Universität Zürich[UZH]

Department of Economics

University of Zurich

FX Data Visualization:

Web-Based Triangular Analysis of Exchange Rates

Working Paper

Submitted by

Felix Vadakumchery

Matriculation Number: 22-703-540

Supervised by

Prof. Dr. Christian Ewerhart

University of Zurich, Department of Economics

Table of Contents

Table of Figures

List of Tables

Abstract

Three currencies define three bilateral exchange rates and tracking how all three move in relation to each other simultaneously is not straightforward with conventional time series charts. This thesis develops a web-based visualization tool that addresses this by projecting any selected currency triple onto a single two-dimensional trajectory, exploiting the mathematical structure of the “triangular no-arbitrage condition” to do so without loss of information. The transformation pipeline and a basis vector projection, reduces three time series to one moving point per time period, where the direction and magnitude of each step encode which currency drove the change.

The tool, deployed at scbt-course.net, supports 30 currencies and allows the user to configure the currency triple, date range, and granularity interactively. Exchange rate data is sourced from the European Central Bank, which publishes daily reference rates for 29 currencies against the euro back to January 1999. A list of macroeconomic events, covering monetary policy decisions by major central banks, geopolitical shocks, and financial crises, can be overlaid on the chart to connect visible features of the trajectory to their economic causes.

An example analysis of the EUR/CHF/USD triple from December 2014 to June 2026 illustrates how the visualization makes monetary policy turning points, geopolitical shocks, and longer-term trends readable in a single chart. The dominant macro trend over the period is a sustained strengthening of the Swiss franc against both the euro and the dollar, punctuated by episodes driven by central bank decisions and external shocks, covering events from the SNB floor removal in January 2015 to the Iran war “safe-haven peak” in early 2026. Events that would require cross-referencing three separate bilateral graphs become readable at one quick glance.

1 Motivation and Goal of the Tool[1]

Exchange rates are central quantities in international macroeconomics and financial markets. Of particular interest are trios of currencies (e.g. EUR-CHF-USD), for which so-called triangular arbitrage relationships arise. The product of the three exchange rates along a closed currency loop should, abstracting from spreads and market inefficiencies, be approximately equal to one. The standard way of presenting exchange rates are time-series data sets plotted as individual graphs which makes it difficult to see how three currencies move in relation to each other simultaneously. When the Swiss franc strengthens against the euro, does it strengthen equally against the dollar, or is the move asymmetric? When the dollar weakens broadly, how is that distributed across the other two legs of the triangle?

The goal of this thesis is to design and implement a web-based application, written in English, that allows the user to select:

- a triple of currencies,
- a time period,
- and a granularity (daily, monthly, etc.)

which allows to generate a visualization of the dynamics of the currency triple and download the resulting charts.

The visualization is based on a mathematically well-defined transformation of the exchange rate data, which I will describe in more detail in the next chapter. This transformation, together with the triangular (no) arbitrage relationships of the exchange rates, allows us to represent three variables in two dimensions rather than three.

Example use cases of the tool could include the following:

- Students who can observe, for instance, how the removal of the CHF-EUR floor affected the currencies involved,
- Researchers who want to examine the influence of different policy regimes across central banks more closely,

The following chapter introduces the theoretical foundations on which both the visualization and the backend of the tool are built.

[1] This Chapter draws on the project proposal provided by Prof. Ewerhart

2 Theoretical Fundamentals[2]

2.1 Triangular (no)-Arbitrage Relationship/Condition

Consider a currency triple (A, B, C) with three exchange rate series, denoted u_t, v_t, and w_t, where each series represents the rate between one pair of currencies at time t. Concretely, in the case of EUR, CHF, and USD:

$$u_t = \frac{\text{CHF}}{\text{EUR}}, v_t = \frac{\text{USD}}{\text{CHF}}, w_t = \frac{\text{EUR}}{\text{USD}}$$

So u_t gives the price of one euro in Swiss francs, v_t gives the price of one franc in US dollars, and w_t gives the price of one dollar back in euros. Together, these three rates form a closed loop.
In a frictionless market, one without transaction costs, bid-ask spreads, or delays, an investor who converts one unit of currency A into B, then B into C, and finally C back into A should end up with exactly one unit of A again. This means that at any point in time t the product of the three rates satisfies:

$$u_t \cdot v_t \cdot w_t = 1$$

This is the "triangular no-arbitrage condition". If the product were to deviate persistently and significantly from one, a trader could exploit the discrepancy by cycling through the three currencies and locking in a risk-free profit. Competitive pressure in liquid markets tends to eliminate such opportunities quickly, which is why the condition holds as a close approximation in practice (Akram et al., 2008).

2.2 Normalization

Real market data rarely satisfies $u_t \cdot v_t \cdot w_t = 1$ exactly. Small deviations arise from bid-ask spreads, fixing-time differences, and other market frictions. Before applying the logarithmic transformation described in the next subchapter, it is useful to symmetrize these deviations around one. This is the purpose of the normalization step.

At each point in time t, the geometric mean of the three rates is computed as:

$$g_t = (u_t \cdot v_t \cdot w_t)^{1/3}$$

If the no-arbitrage condition held perfectly, g_t would equal exactly one at every observation. In practice it hovers close to one and dividing each rate by g_t forces the product back to one exactly. The normalized series are defined as:

$$u'_t = \frac{u_t}{g_t}, v'_t = \frac{v_t}{g_t}, w'_t = \frac{w_t}{g_t}$$

By construction, the normalized rates satisfy:

$$u'_t \cdot v'_t \cdot w'_t = \frac{u_t \cdot v_t \cdot w_t}{g_t^3} = \frac{u_t \cdot v_t \cdot w_t}{u_t \cdot v_t \cdot w_t} = 1$$

The geometric mean is the natural choice here rather than the arithmetic mean, because exchange rates are multiplicative by nature. A rate of 1,20 and a rate of 0,80 are symmetric

[2] This chapter draws on the project proposal provided by Prof. Ewerhart

deviations from 1,00 in a multiplicative sense, meaning their geometric mean is exactly 1,00, while their arithmetic mean is 1,00 only by coincidence. Using the geometric mean therefore treats upward and downward deviations symmetrically, which is the correct behaviour for ratio-scale data of this kind (Aitchison, 1982).

2.3 Log-Transformation

With the normalized rates u_t', v_t', and w_t' in hand, the next step is to apply a logarithmic transformation to each series. The log-transformed values are defined as:

$$u_t'' = \ln{(u_t')}, v_t'' = \ln{(v_t')}, w_t'' = \ln{(w_t')}$$

The key property that motivates this step follows directly from the normalization. Since $u_t' \cdot v_t' \cdot w_t' = 1$, taking the natural logarithm of both sides gives:

$$\ln{(u_t' \cdot v_t' \cdot w_t')} = \ln{(1)} = 0$$

And since the logarithm of a product equals the sum of the logarithms:

$$u_t'' + v_t'' + w_t'' = 0$$

This constraint means that at every point in time, the three log-transformed values sum to zero. Geometrically, this places every observation (u_t'', v_t'', w_t'') on the plane defined by $x + y + z = 0$ in three-dimensional space, a two-dimensional surface passing through the origin. It is precisely this constraint that makes the projection into a two-dimensional chart possible, as there are effectively only two degrees of freedom in the data at any given time. Beyond this geometric consequence, the logarithmic transformation is well motivated on its own terms. In financial economics, log returns are the standard way to measure relative price changes, because they are symmetric and additive over time. A rate that doubles from 1,00 to 2,00 and a rate that halves from 1,00 to 0,50 are equal and opposite movements in log space, with $\ln{(2)} \approx 0.693$ and $\ln{(0.5)} \approx -0.693$ respectively, whereas in levels the two moves are asymmetric (Campbell et al., 1997).

2.4 Projection into 2D-Plane

The logarithmic transformation established that every observation (u_t'', v_t'', w_t'') lies on the plane $u'' + v'' + w'' = 0$ in three-dimensional space. This plane is two-dimensional, meaning any point on it can be described with just two coordinates. The final step is to find a clean and geometrically meaningful way to map these points onto a standard 2D coordinate system so they can be plotted.

The mapping is constructed by assigning each of the three currency axes a unit vector in the 2D plane, placed at angles of 0°, 120°, and 240° from one another:

$$\mathbf{e}_u = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \mathbf{e}_v = \begin{pmatrix} -\frac{1}{2} \\ \frac{\sqrt{3}}{2} \end{pmatrix}, \mathbf{e}_w = \begin{pmatrix} -\frac{1}{2} \\ -\frac{\sqrt{3}}{2} \end{pmatrix}$$

These three vectors are evenly spaced around the origin, reflecting the symmetric role that each currency plays in the triple. A point (u_t'', v_t'', w_t'') is then projected onto the 2D plane by taking a weighted sum of these basis vectors:

$$\binom{x_t}{y_t} = u_t'' \cdot \mathbf{e}_u + v_t'' \cdot \mathbf{e}_v + w_t'' \cdot \mathbf{e}_w$$

Writing this out component by component gives the explicit formulas used in the backend:

$$x_t = u_t'' - \frac{1}{2} v_t'' - \frac{1}{2} w_t''$$
$$y_t = \frac{\sqrt{3}}{2} (v_t'' - w_t'')$$

The resulting coordinates $(x_{t'} y_t)$ are what gets plotted. To make the time dimension visible, consecutive observations are connected by arrows, so the chart reads as a trajectory through the plane rather than a static scatter plot. The direction of each arrow shows the direction of change between two periods, and the length of the arrow reflects how much the exchange rate configuration shifted.

The choice of 120° spacing is not arbitrary. It ensures that the three currency axes contribute symmetrically to the plot, so that no single currency is visually privileged over the others. A movement along the $\mathbf{e}_u$ direction, for instance, corresponds to a strengthening of the EUR/CHF rate relative to the other two, while a movement perpendicular to $\mathbf{e}_u$ reflects changes in the CHF/USD and USD/EUR rates that leave EUR/CHF unchanged.

