



yield_cartography

A user guide to yieldcartography.com

What the site is, what it is for,
and how to navigate the seven dashboards.

Marcin Dec, PhD · Kozminski University · FAME|GRAPE

May 13, 2026

Open methodology · publicly available data · daily refresh

1. What this is

yieldcartography.com publishes daily-refreshed sovereign yield-curve estimates, term-premium decompositions, microstructure liquidity diagnostics, and expectations-hypothesis test outputs for less-liquid sovereign bond markets.

The running empirical case is the Polish sovereign bond market, on which the methodology was developed during a five-year research programme (PhD, Kozminski University, and visiting research with FAME|GRAPE Warsaw). The same dashboards display US Treasury and ECB AAA euro-area curves wherever a cross-country comparison is useful, so that the Polish curve can be read against the canonical benchmarks rather than in isolation.

1.1. Why it exists

Less-liquid sovereign bond markets share three features that make standard term-structure machinery awkward to apply directly:

- **Sparse maturity coverage.** Even the deepest Central European panels rarely carry more than 20 to 25 fixed-coupon bonds at a time, with concentrated coupon dates and short benchmark histories.
- **Microstructure distortions.** Switch auctions, primary-market operations, and venue-share migration introduce systematic pricing dislocations at the short end and at the long end.
- **Inference fragility.** The combination of overlapping holding-period returns and strongly persistent forward-rate regressors makes asymptotic HAC inference unreliable, a well-documented *empirical* feature of the present sample and not a theoretical artefact.

Each of these has a technical fix in the academic literature: liquidity-sensitive weighting in the curve fit, an explicit short-end anchor, and small-sample bootstrap inference. The site implements all three for the Polish panel and exposes the outputs interactively, with the underlying CSVs downloadable.

1.2. Who it is for

- Academic researchers in fixed income, monetary economics, and emerging-market finance.
- Central-bank research staff working on local-currency sovereign curves.
- Sovereign-debt analysts at sell-side institutions and asset managers.
- Treasury risk managers needing transparent, replicable curve fits.
- PhD students looking for a clean replication template that runs end-to-end.

The site is free, no registration is required, and every chart's underlying CSV is downloadable from the same page. The full pipeline (data scrapers, NSS fitter, ACM/BRW term-premium models, EH-test scripts) is mirrored in the working papers listed on the about page.

2. Navigation

The site has a fixed seven-tab header. Each tab opens a single-page dashboard that loads its data from a shared `yields.json` cache and exposes a few interactive controls.

Home

/

The landing page. Carries a short prose introduction, three featured-card links into the most recent shorts, and the *data lineage & freshness* block at the bottom that lists the upstream-source timestamp for every CSV feeding the site. If a particular dashboard looks stale, this is where to check.

Yield curves

/curves/

The Nelson-Siegel-Svensson fitted zero-coupon curves on the Polish sovereign panel from 2005 to today, with the same panel for the US Treasury (GSW) and the ECB AAA euro-area benchmarks. Interactive snapshot viewer with bond bubbles overlaid on the fitted curve, one-year-forward overlay, NBP Survey of Professional Forecasters implied path, full PCA factor loadings, and twelve-tenor cross-country comparison.

Term premia

/term-premia/

Adrian-Crump-Moench (2013) five-factor affine term-structure decompositions for the Polish panel, complemented by Bauer-Rudebusch-Wu (2012, 2014) small-sample bias correction. Term premia at 1y, 2y, 3y, 5y, 7y and 10y horizons, with cross-country PL versus US versus EA comparison and two correlation heatmaps.

Liquidity measures

/liquidity/

Six microstructure liquidity diagnostics computed at the bond-day level: bid-ask spread, zero-trading-day frequency, Amihud (2002) illiquidity, Roll (1984) implied spread, Pastor-Stambaugh gamma, and Corwin-Schultz (2012) high-low spread estimator. Aggregated to monthly cross-sectional medians. A composite z -index across the six measures plus the aggregate price-of-liquidity regression table.

Expectations-hypothesis tests

/eh-tests/

Cross-market Fama-Bliss type-1 β heatmaps for PL, US and EA. Macro-spanning regressions that test whether the forward-spot spread carries information beyond observable macro controls. Cochrane-Piazzesi single-factor and seagull-factor decompositions. Out-of-sample Sarno-Thornton-Valente Diebold-Mariano test against a random-walk benchmark.

Shorts

/shorts/

Two-page summaries of each working paper behind the dashboards. Each short carries a catchy title, a ≈ 200 -word setup, one headline figure, one key table, ≈ 200 words of interpretation, and

a 50-word *what this means for practitioners* line at the bottom. Comments are open under every short. Cadence: roughly one new short per fortnight.

Movies

[/movies/](#)

4K time-lapses of the twenty-one years of daily curve dynamics, embedded from YouTube. The *full-history* movie shows the entire 2005-2026 sweep at three trading days per second; the *five-year* movie zooms in on the post-COVID monetary cycle.

About

[/about/](#)

Author bio, methodology summary, external profiles (Google Scholar, SSRN, FAME|GRAPE), full publications list with DOIs, the canonical reference list with hyperlinks, and a contact form that delivers directly to the author's inbox within 48 hours.

