Time Series Viewer

Time Series Analysis

Loading, transforming, calculating, comparing and exporting data

English

Wolfgang Boehme

wolfgang.boehme@economic-modeling.de

www.economic-modeling.de

8 November 2026

The documentation on this website was developed with the support of the AI Claude (Anthropic).

Table of Contents

1 Introduction **2**

2 Loading Data **2**
2.1 CSV Files 2
2.2 Example Data 2

3 Loaded Series **3**
3.1 Transformations 3
3.2 Formula Calculator 4

4 Charts **5**
4.1 Individual Charts 5
4.2 Comparison Chart and Comparison of HP Trends 5
4.3 The HP Trend 5

5 Export **5**

1 Introduction

The *Time Series Viewer* loads any number of time series, displays them – optionally transformed – individually and together in comparison charts, calculates new series from existing ones on request (formula calculator) and exports a selection as a CSV file. It provides a quick visual overview before more elaborate analyses (time series analyser, time series modelling, periodisation analyser) are started.

The tool runs entirely in the browser as a single HTML file. Only the chart library Chart.js is loaded once from the internet when the page is opened; all other functions work locally, and no data are transmitted to a server.

The interface is divided into five sections, which appear as soon as at least one series has been loaded:

1. Load data (CSV upload or example data)
2. Loaded series (table with transformation, selection and formula calculator)
3. Individual charts
4. Comparison chart and comparison of HP trends
5. Export

2 Loading Data

2.1 CSV Files

"Upload CSV file(s)" allows one or more files to be selected at once. The tool expects the wide format that the data import assistant and the periodisation analyser also use:

```
year,SeriesA,SeriesB,...
1991,100.0,100.0,...
1992,101.4,105.0,...
```

- The header row must contain a column `year` (case-insensitive); every other column becomes a separate series.
- The delimiter is the comma, the decimal separator the point.
- Missing or non-numeric values are allowed: the series then has a gap at that point. A series needs at least two valid values.
- A byte order mark (BOM) at the beginning of the file is removed automatically.

Loading repeatedly does not replace series but extends the list. If a new series has the same name as an existing one, it is numbered (e. g. "GDP (2)").

2.2 Example Data

The button "Load example data" loads three **synthetic** annual series 1991–2024 (not official data): a real GDP index (1991 = 100), a consumer price index (1991 = 100) and the number of persons employed in millions. The slumps of 2009 (financial crisis) and 2020 (pandemic) as well as the inflation surge of 2022/23 are built into the data so that

transformations and the HP trend show visible patterns. All three series are marked for the charts immediately.

Table 1: Example data (extract)

Year	GDP_real_index (Z1)	Consumer_price_index (Z2)	Employment_millions (Z3)
1991	100.0	100.0	38.70
2000	107.0	124.0	38.87
2008	117.6	138.4	39.95
2009	111.1	141.1	39.99
2010	115.6	142.4	40.58
2019	135.4	162.7	43.61
2020	130.0	164.0	43.26
2021	134.5	169.1	43.31
2022	136.5	180.9	43.35
2024	135.8	195.8	43.26

3 Loaded Series

All series appear in a table:

- **Symbol** (`Z1`, `Z2`, ...): consecutive label for the formula calculator. The symbol remains stable even if other series are deleted; clicking it inserts it into the formula.
- **Series**: name; for calculated series the formula is shown as well.
- **Period**: first and last year.
- **Transformation**: selectable per series (section 3.1).
- **Compare**: marks the series for the individual charts and the comparison charts.
- **Export**: marks the series for the CSV export (at most 10 at a time).
- **Bin**: removes the series from the current session.

3.1 Transformations

Original values	y_t
Log returns	$r_t = \ln(y_t/y_{t-1})$, only for $y_{t-1}, y_t > 0$
Rate of change (%)	$g_t = (y_t - y_{t-1})/y_{t-1} \cdot 100$, only for $y_{t-1} \neq 0$
Differences (absolute)	$\Delta y_t = y_t - y_{t-1}$

The last three transformations lose the first observation of the series; points at which a transformation is not defined are skipped. On loading, **log returns** are preset; if a series contains non-positive values, the rate of change is chosen instead.