3 Data

3.1 Exchange Rates

The exchange rate data used in this tool is sourced from the European Central Bank's (ECB) Data Portal, which publishes official daily reference rates for 29 currencies, all quoted against the euro as the base currency (European Central Bank, 2026a). Therefore, 30 available currencies to choose from in the tool (European Central Bank, 2026b). The rates are mid-point quotes, meaning they carry no bid-ask spread, which has a direct consequence for the no-arbitrage product. Since u_t, v_t, and w_t are derived from the same official fixing rather than from independent market quotes, the product $u_t \cdot v_t \cdot w_t$ stays within 0.01% of one on any given trading day. The data is available from January 1999 onwards, covering more than two decades of exchange rate history.

At its core, what the backend receives and processes are time series, namely sequences of exchange rate observations indexed by date. Each of the three rates u_t, v_t, and w_t constitutes one such series, and the entire mathematical pipeline operates on these three series in parallel, transforming them step by step into the two-dimensional coordinates that are eventually passed to the frontend for plotting.

The ECB publishes reference rates on trading days only, typically around 16:00 CET, based on a daily concertation procedure between central banks across Europe that normally takes place around 14:10 CET (European Central Bank, 2026c). Weekends and fixed ECB holidays, including New Year's Day, Labour Day, and Christmas, are absent from the feed entirely (European Central Bank, 2026c). When a user selects a date range whose boundary falls on a non-trading day, the tool uses the nearest available trading day instead. This behaviour is flagged transparently in the data notes displayed below the chart, so the user is always aware of any such adjustment. Isolated gaps within a series can also occur when a country-specific market closure affects one currency but not others. In these cases, the last known rate is carried forward to the missing trading day. Of the three rates that make up the currency triple, only two are fetched directly from the ECB API when EUR is one of the selected currencies. Taking the default triple of EUR, CHF, and USD as an example, the available series are the CHF/EUR rate and the USD/EUR rate, both expressed as the number of foreign currency units per euro. The third rate, CHF/USD, is derived internally as their ratio:

$$w_t = \frac{u_t}{v_t} = \frac{\text{CHF/EUR}}{\text{USD/EUR}} = \frac{\text{CHF}}{\text{USD}}$$

When none of the three selected currencies is EUR, this logic extends one step further: since the ECB only ever publishes rates quoted against the euro, none of the three triangular rates exist as a direct series in this case, and all three must be derived as ratios of the underlying EUR-quoted rates. For a triple such as CHF, USD, and GBP, for instance, each of the three rates is computed analogously from its two relevant EUR-quoted series:

$$\frac{\text{CHF}}{\text{USD}} = \frac{\text{CHF/EUR}}{\text{USD/EUR}}, \frac{\text{USD}}{\text{GBP}} = \frac{\text{USD/EUR}}{\text{GBP/EUR}}, \frac{\text{GBP}}{\text{CHF}} = \frac{\text{GBP/EUR}}{\text{CHF/EUR}}$$

Consequently, the no-arbitrage product $u_t \cdot v_t \cdot w_t = 1$ holds exactly by construction for all ECB-sourced triples, regardless of whether one or all three rates were derived, and the normalization step has no numerical effect in this setting. It remains implemented in the backend for completeness, ensuring the tool behaves correctly if applied to data from other sources where the three rates are quoted independently.

3.2 Economic Events List

In addition to the exchange rate data fetched from the ECB API, the tool displays a curated list of macroeconomic events that can be overlaid on the triangular chart. Unlike the exchange rate data, this list is not sourced from any external database or API. It is a static list hardcoded directly in the backend application, compiled manually during development.

Each event entry contains a date, a short label, a description, a set of currency tags indicating which currencies it is relevant to, a category, and an importance score from one to three. The importance score reflects the degree of public awareness an event generated rather than a measured assessment of its impact on exchange rates. When the user requests a chart, the backend filters the full event list to return only those events that fall within the selected date range and involve at least one of the three currencies in the chosen triple. The filtered list is passed to the frontend as JSON alongside the coordinate data and rendered as labelled markers on the chart.

The list covers events across all 30 currencies supported by the tool, as well as a global category for events that are relevant regardless of which currency triple is selected, such as the collapse of Lehman Brothers in September 2008 or the WHO declaration of the COVID-19 pandemic in March 2020. The scope of events is deliberately broad. Alongside monetary policy decisions and central bank interventions, the list includes geopolitical shocks, financial crises, natural disasters, political elections, and events such as terror attacks that do not have a direct macroeconomic dimension but generated significant public awareness at the time and may have produced short-lived movements in exchange rates.

The compilation followed a three-step process. As a first step, Claude (Anthropic) was used to generate a broad chronological overview of potentially relevant events for each currency, which served as an initial framework for selection. These suggestions were then cross-checked against available sources depending on the nature of the event: monetary policy decisions and central bank interventions were verified against the official press releases and decision archives of the respective central banks, while geopolitical events, crises, and other non-monetary events were verified against other articles if I didn't know the certain events. As a final step, events that were not captured in the initial AI-generated overview but were considered relevant by me were added manually to the list.

The number of entries varies considerably across currencies, which reflects differences in macroeconomic significance rather than any inconsistency in the compilation approach. Currencies such as the USD, EUR, GBP, JPY, and CHF have a larger number of associated events because the monetary policy decisions of their respective central banks carry global relevance and are closely watched by market participants worldwide. Currencies of smaller or emerging economies naturally have fewer entries, as their policy decisions and domestic events tend to have a more limited international footprint, even if they are important within their own context. Since the list is static, it does not update automatically and would require a manual edit to app.py to add new events. This is a known limitation and is discussed in chapter 7.

4 System Architecture

4.1 Overview

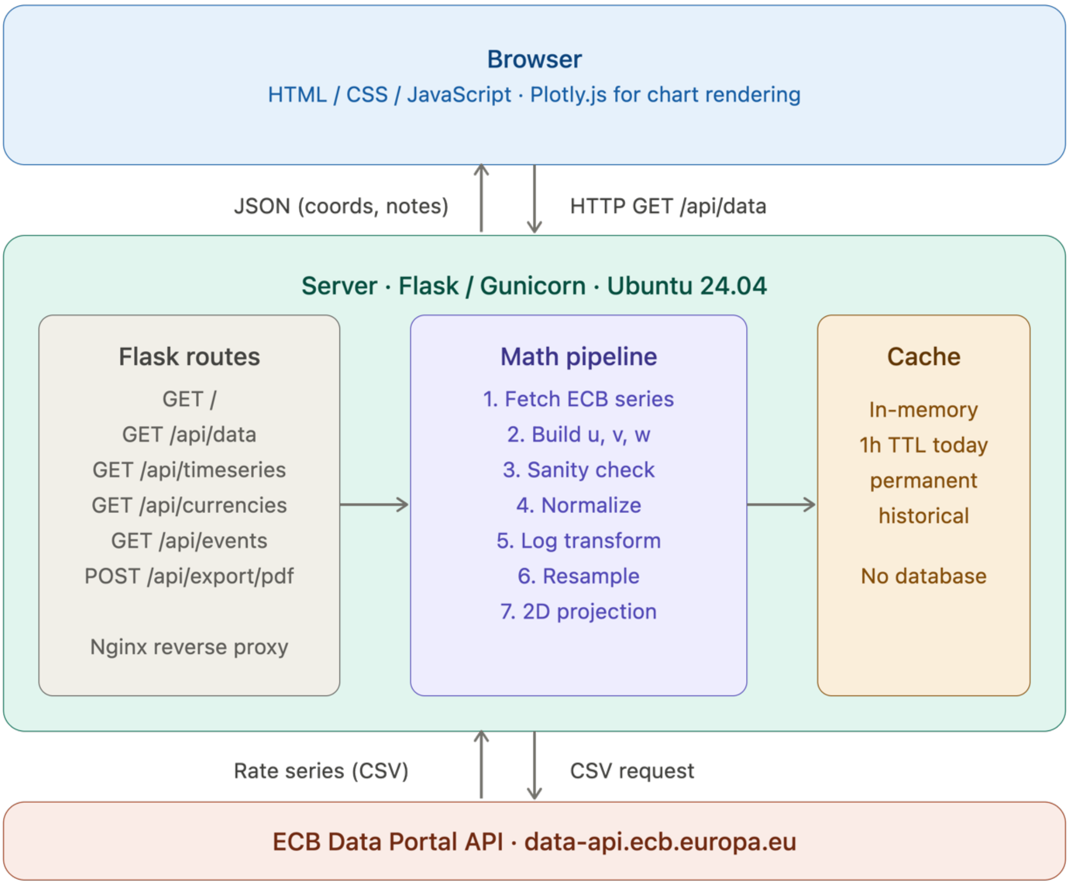

Figure 1: Quick Overview of System Architecture

The implementation of the backend architecture, including the routing structure, the caching layer, the parallel fetching mechanism, and the deployment configuration, was developed with the assistance of Claude Code (Anthropic) alongside the official documentation of the frameworks and libraries used. A full list of the technical resources and documentation pages consulted is provided in the references section.

4.2 Backend

The backend is written in Python using the Flask web framework, served by Gunicorn as the WSGI application server. Flask exposes six routes. The main route, /api/data, accepts the user's currency selection, date range, granularity, and normalization preference as query parameters and returns the full set of projected 2D coordinates as JSON, along with a sanity check value and any data notes. The route /api/timeseries bypasses the mathematical pipeline entirely and returns the raw EUR-quoted rate series, used by the time series view in the frontend. The routes /api/currencies and /api/events return a static list of available currency pairs and a hardcoded list of macroeconomic events respectively. Finally, /api/export/pdf accepts a chart image from the browser and renders it to PDF server-side using Kaleido.

The core of the backend is the mathematical pipeline, called run pipeline, which executes the steps described above: fetching the ECB series in parallel, deriving the third rate, running a sanity check on the no-arbitrage product, normalizing, log-transforming, resampling to the chosen granularity, and projecting to 2D coordinates. Data fetching is handled by a cache-aware wrapper: results are held in an in-memory dictionary, with a one-hour time-to-live for today's data and permanent storage for all historical data. A threading lock prevents duplicate requests when multiple users query the same series simultaneously.