3. Dashboard walkthroughs

3.1. Yield curves /[curves/](#)

The *snapshot viewer* (top of the page) lets you pick any trading day from January 2005 to today. The fitted NSS zero-coupon curve is drawn in solid blue, with the on-day bonds plotted as bubbles whose size scales with the bond-level weight used in the fit (outstanding amount, prior-month turnover, or a 50/50 mix). The slider lets you sweep through time at any speed; the keyboard arrows step day-by-day.

The *one-year-forward overlay* switches the y-axis from the zero rate $y(\tau)$ to the instantaneous forward rate $f(\tau)$, which is more visually responsive to monetary-policy expectations near the front end. The *NBP forecasters* overlay adds the survey-implied path from the NBP Survey of Professional Forecasters (computed from the quarterly survey releases).

The *cross-country panel* below the snapshot viewer plots the Polish zero curve on the same chart as the US Treasury GSW Svensson fit and the ECB AAA euro-area Svensson fit, at twelve standard tenors. This is the cleanest way to see where Polish yields sit relative to the major-sovereign benchmarks at any given moment.

The *PCA loadings* chart at the bottom decomposes the daily yield changes into the canonical level, slope and curvature components. The first three factors typically account for >99 percent of the variance across all three markets.

3.2. Term premia /[term-premia/](#)

Two complementary models are reported. *ACM* (Adrian-Crump-Moench, 2013) is the five-factor affine term-structure model estimated by OLS regression of monthly excess returns on lagged Nelson-Siegel principal components. *BRW* (Bauer-Rudebusch-Wu, 2012, 2014) is ACM with iterated-bootstrap small-sample bias correction on the risky-recursion VAR, with common-random-numbers and Kilian shrinkage. On the Polish panel the BRW correction shifts the 10-year term premium upward by about 1.9 basis points on average; the correction is larger on the US and EA panels.

The *cross-country comparison* (PL BRW vs US ACM vs EA ACM) is monthly-resampled to align the three panels on a common end-of-month observation date. The two heatmaps at the bottom show the Pearson correlation of PL term premia at each horizon with US (left) and EA (right) term premia at every other horizon: off-diagonal cells reveal how Polish curvature co-moves with US slope, and so on.

3.3. Liquidity measures /[liquidity/](#)

Six liquidity measures are computed at the bond-day level and aggregated to monthly cross-sectional medians (or means for zero-trading-day frequency and the composite index). The page reports one panel per measure: **BAS** (bid-ask spread in basis points of yield), **ZTD** (zero-trading-day frequency), **Amihud** (price impact per unit of turnover), **Roll** (implied spread from negative serial covariance of returns), **Pastor-Stambaugh** (return-on-flow regression coefficient), and **Corwin-Schultz** (high-low-derived spread).

Each panel shows the cross-sectional median (solid line) plus the 25th-to-75th percentile band (shaded). A *composite z-score* aggregates the six measures into a single illiquidity index, standardised so that 0 corresponds to the long-run mean of the BondSpot panel.

The *price-of-liquidity regression table* at the foot of the page reports the slopes of regressions of bond yields on each of the six measures, with significance codes and Newey-West HAC standard errors.

3.4. Expectations-hypothesis tests /eh-tests/

The three-market *Fama-Bliss type-1 β heatmap* is the headline diagnostic. Each cell of the (horizon, maturity) grid shows the regression coefficient on the forward-spot spread, with rejection of the PEH null $\beta = 0$ at five percent indicated by cell colour. Polish β values are uniformly close to zero or mildly negative, US values are mildly positive at long horizons, and euro-area values turn strongly positive at long horizons.

The *Cochrane-Piazzesi factor regressions* report the single-factor and seagull-factor decompositions of one-year excess returns. The first-stage R^2 is approximately 20 percent on the Polish panel, comparable in magnitude to the US Treasury benchmark, though the joint significance of the return-forecasting factor under the wild block bootstrap of Bauer-Hamilton (2018) is well above conventional rejection thresholds.

The *out-of-sample STV panel* compares the PEH-implied forecast of the future short rate with a naive random-walk benchmark using Diebold-Mariano statistics on an expanding window. At short maturities the random walk beats PEH on 4 of the 20 applicable pairs at five percent; at longer maturities the two forecasts are statistically indistinguishable.

4. Data lineage and freshness

Every observation on the site traces back to one of eleven upstream sources, each refreshed on a different schedule. The *data lineage & freshness* block at the foot of the home page renders the most recent observation timestamp for each source on every build. The major groups are summarised below.

4.1. Polish-side feeds (daily auto-refresh)

- **BondSpot fixings** (price, best bid, best ask, turnover) are scraped from the BondSpot SA venue each evening and appended to a 21-year daily panel.
- **Ministry of Finance disclosures** (monthly turnover by series, auction register, outstanding amounts) are pulled from gov.pl. The publication lag on monthly turnover can be up to two months; the pipeline forward-fills the most recently published value per series so the panel never shows a spurious zero.
- **Coupon schedules** are extracted from the Ministry's *kalkulatorodsetek* reference workbook, which is the canonical Polish coupon-payment calendar.
- **NBP Monetary Policy Council reference rate** (the *stopa referencyjna*) is pulled from the live Polish-Wikipedia table of MPC decisions, which carries every change since 1998 with the corresponding effective date.

