WORKING NOTE · JUNE 2026

Share-faithful reconstruction of bilateral trade from raw UN Comtrade — and a pre-registered nowcast for critical materials

Varcolacus · independent research · public data only · comtrade-reconcile · interactive atlas

Version 2026-06-29 · cite the dated revision in the commit history · trade years 2002–2024 (BACI HS02 + HS17, V202601)

Abstract. I reconstruct bilateral trade in 32 critical raw materials from raw UN Comtrade using a Gaulier–Zignago-style mirror reconciliation (reliability-weighted inverse-variance averaging of exporter and importer reports), and validate it against the official CEPII BACI dataset on the statistics that matter for supply-concentration analysis: top exporters, trade shares, and Herfindahl concentration. On 2024 the reconstruction recovers the top exporter in **25 of 30** materials with a share mean-absolute-error of **3.5pp** and an HHI correlation of **0.92** (importer side: 22/30, 4.2pp, 0.97). I then (i) document an **origin gap** — in **19 of 32** materials the top *exporter* is not the top *miner*, so import-origin statistics systematically overstate the geographic diversification of supply — and (ii) build a provisional 2025 nowcast and a directional 2026 scenario whose shared assumption (share persistence) I test out-of-sample (**85%** year-over-year top-exporter persistence; a leader's annual share move has P50 3.5pp / P90 8.7pp) and whose accuracy I **pre-register** against the next BACI release. All inputs are public; the raw Comtrade is committed, so all code and fixtures reproduce the reconciliation and validation end-to-end with no API key. Finally I extend the reconciled series back to **2002** and, rather than merely plotting concentration as the literature does, **test** it: **9 of 32** materials show a statistically significant rising export-concentration trend (Mann–Kendall, FDR-corrected), with structural breaks clustering in **2012–2016**.

1. The question

Under customs rules, refining and substantial transformation confer "origin": when a country imports an ore, processes it and re-exports, the traded good takes the processor's nationality and the customs record stops there. Trade statistics therefore measure *where a material was last shipped from*, not *where it came out of the ground*. For critical-material supply-risk analysis — where the policy question is genuine geographic concentration of the upstream — this is a systematic blind spot. This note builds the data to measure it and quantifies how large the gap is.

2. Data and vintages

Table 1 — sources and exact vintages. All public.

LAYER	SOURCE	VINTAGE / COVERAGE
Reconciled bilateral trade (reference)	CEPII BACI	HS17 release **V202601** (Jan 2026), years 2018–2024
Raw bilateral trade (reconstruction input)	UN Comtrade API	annual + monthly, reporter×partner, pulled 2025–26
Mine production shares	USGS Mineral Commodity Summaries	latest published (approx.)
Refining / processing shares	IEA Critical Minerals Outlook	latest published (approx.)
Reserves	USGS Mineral Commodity Summaries	economically recoverable (approx.)
Gravity covariates (CIF/FOB)	CEPII `dist_cepii`	distance, contiguity
Commodity prices (2026 tilt)	World Bank Pink Sheet	monthly, Q1-2026 vs Q1-2025

The 32 materials are mapped to HS6 product codes; three of them — gallium, germanium and hafnium — share a single HS6 code (811292) and so carry *identical* trade columns that cannot be separated. This is flagged wherever those materials appear.

3. Reconciliation method

For each product and country pair *(i, j)* two reports exist: the exporter's free-on-board value x_{ij} and the importer's cost-insurance-freight value m_{ji}. Following Gaulier & Zignago (2010), I (a) strip an estimated CIF/FOB markup from import values, (b) weight each reporter by an estimate of its reliability, and (c) average the two reports as an inverse-variance weighted mean in logs.

CIF/FOB markup

The markup is recovered from the within-product mirror ratio. On the 32-material slice a gravity regression of log(m/x) on distance and contiguity is not identified ($R^2 \approx 0.01$) — BACI estimates this on the full ~5,000-product universe — so I fall back to a robust per-product median markup, the same quantity in reduced form.