4.2.1 Libraries

LIBRARY	ROLE
FLASK	Web framework: routing, HTML rendering, JSON responses
GUNICORN	WSGI production server
PANDAS	Time series alignment, resampling (D/W/M/Q/Y), CSV parsing
NUMPY	Math operations and array arithmetic for the 2D projection
REQUESTS	HTTP client for calling the ECB Data Portal API
PLOTLY (PYTHON)	Server-side PDF export via plotly.io.to_image()
KALEIDO	Static image renderer called by plotly for /api/export/pdf
CONCURRENT.FUTURES	ThreadPoolExecutor for parallel ECB fetches
THREADING	Lock and event for cache safety and in-flight deduplication
IO	Wraps ECB CSV text response so pandas can read it directly
DATETIME	Cache TTL logic (one-hour expiry for today's data)

Table 1: Overview of Backend Libraries

4.3 Frontend

The frontend is built with plain HTML, CSS, and JavaScript, with no additional framework. Plotly.js is used for the interactive chart, which supports zooming, panning, and hovering over individual data points. The user controls the chart through a configuration panel on the left side of the page, where they can select the currency triple, date range, granularity, and whether normalization is applied. A triangular consistency indicator shows the current value of $u_t \cdot v_t \cdot w_t$as a live sanity check. The chart can be downloaded as a PNG or SVG directly from the Plotly toolbar, or as a PDF via the server-side export route.

4.3.1 Libraries

LIBRARY	ROLE
PLOTLY.JS	All chart rendering, animation, download via toImage()
JSPDF	Client-side PDF generation for chart download
MATHJAX	LaTeX formula rendering on the About page
MAIN.JS	All app logic: API calls, chart building, events, theme
THEME.JS	Light/dark mode toggle, shared across both pages
STYLE.CSS	All styling, layout, responsive breakpoints, dark/light theme

Table 2: Overview of Frontend Libraries

4.4 Deployment

The application is deployed on a Linux server running Ubuntu 24.04, provided by Prof. Christian Ewerhart at the University of Zurich. The server is accessible without superuser rights. Gunicorn runs the Flask application on port 5000, and Nginx acts as a reverse proxy in front of it, forwarding incoming HTTPS traffic to Gunicorn and handling TLS termination. The domain and HTTPS certificate were provided alongside the server.

During development, the source code was managed in a private GitHub repository under my own account. The deployment workflow consisted of pushing changes from the local development environment to GitHub, then connecting to the server via SSH, pulling the latest commit, and restarting the Gunicorn process. This workflow is straightforward but tied to my personal GitHub account and SSH credentials, which means a future maintainer would not automatically have the access required to continue deploying updates.

4.5 Handover of the Source Code

For the handover, the repository should be transferred to an account controlled by Prof. Ewerhart or a dedicated university organization on GitHub, using the built-in transfer function available under the repository settings. The server's connection to GitHub is already established via a deploy key. The public half is registered under “Settings”, “Deploy keys” on the repository, and the private half remains on the server at ~/.ssh/id_ed25519. Because deploy keys are attached to the repository itself, not to the current owner's account, the server's ability to pull updates from GitHub is preserved automatically when the repository is transferred.

To simplify the deployment process for future maintainers, a shell script deploy.sh is included in the repository root. Running bash deploy.sh from within the project directory executes the full deployment sequence automatically: activating the virtual environment, pulling the latest code from GitHub, freeing port 5000, and restarting Gunicorn in the background using nohup to ensure the process persists after the SSH session is closed. Before running the script for the first time after a fresh clone, chmod +x deploy.sh must be run once to make it executable. This reduces the risk of errors from manually typing each command and gives any future maintainer a single, reliable entry point for deploying updates. The deployment steps and deploy key setup process are also documented in the README.md file in the repository root.

The application files are located at ~/fx-visualization on the server. The Python virtual environment is installed separately at ~/fxapp/venv. Gunicorn output and error logs are written to ~/fx-visualization/nohup.out and can be inspected after deployment to verify the application started correctly.

5 Frontend and Designing Principles

5.1 Design Principles

The visual design of the tool follows the identity of the University of Zurich. The primary colour used throughout the interface is the UZH institutional blue, which appears in the navigation bar, the “Generate Chart” button, etc. The same typeface and colour palette carry across both the main visualization page and the “About” page, giving the tool a coherent and recognizable appearance. The layout is deliberately restrained. The configuration panel uses a clean, form-based structure with no decorative elements, keeping the user's attention on the chart itself. Labels, tooltips, and directional annotations on the chart are set in a neutral sans-serif font at a size that remains legible at typical screen resolutions without cluttering the visualization. The light and dark themes are both designed to maintain sufficient contrast for all chart elements, including the colour gradient, the orange strong-movement highlights, and the event markers. The active theme (light/dark mode) is detected automatically from the operating system preference on first load, so the tool adapts without requiring any manual configuration. The theme can also be toggled at any time using the button in the navigation bar. Crucially, chart exports respect the currently active theme, so the downloaded PNG, PDF, or SVG matches what the user sees on screen.

5.2 User Interface and Features

5.2.1 Navigation Bar

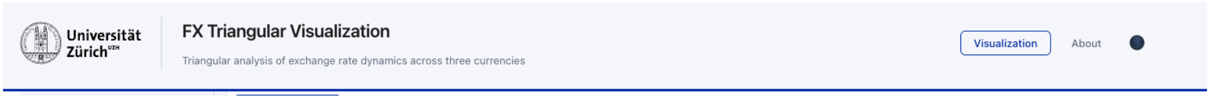

Figure 2: Navigation Bar

The University of Zurich logo in the top left corner is clickable and returns the user to the main visualization page, preserving the last loaded chart. The navigation bar also contains links to the “Visualization” and “About” pages, and the light/dark mode toggle.

5.2.2 Configuration Panel

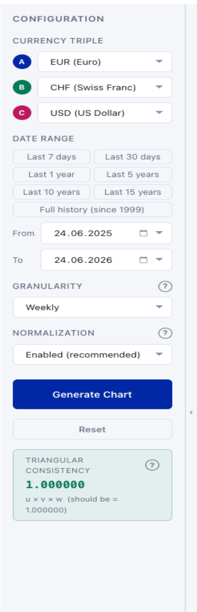

Figure 3: Configuration Panel

The panel on the left is collapsible, allowing the user to maximize the chart area once a configuration has been set. It contains four input groups. The currency triple selector assigns a currency to each of three positions, A, B, and C, from a dropdown list of all 29 ECB-quoted currencies (and EUR). The date range can be set using preset buttons, covering windows from the last seven days to the full history since 1999, or by entering a custom start and end date. When a preset or date range is selected, the granularity is updated automatically based on the length of the selected window, though it can always be overridden manually. The normalization toggle switches between the geometric mean normalization described in Chapter 2.2, Normalization enabled is the recommended default. At the bottom of the panel, the triangular consistency indicator displays the value of $u_t \cdot v_t \cdot w_t$ for the current dataset as a sanity check.

When the page is opened in a new browser tab, the configuration defaults to the EUR/CHF/USD triple over the last year with weekly granularity, providing a meaningful starting point without requiring any input from the user. This default was chosen because EUR, CHF, and USD represent the three currencies most directly relevant in a Swiss context. When the user navigates away from the

“Visualization” page and returns, either by switching between the “Visualization” and “About” pages or by clicking the logo, the last generated chart is restored rather than the default, so the user does not lose their current analysis (cache layer discussed in sections above).

5.2.3 Chart Area

Figure 4: Triangular Chart

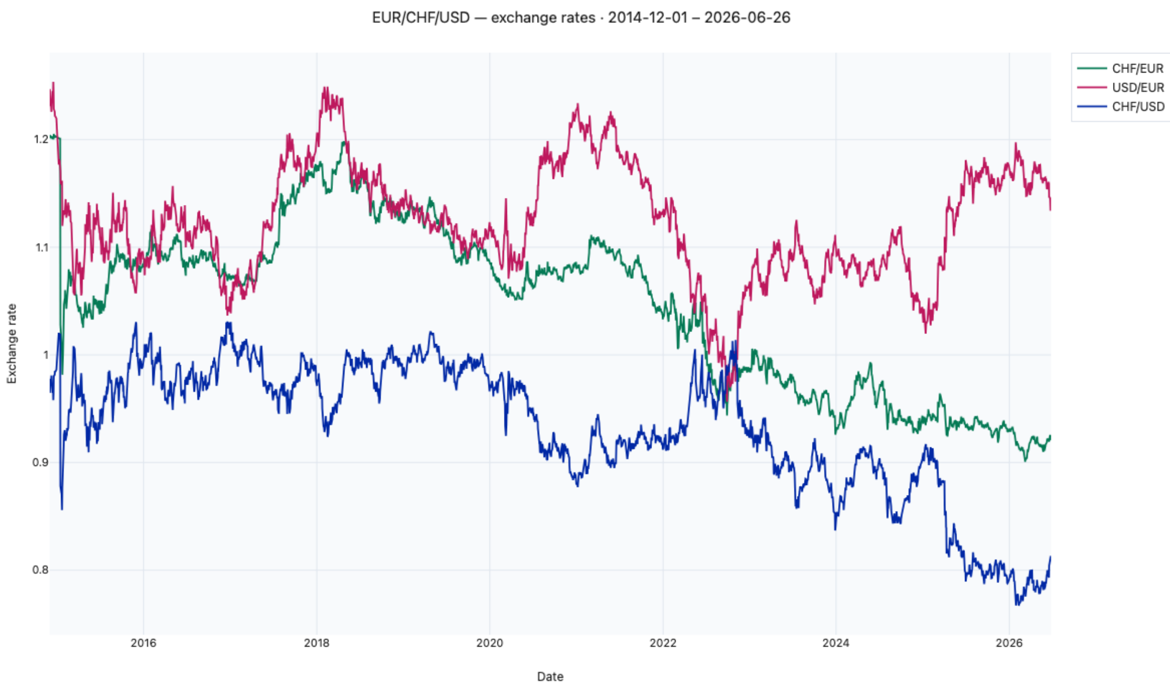

Figure 5: Time-Series Chart

The chart area contains two tabs. The triangular chart tab is the core of this project. Each data point represents one aggregated time period, and consecutive points are connected by arrows showing the direction of movement. Points are coloured along a gradient from purple through blue to red, encoding time so that the direction of travel is immediately readable even when arrows are hidden. This becomes particularly useful when working with large datasets where the number of data points makes rendering arrows impractical or visually overwhelming. In those cases, the colour gradient preserves the chronological dimension of the chart on its own, allowing the user to follow the trajectory through time without relying on directional arrows. Rendering limitations related to large arrow counts are discussed further in Chapter 7. Arrows are coloured orange when a movement exceeds the threshold defined by the median plus 4.5 times the median absolute deviation, flagging statistically anomalous steps relative to the rest of the selected period. Three directional labels appear inside the chart area in the top-left corner, indicating which direction corresponds to a strengthening of each of the three currencies.