Example (example data, GDP 2009):

$$\begin{aligned} r_{2009} &= \ln(111.1/117.6) = -0.0569, \\ g_{2009} &= (111.1 - 117.6)/117.6 \cdot 100 = -5.53\,\%, \\ \Delta y_{2009} &= 111.1 - 117.6 = -6.5. \end{aligned}$$

For small changes $r_t \approx g_t/100$; for larger movements the log return and the rate of change differ visibly. For the consumer price index, the rate of change in 2022 shows the inflation surge: $(180.9 - 169.1)/169.1 \cdot 100 = 6.98\,\%$.

3.2 Formula Calculator

Below the table, new series can be calculated from existing ones. The following are allowed:

- series via their symbol `Z1`, `Z2`, ...
- numbers (decimal point, or decimal comma between digits, e. g. `0,5`)
- operators `+ - * /`, power `^`, brackets, unary minus
- placeholders `a1`, `a2`, ... for coefficients – an input field appears for each (initial value 1)
- functions `ln()`, `exp()`, `sqrt()`, `abs()`

The calculation always uses the **original values**, year by year. A year only receives a value if all series used have a value there and the result is finite (no division by 0). While typing, the tool shows whether the formula is valid, how many years are calculated and the first and last values; errors (unknown symbol, missing bracket, invalid character ...) are reported in plain language. The formula is evaluated by a dedicated parser without `eval`.

The new series receives the next free symbol and the transformation "Original values" and is marked for the charts. If no name is entered, the formula (with the coefficients substituted) is used as the name.

Examples with the example data:

Table 2: Formula examples

Formula	Meaning	2009	2024
`Z1*Z2/100`	nominal GDP index (1991 = 100)	156.76	265.90
`Z1/Z3*a1`, $a_1 = 38.7$	labour productivity (GDP per person employed, 1991 = 100)	107.52	121.49

Calculation for 2009: $111.1 \cdot 141.1/100 = 156.76$ and $111.1/39.99 \cdot 38.7 = 107.52$. The coefficient $a_1 = 38.7$ (persons employed in 1991) normalises productivity to 1991 = 100. Productivity falls from 113.92 to 107.52 in 2009 because employment did not decline despite the slump in GDP.

4 Charts

4.1 Individual Charts

For each series marked under “Compare”, a separate chart with the selected transformation appears. If the transformed series has at least five values, the HP trend is additionally shown as a dashed line. Each chart can be exported as a PNG (with a white background).

4.2 Comparison Chart and Comparison of HP Trends

The comparison chart shows all marked series together, each with its own transformation and colour on a common y-axis. With different units this display is of limited informative value; a uniform transformation (log returns or rate of change) makes series comparable. Below it, a second chart shows only the HP trends of the same series in the same colours – this makes it easier to compare long-term developments.

4.3 The HP Trend

The Hodrick-Prescott filter determines a smooth trend τ_t of the (transformed) series y_t:

$$\min_{\{\tau_t\}} \sum_{t=1}^{T} (y_t - \tau_t)^2 \;+\; \lambda \sum_{t=2}^{T-1} \left[(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1}) \right]^2 \qquad (1)$$

The smoothing parameter is fixed at the value $\lambda = 100$ customary for annual data. The solution follows from the linear system $(\mathbf{I} + \lambda \mathbf{D}^\top \mathbf{D})\,\boldsymbol{\tau} = \mathbf{y}$ ($\mathbf{D}$: matrix of second differences), which is solved in the browser by Cholesky decomposition. The trend is always calculated on the displayed, i. e. transformed, series.

With the example data and the log-return transformation, the HP trend of GDP is about 0.008 in the 1990s (roughly 0.8 % growth per year), rises to about 0.014 by the mid-2010s and then falls to about 0.006. The slumps of 2009 and 2020 appear as short-term deviations around this trend.

5 Export

Section 5 exports all series marked under “Export” (at most 10) as the CSV file `time_series_viewer_export_YYYY-MM-DD.csv`.

- The export always contains the **original values**, regardless of the transformation selected for display.
- The file contains **all years** of the selected series (union). If a series has no value in a year, the field remains empty.
- Format: wide format with the header row `year,Series1,Series2,...` – the same structure as for loading. The file can therefore be loaded again directly in the viewer, in gap filling or in the periodisation analyser.
- A counter shows the number of selected series and changes colour at 10.

Typical workflow

1. Load CSV files (or the example data).
2. Select the transformation per series and mark the series of interest for comparison.
3. If required, create derived series with the formula calculator.
4. Inspect individual and comparison charts as well as HP trends; save charts as PNG.
5. Mark up to 10 series for export and save them as CSV – if there are gaps, close them afterwards in gap filling.