Reliability weights (variance components)

Let d_{ij} = log x_{ij} − log m_{ji} be the mirror discrepancy. Modelling it as the sum of a reporter-i and reporter-j error, an OLS regression of d^2 on reporter dummies yields each reporter's error variance v_i (a variance-components estimator). The reconciled value averages the two reports with weights inversely proportional to those variances:

$$\log t_{ij} = (w_i \cdot \log x_{ij} + w_j \cdot \log \tilde{m}_{ji}) / (w_i + w_j), \quad w_i = 1 / v_i$$

where $\tilde{m}$ is the import value after removing the CIF/FOB markup. The result is a single reconciled bilateral matrix per product per year — the same object BACI produces, built here from the raw reports.

---

4. Validation against BACI

The reconstruction is validated on the statistics the atlas actually uses — shares, ranks, concentration — not on absolute levels.

Table 2 — reconstruction vs official BACI, both sides of the market.

YEAR · SIDE	TOP-1 EXPORTER/IMPORTER CORRECT	SHARE MAE	HHI CORRELATION
2024 · exporter	25 / 30	3.5 pp	0.92
2024 · importer	22 / 30	4.2 pp	0.97
2022 · exporter	22 / 30	3.9 pp	0.89

On levels. Reconstructed totals run ~1.5–1.8× BACI. This is *not* a reconciliation artefact: current raw Comtrade already sits ~1.8–1.9× above BACI's published snapshot before any reconciliation (Comtrade absorbs continuous revisions; BACI applies additional outlier/quality filtering this note does not replicate), and the reconstruction sits *between* the two raw mirror reports. A settled year (2022) still shows ~1.5×, so the offset is partly persistent filtering, not only vintage. The claim is therefore explicitly **share-faithful**, not level-faithful — which is the relevant property for concentration analysis.

5. Nowcast and pre-registration

BACI lags ~18 months. For the years it has not released I extend the series:

- **2025 (provisional).** The full reconciliation applied to the partial 2025 Comtrade that has been filed (~half of reporters by mid-2026), with levels calibrated per material to BACI 2024.
- **2026 (directional).** Only ~Q1 monthly Comtrade exists, so 2025's reconciled structure is carried forward and only *levels* are tilted, per material, by reporter-matched Q1 export momentum blended with the World Bank Pink Sheet price change. Shares are held at 2025: this is a directional scenario, not a bilateral measurement.

Both rest on one assumption — that last year's trade structure predicts this year's. I test it directly on the measured 2018–2024 series, scoring year $T-1$'s shares as a nowcast for year T:

Table 3 — out-of-sample persistence (192 material-years, 2018–2024).

METRIC	VALUE
Top-exporter unchanged year-over-year	85%
Mean share MAE across exporters	0.37 pp
Leader's annual share move — median (P50)	3.5 pp
Leader's annual share move — P90	8.7 pp

The P50/P90 are the nowcast's empirical uncertainty band, shown on the interactive year slider. Most-predictable materials: magnets, coking coal, manganese, niobium; least: arsenic, beryllium, hafnium, gallium, germanium, fluorspar (thin or non-reporting-producer markets).

Pre-registration. Before BACI 2025 exists, I have committed (`PREREGISTRATION.md`, in git history at the relevant commit) to how the frozen `flows_2025.json` will be scored when it releases — `validate.py 2025`, same metrics — with numeric thresholds (≥ 22/30 top-1, ≤ 5pp share MAE, ≥ 0.88 HHI correlation; scored on the 30 validatable materials) and a commitment to publish the result pass or fail. The bet is written down first; the test is falsifiable on a fixed date by one command.

6. Finding: the origin gap

With reconciled trade and the mine-production layer side by side, define, per material in a given year:

$$\textit{origin gap} = \textit{(top exporter's share of world trade)} - \textit{(that same country's share of}$$

world mine output)

A large positive gap means a country sells far more than it mines — it processes or trans-ships someone else's ore. The headline result, on 2024: in **19 of 32** materials the top exporter is not the top miner; in **4 of 32** a country mining under 5% of world supply exports more than a quarter of it.

Table 4 — largest origin gaps, 2024 (where exporter ≠ miner).

