Fourier terms: 2.21% SMAPE
    - Seasonal differencing: 3.2% SMAPE
    - Monthly time features: 3.99% SMAPE

Seasonal Sub-series Plot (Quarterly)

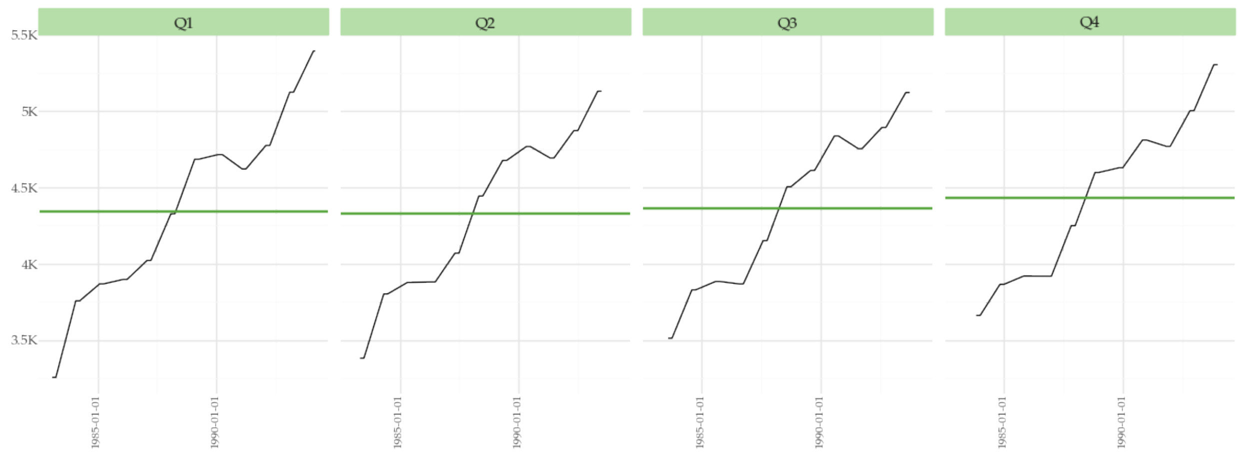

Figure 8: Quarterly seasonal sub-series. This plot helps to understand how the data varies within and across quarterly groups.

- 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 (4.3% SMAPE):
    - Fourier terms: 3.99% SMAPE
    - Quarterly seasonal differencing: 2.63% SMAPE
    - Quarterly time features: 4.63% SMAPE

Change Detection

Change points denote significant shifts in the underlying distribution of time series. These structural changes can manifest as sudden shifts in level, trend, variance, or seasonal patterns. Detecting and understanding these points is crucial as they often indicate important events or regime changes that affect modeling decisions. This section identifies potential change points and assesses their impact on the overall analysis strategy.

Change Points

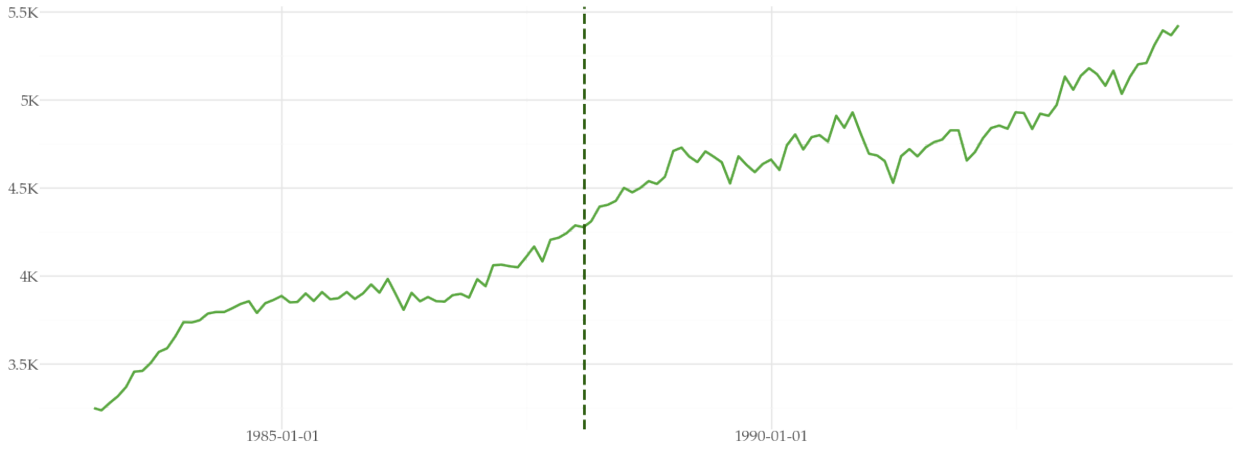

Figure 9: Time series plot with marked change points according to the PELT method.

- A single change point was found in the time series.
- The change point was found at January 1988 where the time series shows an increasing level.

Effect on Model Parameters

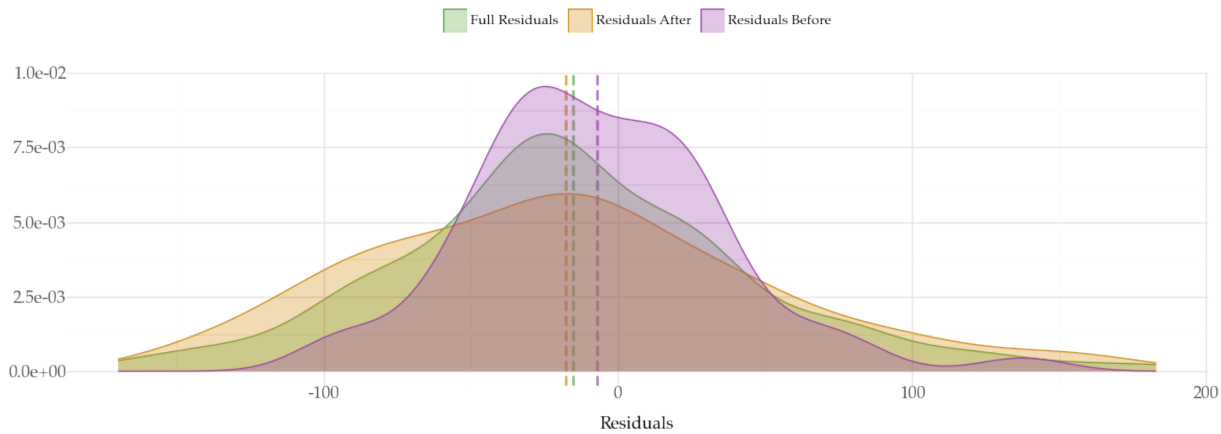

Figure 10: Distribution of the residuals of an ARIMA model before and after the first detected change point. The plot compares three kernel density estimates: residuals from the pre-change model, post-change model, and full series model. This comparison helps assess whether the structural break affects model adequacy and error distribution properties.

- A Chow test was conducted using an ARIMA(2, 1, 2) model. The test rejects the null hypothesis of parameter stability. This suggests that the ARIMA parameters are significantly different before and after the first detected change point, indicating a structural change in the underlying process.
- **Preliminary experiments:** Adding a step intervention at the change point did not affect the model performance. The baseline SMAPE of 4.3% remained the same when including the intervention (4.3%).