The time series tab shows the exchange rates for the selected triple over the chosen date range. When EUR is one of the three selected currencies, two of the three rates are sourced directly from the ECB as EUR-quoted series. The third rate, which has no direct EUR-quoted counterpart and is instead derived as the ratio of the two available series. When none of the three selected currencies is EUR, none of the triangular rates exist as direct ECB series, since the ECB only publishes rates against the euro. In this case all three rates are derived as ratios of the underlying EUR-quoted series.

Below the triangular chart, a time slider allows the user to animate the trajectory step by step. A "Play/Pause" button starts and stops the animation, and the slider can be dragged manually to any point in the timeline. For datasets with more than 300 data points the slider is subsampled to at most 300 steps, so each step may advance the animation by more than one data point. Data notes appear below the chart whenever the requested date range includes days for which the ECB did not publish rates, explaining exactly how those cases were handled.

5.2.4 Explaining Strong Movements

A visual feature of the chart that required deliberate methodological attention is the highlighting of statistically anomalous movements. When a step between two consecutive projected points is large relative to the rest of the selected period, the corresponding arrow is coloured orange. The initial implementation used the classical three-sigma rule: a step was flagged as strong if it exceeded the mean step size by more than three standard deviations,

$$\tau_{\mu\sigma} = \bar{d} + 3\sigma_d \text{ where } \bar{d} = \frac{1}{n}\sum_{i=1}^{n} d_i \,, \sigma_d = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(d_i - \bar{d})^2}$$

Under a normal distribution this threshold captures approximately 99.7% of observations, leaving only the most extreme 0.3% flagged as outliers. Two properties of exchange rate data make this approach problematic.

Exchange rate changes do not follow a normal distribution. The empirical evidence is well established that their distributions exhibit excess kurtosis, meaning extreme moves occur far more often than a Gaussian model predicts (Boothe & Glassman, 1987). As a result, $\bar{d} + 3\sigma_d$is exceeded more frequently than the nominal 0.3%, making the threshold less selective than intended. Furthermore, exchange rate movements exhibit volatility clustering. Turbulent periods tend to follow turbulent periods (Mandelbrot, 1963), so the assumption of independent observations underlying the classical rule does not hold. A σ_d computed over the entire selected

range can be inflated by a crisis episode, causing ordinary moves in calmer sub-periods to never be flagged while during the crisis itself many consecutive arrows exceed the threshold at once.
To address both problems, the implementation was revised to use the Median Absolute Deviation (MAD), a robust statistics alternative that makes no distributional assumptions (Pham-Gia & Hung, 2001):

$$\tilde{d} = \text{median}\,(d_1, \ldots, d_n)$$
$$\text{MAD} = \text{median}\,(|\, d_i - \tilde{d}\, |)$$
$$\tau_{\text{MAD}} = \tilde{d} + 4.5 \cdot \text{MAD}$$

The median and MAD are both resistant to fat-tailed distributions, since a small number of extreme values cannot pull the threshold upward the way a large σ_d can. Orange arrows therefore only appear when at least one step is genuinely anomalous relative to the typical step size in the selected range.
The multiplier $k = 4.5$ was chosen because under a normal distribution, MAD and the standard deviation σ are related by a known scaling constant:

$$\text{MAD} \approx 0.6745\,\sigma \Leftrightarrow \sigma \approx 1.4826 \cdot \text{MAD}$$

This means $k = 4.5$ multiples of MAD correspond to approximately $4.5/1.4826 \approx 3.0\sigma$, which is precisely the three-sigma threshold of the original rule. The MAD-based approach therefore inherits the familiar nominal flagging rate of the classical rule while remaining valid for the fat-tailed, autocorrelated distributions that characterize exchange rate data.

5.2.5 Events Panel

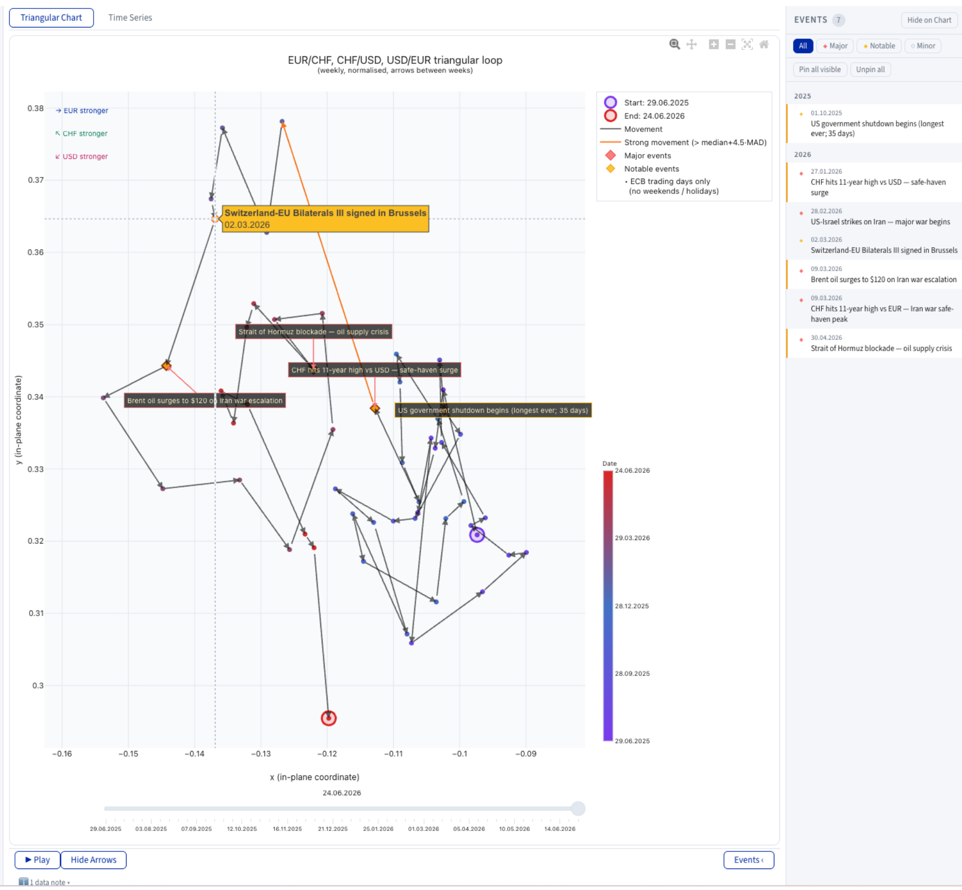

Figure 6: Demonstration of Events Panel

The events panel is collapsible and opens from the right side of the chart area. It lists all events from the list that fall within the selected date range and involve at least one of the chosen currencies. Events are classified into three tiers: major events, shown with a red diamond, cover central bank rate decisions and large-scale interventions; notable events, shown with an amber diamond, cover significant geopolitical developments and elections; and minor events, shown with a grey open diamond, cover contextual reference events. The user can filter the list by tier using the “All”, “Major”, “Notable”, and “Minor” buttons at the top of the panel.

Events can be shown or hidden on the chart collectively using the “Show on Chart” button, which toggles all currently visible events as markers. Individual events can also be pinned, which attaches a labelled annotation directly to the corresponding data point on the chart. Pinned events remain visible regardless of the global show/hide toggle, allowing the user to keep selected reference points visible while hiding the rest. The “Pin all” visible and “Unpin all” buttons apply these actions to all events in the current filtered list at once. As visible in the screenshot above, pinned annotations display the event date and label directly on the chart, giving the visualization a self-contained “narrative”. When multiple events fall within the same aggregation period and therefore map to the same projected chart point, their markers overlap with different sizes per tier, and hovering over the marker reveals all events in a stacked tooltip listed chronologically by date.

5.2.6 Downloads

Figure 7: Download Possibilities

The charts can be exported as a PNG, PDF, or SVG using the download buttons below the chart area. Each export is automatically labelled with the currency triple and chart type. The underlying exchange rate data can also be downloaded as a CSV file, allowing users to work with the raw series independently of the tool.