MATERIAL	TOP EXPORTER	EXPORTS	IT MINES	GAP	LEAD MINER
Beryllium, unwrought	Kazakhstan	89%	0%	+89	United States (65%)
Strontium carbonate	Germany	60%	0%	+60	Spain (30%)
Lithium carbonate	Chile	75%	20%	+55	Australia (37%)
Bauxite / aluminium ore	Guinea	72%	25%	+47	Australia (27%)
Cobalt oxides & hydroxides	Finland	29%	0%	+29	DR Congo (76%)
Tantalum, unwrought	United States	22%	0%	+22	DR Congo (40%)
Nickel, unwrought	Norway	18%	0%	+18	Indonesia (60%)

The twist. The intuitive story is "China hides behind refineries." The data only half-supports it: for many materials China is *both* the lead miner and the lead exporter — its position is largely genuine, not an artefact. The materials where exporter and miner diverge are instead fronted by industrial refiners and entrepôts — Finland for Congolese cobalt, Japan for largely Chinese-mined titanium, Germany for Spanish strontium, Norway for Indonesian nickel. The corrective therefore cuts both ways: it deflates apparent dependence on refiner countries, and it reveals that genuinely concentrated upstreams (Congo cobalt, Indonesian nickel) are *more* concentrated than the diversified-looking trade ledger suggests.

7. Trends, breaks and tested concentration (2002–2024)

The reconciled series is extended back to **2002** using CEPII BACI's HS2002 release, spliced to HS17 from 2017. The 32 materials' HS6 codes are stable across that window with two documented exceptions: boron, whose code was merged in HS2012 (recovered by summing its two predecessor codes), and hafnium, pooled inside 811292 before HS2022 (as gallium and germanium remain). This yields a continuous 22-year, share-faithful panel — long enough to *test*, not merely plot, how concentration evolved. Most critical-minerals work stops at plotting Herfindahl trajectories; here each material's export-HHI, China-share and origin-gap series is put through standard trend statistics:

- **Monotonic trend.** The Mann–Kendall rank test (Mann 1945; Kendall 1975) for significance and the Theil–Sen estimator (Sen 1968) for a robust slope, with Benjamini–Hochberg (1995) false-discovery-rate correction across the 32 materials.
- **Structural breaks.** The Pettitt (1979) non-parametric change-point test dates the year a series shifts regime (cf. Bai & Perron 1998 for multiple breaks).
- **Decomposition.** The 2002→2024 change in export-HHI is attributed by country (contribution = $share^2_{2024} - share^2_{2002}$), identifying which exporter drove each material's concentration.

The result: **9 of 32** materials show a statistically significant rising export-concentration trend (FDR<0.05); their structural breaks **cluster in 2012–2016**, the period in which export-control measures and critical-minerals-security policy intensified. The steepest are beryllium, tungsten, bauxite and rare-earth magnets. The aggregate **origin gap** — averaged across materials — is wide and somewhat higher at the end of the period (~15→~18pp), but the index steps at the **2017 HS-vintage join**, so this is read as trade-side drift, not a clean trend — an HS17-only confirmation is the proper test. Because USGS mine shares enter as a current fixed reference, this isolates how the *trade* map drifted from today's mining geography; it is a trade-side, not a geological, trend. I did not find a peer-reviewed study applying this combined Mann–Kendall / Theil–Sen / Pettitt screening to critical-material export-concentration panels; formal trend and change-point inference on these series appears uncommon in the policy literature. The contribution is the testing, not the plotting.

---

8. Further analytical layers

Four additional lenses are built on the same reconciled trade, each deliberately framed as a *screening indicator*, not a structural claim:

- **Trade-routing centrality.** Treating each material's bilateral flows as a directed network, betweenness and PageRank (Freeman 1977) plus a node-removal fragility test surface the brokers and processing hubs that share-concentration misses. Read as routing centrality — betweenness on a trade network also reflects import size, so a large buyer scores high without "controlling" anything.
- **Governance-weighted criticality.** An EU/SCRREEN-shaped supply-risk proxy (Blengini et al. 2017) weighting concentration by World Bank Worldwide Governance Indicators, plus a Graedel-style supply-risk axis (Graedel et al. 2012). These are proxies on public approximations, not the official EU or Yale scores; they correlate with the project's transparent index at $\rho \approx 0.79$ — partly mechanically, since they share inputs.
- **Economic complexity.** Balassa (1965) revealed comparative advantage over the 32-material export matrix — specialization *within* critical materials, not economy-wide RCA — with material ubiquity and relatedness (Hidalgo & Hausmann 2009).
- **Origin trace.** A first-order re-attribution of refiner-fronted imports to the dominant mine, reported explicitly as an *upper bound* on single-origin concentration.