4.2. US and euro-area feeds (daily auto-refresh)

- **FRB GSW US Treasury Svensson parameters** are pulled from the Federal Reserve Board canonical CSV at [federalreserve.gov/data/yield-curve-tables/feds200628.csv](https://www.federalreserve.gov/data/yield-curve-tables/feds200628.csv), then re-fitted onto the same twelve-tenor zero-coupon grid as the Polish curve.
- **ECB AAA euro-area Svensson parameters** are pulled from the ECB Data Portal SDMX API endpoint and reduced to the same twelve-tenor zero-coupon grid.

4.3. Pipeline outputs (computed daily)

- LW-NSS fitted Polish zero-coupon curves and the corresponding NSS parameter history are computed by the daily NSS fitter that runs at 18:30 Warsaw time.
- ACM and BRW term-premia panels are computed weekly on Friday closes from the LW-NSS Polish panel.
- The site is built from these inputs and pushed to Cloudflare Pages on every cron completion.

5. Methodology and working papers

The site is the empirical output of three working papers, all openly available with DOIs.

5.1. Closer to New York than to Frankfurt?

The Expectations Hypothesis in Poland, the US and the Euro Area. Working paper, 2026. SSRN [6695444](#).

Applies the five-test PEH battery (Fama-Bliss type-1 and type-2, Thornton conventional and contrarian, Cochrane-Piazzesi) to three sovereign curves on a common 2005–2026 monthly panel. Headline finding: the Polish curve is mildly anti-PEH, the US sits roughly consistent with the PEH null at long horizons, and the euro-area AAA panel rejects strongly under asymptotic Newey-West but barely under the Bauer-Hamilton (2018) wild block bootstrap. The cleanest evidence in our sample of the small-sample inference concern about overlapping-return regressions with persistent regressors.

5.2. Are Survey-Based Rate Expectations Informative?

Evidence from Less-Liquid Markets. Working paper, 2026. SSRN [6644222](#).

Tests whether the NBP Survey of Professional Forecasters carries information beyond what is already in the Polish term structure. Uses ACM and BRW expected-rate paths from the LW-NSS panel as the model benchmark, runs Diebold-Mariano, Clark-West and forecast-encompassing tests at horizons of 1 to 60 months. The model wins at 3y and 5y, the survey wins at 1y, and encompassing tests reject the survey at long horizons.

5.3. Parsimonious Yield Curve Modelling in Less-Liquid Markets

FAME|GRAPE working paper, extended LW-NSS version 2026. grape.org.pl/publications/wps.

Develops the liquidity-weighted Nelson-Siegel-Svensson (LW-NSS) framework for sovereign curves where bond-by-bond observation noise is heterogeneous. Information-matrix derivation gives an explicit weight-matrix structure derived from BondSpot turnover and outstanding amounts. Refit of the Polish panel reduces 21-year mean fit MAE by 1.6 basis points relative to equal-weight NSS while preserving curvature.

6. Data and code

Every chart on the site exposes a *download CSV* button that delivers the data behind the chart in machine-readable form. The CSV schema is documented inline on the about page (</about/>) under *Methodology and data*. Common conventions:

- Dates are ISO-format YYYY-MM-DD (no timezone).
- Yields are continuously-compounded annual decimals ($0.0356 = 3.56\%$).
- Term premia are in basis points ($45.2 \text{ bp} = 0.452\%$).
- Bond identifiers follow the Polish-series convention (two-letter prefix PS, DS, OK etc. plus MMY maturity suffix) on the Polish panel, and ISIN on the cross-country panels.

The pipeline source code (BondSpot scraper, LW-NSS fitter, ACM/BRW estimator, daily orchestrator) is published alongside the working papers. The methodology section of each paper carries the canonical equation form; the code mirrors that exactly so that any reader can replicate the results from the published CSVs without intermediate processing.

7. Contact and collaboration

Direct email. mdec@kozminski.edu.pl. Responses to substantive questions arrive within 48 hours.

Contact form. The *Get in touch* block on the about page delivers to the same inbox with a short subject prefix that helps with sorting.

Collaboration. The framework generalises beyond Poland. Active interest in extending coverage to Central and Eastern Europe (Czech Republic, Hungary, Romania, Slovakia, Bulgaria, Croatia), Latin America (Mexico, Brazil, Chile, Colombia, Peru), Asia (India, Indonesia, Philippines, Thailand), Africa (South Africa, Nigeria, Kenya), and the Middle East. Mutual credit on extensions; co-authorship on the resulting working papers.

Citation. When citing the site for academic work, please use the working-paper DOIs listed above rather than the URL. The DOIs are stable; the site itself is a moving target with new dashboards added every quarter.

This guide is generated from the `yc_site_guide.tex` source in the project repository and is refreshed when the site architecture changes. The site itself is updated daily; this guide every few months. The version number and build date appear in the page footer.