6 Example Analysis with the Tool

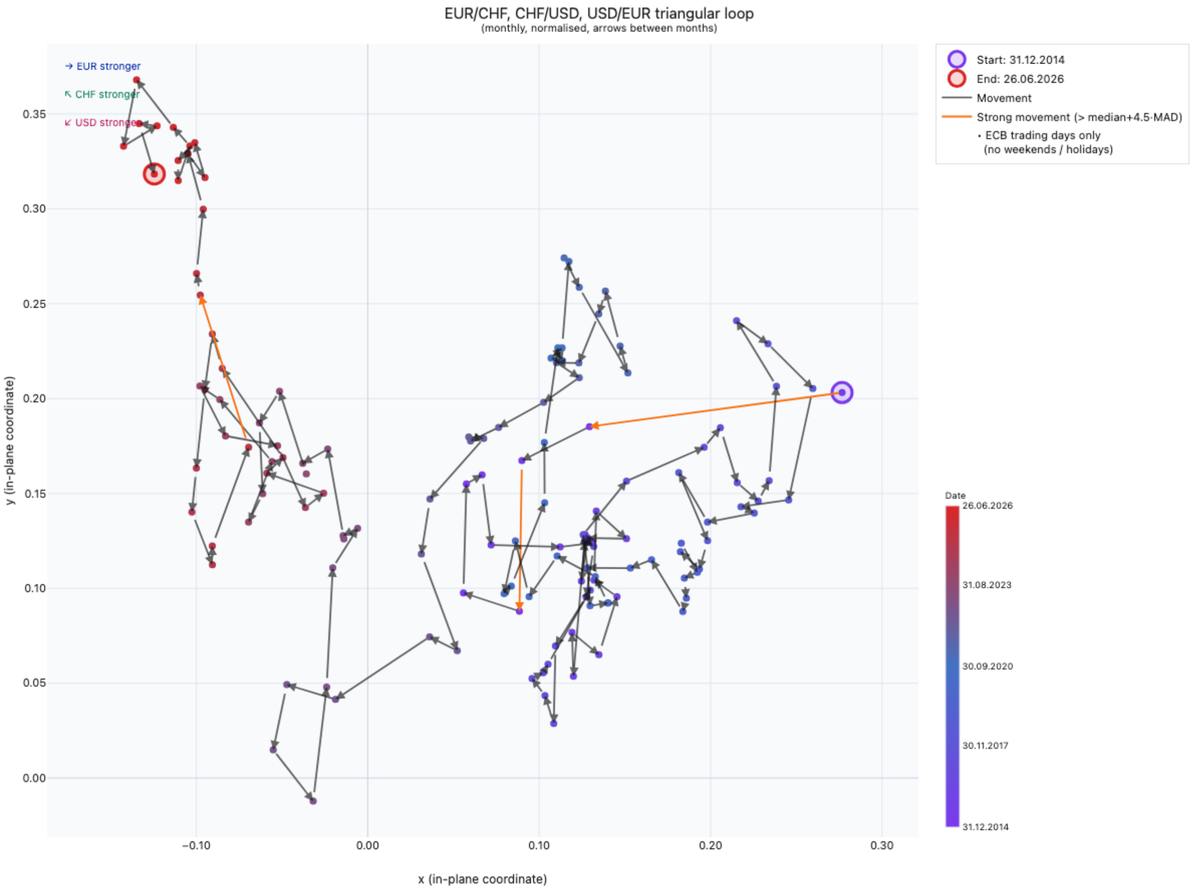

Figure 8: Triangular Chart from Dec 2014 – Today

6.1 Reading the Configuration

The starting point of the trajectory, marked by the open purple circle at approximately (0.28, 0.2), represents the relative configuration of the three currencies in December 2014. As the directional labels in the upper left of the chart indicate, a rightward position corresponds to EUR strength, an upper-left position to CHF strength, and a lower-left position to USD strength. The starting point sitting to the right and slightly above center therefore encodes a configuration where the EUR was the strongest of the three currencies at the opening of the period, with the CHF held in an intermediate position and the USD the weakest of the three. This reflects the macroeconomic context of late 2014: the SNB's EUR/CHF floor at 1.20 was still in place, artificially suppressing franc strength, the ECB had not yet launched its quantitative easing programme, and while the dollar had been gaining momentum through the second half of 2014 on expectations of Fed tightening, that strengthening cycle had not yet run its course. The trajectory that follows shows exactly how this configuration shifted in the course of the following years.

Eleven years later, the open red circle in the upper-left corner shows a different picture. The clearest macro trend running through the entire period is a strengthening of the CHF against both the EUR and the USD. The EUR weakened persistently against the franc throughout the eleven years. The USD tells a more nuanced story: it was broadly stable or even strong against the EUR after the period started but weakened noticeably against both counterparts in recent times from 2025 onwards, driven by trade policy uncertainty and safe-haven flows into the franc.

Three orange arrows mark movements that exceeded the statistical threshold of median plus 4.5 times the median absolute deviation. These three anomalous steps correspond to the three largest single-month shocks visible in the data, and each of them is perfectly explainable.

6.2 2015: The SNB Shock and ECB QE

The starting position of the trajectory is immediately relevant. In December 2014 the Swiss franc was already under upward pressure, as investors anticipated that the SNB's defence of the EUR/CHF minimum rate of 1.20 was becoming increasingly expensive against the backdrop of a weakening euro. On 15 January 2015, the SNB abandoned the floor without warning and simultaneously cut its policy rate to -0.75%, the lowest in the world at that time, to cushion the impact of the franc's appreciation (Swiss National Bank, 2015). The franc surged by as much as 30% against the euro within minutes of the announcement before settling around 13% higher, one of the largest single-day moves ever recorded in a major currency pair (Bishop, 2015). One week later, on 22 January, the ECB announced its €1.1 trillion quantitative easing programme, adding further downward pressure on the euro (European Central Bank, 2015).

In the triangular chart, these events produce the first orange arrow, running almost horizontally to the left. A purely horizontal leftward movement means that CHF and USD gained against the EUR in roughly equal measure, which is precisely what the floor removal produced. The EUR weakened simultaneously against both counterparts, while the bilateral relationship between the CHF and the USD remained relatively undisturbed in that first month.

The arrow between February and March 2015 also appears in orange, but its direction is notably different, running more vertically downward. A downward movement means that USD strengthens against the CHF, with the EUR remaining relatively stable. This reflects the continuation of market repricing. Once the initial EUR weakness had been absorbed, the subsequent adjustment shifted toward a reassessment of the CHF/USD relationship, as the SNB's aggressive negative rate of -0.75% made holding francs less attractive and the dollar began to benefit from growing expectations of Fed tightening, with the Federal Open Market Committee signalling as early as March 2015 that a rate increase would follow once labor market and inflation conditions were met (Board of Governors of the Federal Reserve System, 2016). The EUR/CHF rate had moved 14% by the end of March relative to January 14, and then moved even more, meaning the full repricing of the combined SNB and ECB decisions extended across several months rather than resolving in a single step.

6.3 2015 – 2019: Fed Tightening and Diverging Monetary Policies

After the SNB shock settled, the trajectory in the second half of 2015 moved downward and to the right. The downward component reflects USD strengthening against the CHF specifically, driven by the Fed's first rate increase since 2006 on 16 December 2015, raising its target range to 0.25–0.5% (Board of Governors of the Federal Reserve System, 2015). The rightward component reflects the EUR partially recovering against the CHF, as the initial shock from the floor removal began to fade and markets stabilised. The SNB held its rate at -0.75% throughout this period, but the acute pressure on the franc eased as the most dramatic repricing had already occurred.

Through 2016 and into 2017, the EUR continued recovering against the CHF, with the bilateral rate approaching the old floor level of 1.20 again by early 2018. In the chart this appears as a continued rightward movement, reflecting EUR strength against CHF. The Fed continued hiking slowly, three further 25 basis point increases through 2016 and 2017 bringing the rate to 1.25–1.5% by December 2017, and four more hikes in 2018 took it to a peak of 2.25–2.5% by December

2018 (Anderson et al., 2017; Foster & Lowery, 2026). Political uncertainty around Trump's first term briefly weighed on the dollar through 2017, visible as temporary oscillations in the trajectory.

From late 2018 onwards the picture reversed. The EUR began weakening again against both the CHF and the USD, and the trajectory turned leftward and downward. The Fed signalled a pause and then delivered three cuts in 2019, bringing the rate back to 1.5–1.75% by October (Bahr 2025), but this did not help the EUR, as the ECB on 12 September 2019 cut its deposit rate further to -0.5% and restarted asset purchases at a monthly pace of €20 billion, putting renewed downward pressure on the euro against both counterparts (European Central Bank, 2019). By early 2020, as the trajectory approached the COVID pandemic, the configuration had drifted back to a position not far from where it stood in the months immediately after the SNB floor removal in early 2015, suggesting that the EUR recovery of 2016 to 2018 had fully unwound and the three currencies had returned to a similar trilateral configuration after five years of fluctuation.

6.4 2020 – 2021: COVID Pandemic

In January, the WHO declared COVID-19 a global health emergency, and in March a pandemic. The Fed responded with an emergency cut to 0–0.25% on 15 March and launched unlimited quantitative easing on 23 March (Board of Governors of the Federal Reserve System, 2020). The ECB launched its €750 billion Pandemic Emergency Purchase Programme on 18 March (European Central Bank, 2020). The SNB held its rate at -0.75% and intensified foreign exchange interventions to resist safe-haven appreciation of the franc.

In the triangular chart, the COVID shock is visible as a sharp vertical upward movement in early 2020. A vertical upward direction means CHF and EUR both strengthening against the USD simultaneously, while the bilateral CHF/EUR relationship remained relatively undisturbed. This pattern suggests the dollar bore the brunt of the shock in the triangular configuration, while both European currencies held their ground. The franc in particular likely benefited from its traditional safe-haven status, drawing capital inflows as global uncertainty spiked. Through 2021 the trajectory moved relatively tightly around the same zone, reflecting the broadly aligned ultra-loose stances of all three central banks: the Fed at 0–0.25%, the ECB deposit rate at -0.5% the SNB at -0.75%. With policy rates anchored near zero across all three and coordinated fiscal stimulus dampening exchange rate volatility, there was little rate-differential signal to drive sustained directional moves.

Beneath the surface, however, inflationary pressures were building. The combination of supply chain disruptions, surging energy prices, and pent-up consumer demand produced the largest period of inflation since 2008 (U.S. Bureau of Labor Statistics, 2022), with individual monthly readings by late 2021 reaching levels not seen since 1982 (Cox, 2021). By late 2021 inflation had already exceeded 5% in the US (European Central Bank, 2022b) and was accelerating in Europe, yet all three central banks were still holding rates at or near zero. The Fed began tapering its asset purchases in November 2021 (Board of Governors of the Federal Reserve System, 2024) but had not yet moved on rates. Switzerland was a notable exception to the inflationary pattern: the SNB's longstanding policy of franc strength effectively acted as a natural buffer against imported inflation, keeping Swiss consumer prices far more contained than in the eurozone or the US. This came at a cost, however, as a persistently strong franc weighed on the competitiveness of Swiss exporters, particularly in the machinery, watchmaking, and pharmaceutical sectors, for whom a strong currency directly erodes margins on foreign sales. The policy response that followed would break the triangular stability visible in this period.

6.5 2022 – 2023: Russia, Inflation, and Coordinated Tightening

Russia's invasion of Ukraine in February 2022 sent energy prices surging, with Brent oil prices surging in March, compounding the post-pandemic supply disruptions that were feeding inflation. The policy response was an aggressive synchronized tightening cycle.