9. Limitations and non-claims

- The origin gap is a **measurement gap** between two public datasets (reconciled trade shares vs USGS mine shares), *not* a fully traced physical supply chain. Some gap is legitimate — a refiner genuinely exports a *different product* than the mined ore (cobalt chemicals vs cobalt ore), which sits under a different code.
- Customs data cannot separate a true refiner from a re-export hub; trade-derived country roles are *trade exposure*, not proof of physical processing. Entrepôts are flagged.
- gallium / germanium / hafnium share one HS6 code and cannot be separated in trade.
- Figures are share-faithful, not level-faithful; absolute values are not claimed. Mine and refine layers are approximate USGS/IEA reference shares, undated. 2025 is provisional; 2026 is directional.

10. Reproducibility

Two public repositories, **no API key required to reproduce the pipeline end-to-end**: the interactive atlas (Varcolacus/critical-materials-atlas) and the reconciliation engine (Varcolacus/comtrade-reconcile, CI green). The raw 2024 Comtrade is committed (gzipped) under `fixtures/raw/`, so the reconciliation regenerates from raw rather than from a pre-computed file. On a fresh clone of the engine: `pip install -r requirements.txt`, then `ATLAS_ROOT=fixtures python reconcile.py 2024` (regenerates the reconciliation), `… validate.py 2024` (Table 2), `python backtest.py` (Table 3), `python findings.py` (Table 4). CI runs the `reconcile → validate` chain on every push. The Comtrade key is needed only for the initial network pull and is read from the environment, never committed.

References

Data.

- Gaulier, G. & Zignago, S. (2010). *BACI: International Trade Database at the Product-Level*. CEPII Working Paper 2010-23.
- UN Comtrade Database. United Nations Statistics Division.
- U.S. Geological Survey. *Mineral Commodity Summaries* (annual).
- International Energy Agency. *Global Critical Minerals Outlook* (2025).
- World Bank. *Worldwide Governance Indicators*; *Commodity Markets ("Pink Sheet")*.

Methods.

- Mann, H.B. (1945). Nonparametric tests against trend. *Econometrica* 13(3).
- Kendall, M.G. (1975). *Rank Correlation Methods*. Griffin.

- Sen, P.K. (1968). Estimates of the regression coefficient based on Kendall's tau. *JASA* 63(324).
- Pettitt, A.N. (1979). A non-parametric approach to the change-point problem. *Applied Statistics* 28(2).
- Bai, J. & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica* 66(1).
- Benjamini, Y. & Hochberg, Y. (1995). Controlling the false discovery rate. *J. Royal Stat. Soc. B* 57(1).
- Balassa, B. (1965). Trade liberalisation and "revealed" comparative advantage. *Manchester School* 33.
- Hidalgo, C.A. & Hausmann, R. (2009). The building blocks of economic complexity. *PNAS* 106(26).
- Freeman, L.C. (1977). A set of measures of centrality based on betweenness. *Sociometry* 40(1).
- Graedel, T.E. et al. (2012). Methodology of metal criticality determination. *Environ. Sci. Technol.* 46(2).
- Blengini, G.A. et al. (2017). EU methodology for establishing the list of critical raw materials. *Resources Policy* 53.

Recent critical-mineral trade-network & criticality (context).

- Nassar, N.T., Brainard, J. & Gulley, A. et al. (2020). Evaluating the mineral commodity supply risk of the U.S. manufacturing sector. *Science Advances* 6(8).
- Schrijvers, D. et al. (2020). A review of methods and data to determine raw material criticality. *Resources, Conservation & Recycling* 155.
- Liu, Y., Ouyang, X., Gao, J. et al. (2024). Multilayer material-flow and trade network of cobalt. *Environ. Sci. Technol.* 58.
- Guo, Q. & Wang, C. (2024). Evolution of the rare-earth trade-dependence network. *Resources Policy* 91.
- Wang, X. et al. (2019). The graphite trade network. *Resources Policy* 60.