The Fed moved first, raising rates in March 2022 for the first time since December 2018 (Board of Governors of the Federal Reserve System, 2022; Cox, 2022), and then delivering four consecutive 75 basis point increases between June and November 2022 (Rodini & Salvucci, 2026). The federal funds rate reached 5.25–5.5% by July 2023, a 22-year high (Rodini & Salvucci, 2026). The ECB followed, raising rates for the first time since 2011 in July 2022, delivering a larger-than-expected 50 basis point hike (European Central Bank, 2022a). The SNB acted most dramatically relative to its own recent history. After eight years in negative territory, it raised rates and ened up with 0.5% in September 2022.

The initial phase of 2022 sees the trajectory drift toward USD strength as the Fed moved earliest and most aggressively, reflecting the rate differential advantage the dollar was accumulating. As the ECB and SNB caught up through late 2022, the trajectory partially reversed, with the EUR and CHF recovering some ground against the dollar.

Silicon Valley Bank failed in March 2023, triggering a wave of concern about interest rate risk embedded in bank balance sheets after the rapid hiking cycle. Signature Bank was closed two days later, and First Republic Bank subsequently failed on 1 May, with JPMorgan Chase acquiring the substantial majority of its assets and deposits from the FDIC (JPMorgan Chase & Co., 2023). The stress was not confined to the US. Credit Suisse, which had been under pressure for years due to a series of risk management failures, became the focal point of European banking anxiety and was absorbed by UBS in a forced takeover engineered by the Swiss authorities over a single weekend in March 2023. The combination of US bank failures and the Credit Suisse collapse created a brief but sharp flight to safety across markets. In the triangular chart this is visible as a short leftward movement, reflecting safe-haven demand for the franc despite Switzerland being the source of one of the shocks.

6.6 2024 – 2025: Diverging Cutting Cycles and Trump Uncertainty

As inflation fell back toward target levels across all three economies, the cutting cycles began but diverged sharply in timing and pace. The SNB moved first among G10 central banks, cutting by 25 basis points in March 2024, followed by the ECB in June 2024. The Fed did not begin cutting until September 2024 with a larger-than-expected 50 basis point reduction, followed by 25 basis point cuts in November and December 2024, bringing the rate to 4.25–4.5%. The SNB cut aggressively through 2024 and 2025, moving through five reductions to reach 0.25% by March 2025 and then cutting to 0% in June 2025, returning to the zero lower bound. The ECB completed eight cuts between June 2024 and June 2025, bringing its deposit rate from 4.0% to 2.0%. The Fed was far more cautious, pausing for five consecutive meetings in early 2025 before delivering three more cuts in September, October, and December 2025, bringing its rate to 3.5–3.75% (Swiss National Bank, 2026).

Trump's election victory in November 2024 introduced a new and significant source of uncertainty. Beyond the immediate policy implications of his return to office, Trump repeatedly and publicly pressured Federal Reserve Chair Jerome Powell to lower interest rates, at times questioning Powell's position. Powell held firm, insisting that monetary policy decisions would remain data-driven and independent of political influence. This tension raised concerns among market participants about the institutional integrity of the Fed. From Trump's election onwards, the trajectory shows the dollar losing ground against both the franc and the euro, visible as an upward and slightly leftward movement from early 2025 onward. One reason for that could be “Liberation Day” tariff announcement in April 2025. Trump announced sweeping tariffs on nearly all trading partners, declaring a national emergency over the US trade deficit. All G10 currencies appreciated against the dollar in the 24 hours following the announcement because the tariffs were seen as harmful to the US economy, raising fears of slower growth. These uncertainties produce the last strong movement visible on the very left side of the chart. For the rest of the year the trajectory remained more or less at the same point.

6.7 2026: Iran War and Safe-Haven Peak

The most recent phase of the trajectory brings the CHF to its highest relative position of the entire eleven-year period. In February 2026, the United States and Israel launched strikes on Iran, sending oil prices sharply higher and intensifying safe-haven demand for the franc. The point sitting at the very top-left of the entire trajectory at approximately $(-0.13, 0.37)$, marks the franc reaching its strongest position against both counterparts simultaneously over the full period covered. The franc reached an 11-year high against the euro in March 2026 and the dollar in January 2026 (Trading Economics, 2026). Following the peak, the trajectory turns partially back toward the right as oil prices stabilised somewhat and the ECB raised rates for the first time since 2023 in June 2026 (Swiss National Bank, 2026), responding to renewed inflation pressure from the energy shock. The Fed, now under its new chair Kevin Warsh, held rates at 3.5–3.75% through four consecutive meetings (Swiss National Bank, 2026) while signalling that rate hikes were back on the table if inflation did not subside (Gould, 2026).

The trajectory sitting in the upper-left quadrant, far from where it started, representing eleven years of persistent franc strengthening against both the euro and the dollar, a trend that kicked off with the SNB floor removal in January 2015 and accelerated through each successive episode of global uncertainty that followed. What stands out across the full chart is not any single shock but the consistency of the direction: every major disruption discussed in this chapter, the SNB floor removal, the COVID pandemic, the 2022 inflation surge, the banking stress of 2023, and the Iran war of 2026, pushed the trajectory the same way rather than reversing it. The euro and the dollar each had periods of relative strength against one another, visible in the oscillations of the chart, but neither currency managed to gain lasting ground against the franc. This pattern is consistent with the franc's long-standing role as a safe haven currency, where uncertainty of almost any kind, whether a domestic Swiss policy decision or a geopolitical conflict thousands of kilometres away, tends to attract capital into Swiss assets rather than away from them. The triangular visualization makes this structural asymmetry between the three currencies visible in a way a bilateral chart would not: rather than three separate stories of EUR/CHF, USD/CHF, and EUR/USD each moving independently, the single trajectory shows one coherent eleven-year narrative of franc strength reasserting itself again and again, regardless of which event triggered the most recent move.

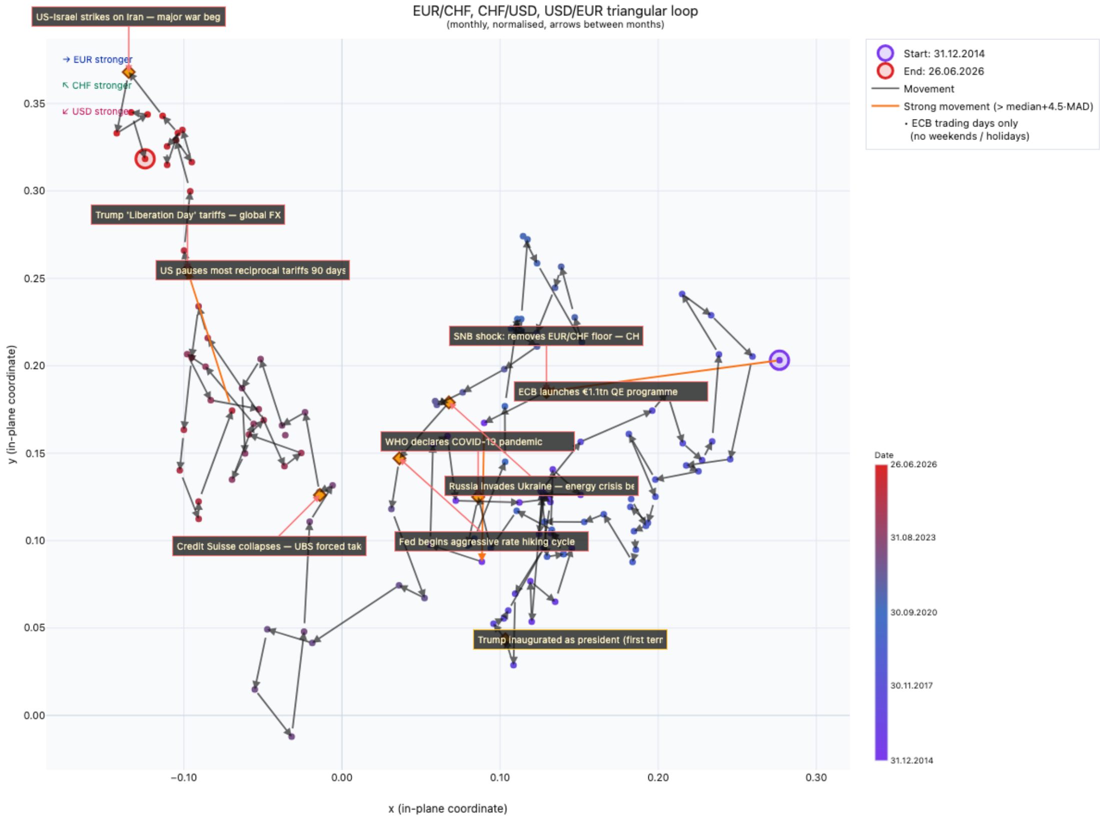

Figure 9: Same Figure as Figure 8 but with labelled Events

6.8 Time-Series Chart

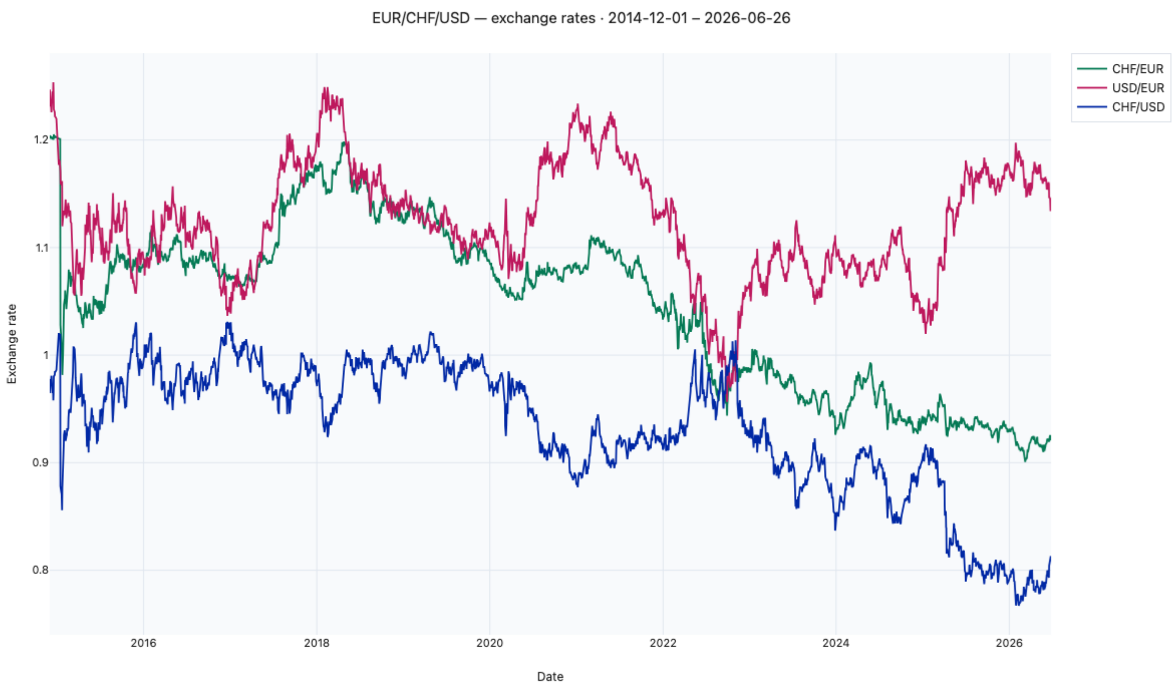

Figure 10: Regular Chart with plotted Time Series Data

The time series chart above plots the three bilateral exchange rates directly. CHF/EUR in green, USD/EUR in red, and the derived CHF/USD in blue. Several of the macro trends discussed in this analysis are visible here, including the sharp CHF/EUR drop in January 2015, the broad USD strength through 2015 to 2018, and the persistent CHF appreciation from 2022 onwards.

What the time series does not convey easily is the joint dynamic between all three currencies at any given moment. When multiple rates move simultaneously, one must cross-reference several lines to identify which currency was the primary driver and which were passive counterparts. The triangular chart compresses this into a single trajectory where direction encodes attribution immediately. A horizontal leftward arrow means EUR weakness against both others, a vertical upward move means CHF strength against the USD specifically, and a diagonal reflects a more complex redistribution across all three. For anyone interested in a single bilateral rate in isolation, the time series remains the natural format. For questions about how the three-way configuration shifted as a whole, particularly during episodes like the SNB shock, the COVID period, or the coordinated tightening of 2022, the triangular chart offers a level of geometric clarity that three separate lines cannot match.

It is also worth acknowledging that the time series chart has a different kind of accessibility. Exchange rates expressed as simple bilateral numbers are something most people have encountered before, whether checking a rate before a holiday or watching a headline about a currency's value, and that familiarity can make the time series feel more immediately readable to a general audience. The triangular chart, by contrast, asks the reader to learn a new visual language. The triangular chart was built in way that the learning curve shouldn't be too steep, with explanations on the web-page and also annotations on the graph.

7 Problems and Limitations along the Way

The theoretical foundations and the data pipeline presented relatively few difficulties during development. The mathematical transformation from three exchange rate series to a two-dimensional trajectory follows a well-defined sequence of steps, and the ECB Data Portal provides a clean and consistently formatted API. The more significant challenges were rather technical.

The most significant performance problem was the initial approach to fetching data from the ECB API. The original implementation fetched the required series sequentially, one request at a time, meaning a three-currency query involving full history since 1999 could take 60 to 90 seconds before the chart appeared. The solution was to switch to parallel fetching using a ThreadPoolExecutor, so that all required series are requested simultaneously. The total wait time then collapses to the duration of the single slowest request rather than the sum of all requests, reducing typical load times. An in-memory caching layer was added alongside this, storing completed fetches keyed by series and date range, so repeated requests for the same data return instantly without hitting the ECB API again.

Another problem that occurred concerned the time slider animation. Plotly renders arrows as individual annotation objects attached to data points, and when the animation advances frame by frame through a large dataset, the arrow rendering lagged significantly behind the moving points, producing a visually incorrect result. I decided to keep arrows as static elements fixed to the full dataset rather than animating them frame by frame. As a practical consequence, the animation performs best or is most fluent when arrows are hidden. For datasets with a large number of data points, the slider is additionally subsampled to at most 300 steps to keep the animation smooth.

One limitation that persists is arrow bleeding when the user zooms into the chart. Plotly clips the visible plot area but arrows, being annotation objects rather than trace elements, are attached to their source data point rather than to the viewport. When zooming in to the point where a data point moves outside the visible area, its arrow may remain partially visible inside the plot boundary, creating a visual artefact. The arrow disappears entirely only once its source data point is fully outside the viewport. This causes exported and zoomed-in chart that may have some residuals of arrows on the chart title.

A further limitation concerns the economic event list. Because the list is compiled manually and hardcoded directly in the application, it does not update automatically as new events occur. Keeping it current requires periodic manual edits to the backend code, which in turn requires a new deployment to the server. Beyond the maintenance burden, the selection of events is inherently subjective. The list reflects the judgment of me as to which monetary policy decisions, geopolitical shocks, and macroeconomic events were significant enough to include. Another user or researcher working with the tool may find events they consider relevant to be missing or may disagree with the importance ratings assigned. Ideally the long-term solution could be to move the event list out of the application code and into a configuration file or lightweight database that an authorised user could edit without requiring a full redeployment, and to provide a mechanism for users to suggest additions.

8 Use of AI

The most substantial contribution was in the area of coding. Claude Code was used throughout the development of the application. This shifted the nature of the work considerably. Rather than spending most of the time on the syntax and mechanics of writing code, it became possible to focus on the conceptual questions: how to structure the pipeline, which design decisions to make at the frontend, how to handle edge cases in the data. The code itself became something to review, test, and direct rather than something to produce line by line.

Deployment was an area where AI assistance proved particularly valuable given a complete lack of prior experience. Setting up a Linux server, configuring Nginx as a reverse proxy, managing a Python virtual environment, handling SSH authentication, and troubleshooting Gunicorn errors are not skills typically acquired in an economics degree. Working through these problems with AI guidance made it possible to build and maintain a production deployment independently, and the knowledge gained in doing so is transferable to future projects. It is self-explanatory that besides AI help also the guides the referring frameworks were used as a help source.

In writing, Claude's Sonnet 4.6 was used to improve precision and clarity, checking grammar, and suggesting more differentiated vocabulary. The arguments, the analytical judgments, and the final editorial decisions were mine.

A few years ago, it would not have been realistic for me, with no real background in web development or server administration to build, deploy, and maintain a full-stack web application while simultaneously working through the theoretical and analytical content of a thesis. The combination of tasks would simply have taken too long, and the learning curve on the technical side could've crowded out everything else.

9 Discussion and Conclusion

I was provided with a clear foundation to build on. The core idea, projecting three exchange rate series onto a two-dimensional plane by exploiting the triangular no-arbitrage condition, was given. What developed from there involved a substantial number of additions and refinements. The colour gradient encoding time along the trajectory, the MAD-based detection of statistically anomalous movements highlighted in orange, the economic event list covering all 30 supported currencies, the parallel data fetching mechanism, the collapsible interface panels, the animation slider, the light and dark mode theming, etc. were all decisions made in the course of the project rather than prescribed in advance.

9.1 Usability

For researchers, the tool offers a compact way to survey a decade or more of trilateral exchange rate dynamics in a single chart. The starting point that the trajectory encodes the initial configuration of the three currencies, and the direction of each arrow that shows which currency drove any given movement, means that hypotheses about policy transmission, contagion, or safe-haven behaviour can be visually tested against the data before any formal econometric analysis begins. The event list provides a first layer of contextual annotation, connecting visible features of the trajectory to monetary policy decisions, geopolitical shocks, and financial crises, which helps identify periods worth examining more closely. The CSV export makes the underlying data immediately available for further analysis in any standard tool.

For students and instructors, the interactive nature of the tool is its main strength. Selecting a currency triple, adjusting the date range, and watching the trajectory shift is an intuitive way to develop intuition for how exchange rates move together rather than in isolation. The event overlay turns this into something closer to a guided exercise. A student can be asked to look at a ten-year chart, identify the most prominent movements, and form a hypothesis about their causes before checking the event list for confirmation or correction. The SNB floor removal of January 2015 is a particularly clear example, where the chart shows a sharp horizontal arrow that visually shows the simultaneous weakening of the EUR against both the CHF and the USD in a way that is immediately legible even without prior knowledge of the event. One could also test theory against what happened empirically after e.g. central banks made certain decisions. An instructor/lecturer can use the same feature in reverse, displaying the chart with events hidden and asking the class to identify the shock.

For market practitioners and analysts, the tool provides a quick overview of the current and historical configuration of any currency triple across a broad set of 30 currencies. Pairs or triples that have drifted far from their historical equilibrium zones, or where a recent movement appears anomalous relative to the surrounding period, can be identified visually in seconds. The tool does not replace quantitative risk models, but it provides a fast and accessible first look at trilateral dynamics that bilateral charts do not offer. In the end market-analysts can leverage similar use as researchers.

9.2 General Assessment

The tool is designed to be accessible without prior technical knowledge. Loading a chart requires selecting three currencies and clicking a button, and the default configuration already provides a meaningful starting point. The configuration panel, the animation, the event overlay, the export options, and the “About” page are all discoverable through the interface without requiring a manual. To get a sense of whether this intuitive intent translated into practice, the tool was shared informally with a small group of friends and family with varying backgrounds, ranging from economics students to people with no particular interest or prior exposure to financial markets. Their feedback was used to identify points of confusion and inform further corrections to the interface. Whether the result fully meets the standard of intuitive usability is ultimately for the user to judge, but the design principle throughout was to keep the interface clean and the controls self-explanatory.

The range of possible configurations is broad. With 30 currencies, there are over 3’000 possible currency triples, each explorable across any date range from 1999 to the present, at five levels of granularity, and with event overlay. The tool does what it set out to do: it makes the joint dynamics of three currencies readable in a single chart, and does so in a way that is accessible, configurable, and exportable for use beyond the browser.

9.3 Extensions and Improvements

Several natural extensions suggest themselves. The economic event list would benefit from being moved out of the application code into a configuration file or lightweight database, allowing it to be updated without a full redeployment and potentially opened for user contributions. The ability for users to add their own annotations directly to the chart, marking a date of interest with a custom label, would extend the tool's usefulness where the user has knowledge about events that were impactful but are not listed.

A more ambitious extension would involve using the historical triangular trajectory, potentially combined with macroeconomic indicators or the event data itself, as input to a regression or machine learning model aimed at forecasting future currency movements rather than only visualizing historical ones. Such an extension would need to be approached with caution, however, since exchange rate forecasting is a long-standing and largely unresolved challenge in international finance. The seminal result by Meese and Rogoff (1983) showed that structural exchange rate models fail to systematically outperform a simple random walk in out-of-sample forecasts, a finding that has proven remarkably robust across subsequent decades of research. Any predictive extension to this tool would therefore need to be framed honestly as an exploratory or illustrative feature rather than a reliable forecasting instrument.

9.4 Reflection

Looking back at the project as a whole, what stands out is how much was learned across domains that would not normally be part of an economics curriculum. The mathematical properties of the triangular no-arbitrage condition and the geometry of the log-normalization and projection gave a rigorous foundation that made the visualization intellectually defensible rather than merely visually appealing. The statistical literature on fat-tailed distributions in exchange rate data and robust outlier detection led to a methodological choice, the MAD-based threshold, that is better grounded than the standard sigma rule it replaced.

On the technical side, building and deploying a full-stack web application from scratch, including the server configuration, the API design, the caching layer, and the frontend, was a practical education in systems thinking that no course had provided.

The use of AI tools throughout this process fundamentally changed the nature of the work. Delegating the implementation of well-defined coding tasks to Claude Code shifted the time budget away from writing code and toward making decisions and understanding the underlying concepts in depth, which in turn allowed the project as a whole to proceed more efficiently.

References

- Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological), 44*(2), 139–177.
- Akram, Q. F., Rime, D., & Sarno, L. (2008). Arbitrage in the foreign exchange market: Turning on the microscope. *Journal of International Economics, 76*(2), 237–253.
- Anderson, A., Ihrig, J., Styczynski, M.-F., & Weinbach, G. C. (2017, June 2). How have the Fed's three rate hikes passed through to selected short-term interest rates? *Federal Reserve Bank of New York Liberty Street Economics*. https://www.federalreserve.gov/econres/notes/feds-notes/how-have-the-feds-three-rate-hikes-passed-through-to-selected-short-term-interest-rates-20170602.html
- Bahr, K. (2025). The Federal Reserve and interest rate changes. University of Wisconsin-Stevens Point. https://blog.uwsp.edu/cps/2025/03/25/the-federal-reserve-and-interest-rate-changes
- Bishop, K. (2015, January 15). Swiss franc soars, stocks tank as euro peg scrapped. *CNBC*. https://www.cnbc.com/2015/01/15/swiss-franc-sours-stocks-tank-as-euro-peg-scrapped.html
- Board of Governors of the Federal Reserve System. (2015, December 16). *Federal Reserve issues FOMC statement* [Press release]. https://www.federalreserve.gov/newsevents/pressreleases/monetary20151216a.htm
- Board of Governors of the Federal Reserve System. (2016). *Monetary policy and economic developments* [Annual Report 2015]. https://www.federalreserve.gov/publications/annual-report/2015-monetary-policy-and-economic-developments.htm
- Board of Governors of the Federal Reserve System. (2020, March 15). *Federal Reserve issues FOMC statement* [Press release]. https://www.federalreserve.gov/newsevents/pressreleases/monetary20200315a.htm
- Board of Governors of the Federal Reserve System. (2022, March 16). *Federal Reserve issues FOMC statement* [Press release]. https://www.federalreserve.gov/newsevents/pressreleases/monetary20220316a.htm
- Board of Governors of the Federal Reserve System. (2024). *The Federal Reserve's responses to the post-Covid period of high inflation*. https://www.federalreserve.gov/econres/notes/feds-notes/the-federal-reserves-responses-to-the-post-covid-period-of-high-inflation-20240214.html
- Boothe, P., & Glassman, D. (1987). The statistical distribution of exchange rates: Empirical evidence and economic implications. *Journal of International Economics, 22*(3–4), 297–319.
- Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton University Press.
- Cox, J. (2021, December 10). CPI November 2021: Inflation surged 6.8% in November, more than expected, to fastest since 1982. *CNBC*. https://www.cnbc.com/2021/12/10/consumer-price-index-november-2021.html
- Cox, J. (2022, March 16). Federal Reserve approves first interest rate hike in more than three years, sees six more ahead. *CNBC*. https://www.cnbc.com/2022/03/16/federal-reserve-meeting.html
- European Central Bank. (2015, January 22). *ECB announces expanded asset purchase programme* [Press release]. https://www.ecb.europa.eu/press/pr/date/2015/html/pr150122_1.en.html
- European Central Bank. (2019, September 12). *Monetary policy decisions* [Press release]. https://www.ecb.europa.eu/press/pr/date/2019/html/ecb.mp190912~08de50b4d2.en.html
- European Central Bank. (2020, March 18). *ECB announces €750 billion Pandemic Emergency Purchase Programme (PEPP)* [Press release]. https://www.ecb.europa.eu/press/pr/date/2020/html/ecb.pr200318_1~3949d6f266.en.html
- European Central Bank. (2022a, July 21). *Monetary policy decisions* [Press release]. https://www.ecb.europa.eu/press/pr/date/2022/html/ecb.mp220721~53e5bdd317.en.html
- European Central Bank. (2022b). Recent inflation developments in the United States and the euro area – an update. *ECB Economic Bulletin*. https://www.ecb.europa.eu/press/economic-bulletin/focus/2022/html/ecb.ebbox202201_01~4bb2c93b96.en.html
- European Central Bank. (2026a). *Euro foreign exchange reference rates*. https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/index.en.html
- European Central Bank. (2026b). *Exchange rates* [Data set]. ECB Data Portal. https://data.ecb.europa.eu/key-figures/ecb-interest-rates-and-exchange-rates/exchange-rates
- European Central Bank. (2026c). *Working hours and public holidays*. https://www.ecb.europa.eu/ecb/contacts/working-hours/html/index.en.html
- Foster, S. (2026). Federal funds rate history: 1980 through the present. *Bankrate*. https://www.bankrate.com/banking/federal-reserve/history-of-federal-funds-rate
- Gould, H. (2026, June). What happened at Kevin Warsh's first Fed meeting as chair? 3 key takeaways from the June 2026 FOMC decision. *Chase*. https://www.chase.com/personal/investments/learning-and-insights/article/kevin-warsh-june-2026-federal-reserve-meeting-key-takeaways

- JPMorgan Chase & Co. (2023, May 1). *JPMorgan Chase acquires substantial majority of assets and assumes certain liabilities of First Republic Bank* [Press release]. https://www.sec.gov/Archives/edgar/data/0000019617/000001961723000304/jpmc2023exhibit99.htm
- Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business, 36*(4), 394–419.
- Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics, 14*(1–2), 3–24.
- Pham-Gia, T., & Hung, T. L. (2001). The mean and median absolute deviations. *Mathematical and Computer Modelling, 34*(7–8), 921–936.
- requests contributors. (n.d.). *Requests: HTTP for Humans*. https://requests.readthedocs.io
- Rodini, L. (2026). The Fed's '22–'26 rate hikes & cuts (& what caused them). *The Street*. https://www.thestreet.com/fed/fed-rate-hikes-2022-2023-timeline-discussion
- Swiss National Bank. (2015, January 15). *Swiss National Bank discontinues minimum exchange rate and lowers interest rate to −0.75%* [Press release]. https://www.snb.ch/en/publications/communication/press-releases/2015/pre_20150115
- Swiss National Bank. (2026). *Official interest rates* [Data set]. SNB Data Portal. https://data.snb.ch/en/topics/snb/cube/snboffzisa
- Trading Economics. (2026, January 23). Swiss franc hits 11-year high. https://tradingeconomics.com/switzerland/currency/news/519770
- U.S. Bureau of Labor Statistics. (2022). Exploring price increases in 2021 and previous periods of inflation. *Beyond the Numbers, 11*(12). https://www.bls.gov/opub/btn/volume-11/exploring-price-increases-in-2021-and-previous-periods-of-inflation.htm

Technical Resources

- F5 Networks. (n.d.). *NGINX documentation*. https://nginx.org/en/docs
- Gunicorn. (n.d.). *Gunicorn – WSGI server*. https://gunicorn.org/deploy/
- MathJax Consortium. (n.d.). *MathJax*. https://www.mathjax.org
- Pallets Projects. (n.d.). *Flask*. https://flask.palletsprojects.com
- parallax. (n.d.). *jsPDF*. https://github.com/parallax/jsPDF
- Plotly Technologies Inc. (n.d.-a). *Getting started with Plotly in JavaScript*. https://plotly.com/javascript/getting-started/
- Plotly Technologies Inc. (n.d.-b). *Getting started with Plotly in Python*. https://plotly.com/python/getting-started/
- Plotly Technologies Inc. (n.d.-c). *Kaleido*. https://github.com/plotly/Kaleido
- Python Software Foundation. (n.d.-a). *concurrent.futures — Launching parallel tasks*. https://docs.python.org/3/library/concurrent.futures.html
- Python Software Foundation. (n.d.-b). *datetime — Basic date and time types*. https://docs.python.org/3/library/datetime.html
- Python Software Foundation. (n.d.-c). *io — Core tools for working with streams*. https://docs.python.org/3/library/io.html
- Python Software Foundation. (n.d.-d). *threading — Thread-based parallelism*. https://docs.python.org/3/library/threading.html
- The pandas development team. (n.d.). *pandas documentation*. https://pandas.pydata.org/docs

Declaration of Authorship

I hereby declare that the thesis with title

FX Data Visualization: Web-Based Triangular Analysis of Exchange Rates

has been composed by myself independently and only with the aid referred to in References/Bibliography and that no means other than those declared were used.

In every single case, I have marked parts that were taken out of published or unpublished work, either verbatim or in a paraphrased manner, as such through a quotation.

This thesis has not been handed in or published before in the same or similar form. A review of the work for plagiarism using appropriate software may be made.

Where applicable, the use of artificial intelligence tools has been disclosed in accordance with the relevant guidelines.

Felix Vadakumchery