Concentration, Chokepoints, and Supply-at-Risk: A Structural Stress Test of Critical Mineral Vulnerability, 2015-2026

A reproducible quantitative research report

Author: Akshay Byagathvalli

8 August 2026

DATA	1,548 rows | 24 minerals | 35 countries
METHODS	Concentration | Network | Simulation | Econometrics | ML
SCOPE	Synthetic-panel structural stress test

The source data dictionary identifies principal quantitative fields as simulated. Numeric findings characterize the supplied panel and declared stress scenarios; they are not official production statistics or geopolitical forecasts.

Contents

Author: Akshay Byagathvalli

Research report: Critical-minerals structural vulnerability project

Data access date: 8 August 2026

Reproducibility seed: 42

Scope note. *The supplied data dictionary explicitly identifies mine production, reserves, and prices as simulated. This paper is therefore a reproducible structural stress test and methodological research study of the supplied synthetic panel. Its numeric rankings are findings about that panel, not audited estimates of observed world production or geopolitical event probabilities.*

Abstract

Critical-mineral security depends not only on how much material is produced but also on where production is located, how persistent dominant suppliers are, and how alternative producers connect across mineral markets. This study asks how geographic concentration evolved from 2015 through 2026 in a 1,548-row country-mineral-year panel covering 24 minerals and 35 producer countries, and whether concentration, production volatility, growth, reserves, and country-mineral network structure explain and predict structural supply-disruption exposure. Because the source dictionary labels key values as simulated, the analysis is a structural stress test rather than a historical audit. Production shares and all concentration metrics were independently recomputed from tonnage after supplied shares failed adding-up checks in 268 of 288 mineral-years. The study integrates HHI, concentration ratios, entropy, effective producer counts, fixed-effects regression, bipartite network centrality, targeted node removal, a six-component Critical Mineral Vulnerability Index (CMVI), principal components, 10,000-draw Monte Carlo stress regimes, temporal machine learning, SHAP, and 10,000 alternative index-weight vectors. Mean HHI fell by 0.014 ($p = 0.017$), with 16 of 24 minerals diversifying and eight concentrating. Yet concentration remained extreme for several minerals: Gallium's 2026 HHI was 0.900, implying only 1.111 effective producers. Across six ranking families, Gallium, Terbium, Dysprosium, Tungsten, and Germanium were most vulnerable; Nickel, Copper, Manganese, Tantalum, and Tin were least vulnerable. China led 15 minerals and had a normalized systemic score of 1.000. A 50% dominant-producer shock removed 31.2% of median mineral supply. Gallium's medium-regime CVaR95 reached 61.6%. Linear Regression predicted next-year structural CVaR with test MAE 0.008, R-squared 0.994, and rank correlation 0.995; network exposure, CR1, and HHI dominated SHAP importance. The central result is that modest panel-wide diversification coexists with a small, stable set of highly concentrated mineral structures whose tail exposure is amplified by a shared systemic country node.

1. Introduction

Critical minerals sit upstream of batteries, electricity networks, permanent magnets, semiconductors, catalysts, aerospace systems, and many other advanced technologies. Their strategic significance follows from a conjunction of essential use, limited substitutability, production constraints, and exposure to disruption. The clean-energy transition makes this conjunction more consequential because material demand can rise faster than mining, processing, substitution, recycling, and permitting systems adapt. The World Bank estimated that a below-2-degree energy transition would require very large additions of mineral and metal production, while the International Energy Agency (IEA) has repeatedly emphasized that apparently adequate aggregate supply can still be vulnerable when extraction or refining is concentrated in a few jurisdictions (Hund et al., 2020; IEA, 2025).

Volume alone is therefore a poor definition of security. A mineral can have rapidly growing output yet remain structurally fragile if one country controls most production. Conversely, a smaller or slower-growing market can be resilient when supply is distributed among independent producers. The key analytical objects are shares, alternative sources, temporal persistence, and the network position of supplier countries. A country that is dominant in one mineral creates a commodity-specific bottleneck; a country that is important across many minerals is a systemic node whose disruption propagates through multiple technology chains. This distinction motivates a research design that combines conventional industrial-organization concentration measures with entropy, network science, stress simulation, and temporal prediction.

Authoritative policy has increasingly adopted the same structural framing. The European Union's Critical Raw Materials Act identifies high supply concentration in third countries as a core source of risk and sets a 2030 diversification benchmark under which no single third country should supply more than 65% of annual consumption of a strategic raw material at a relevant processing stage (European Parliament and Council, 2024). The IEA's 2025 outlook reports that top-three mining and refining shares remain high for major energy-transition minerals and shows why the loss of a leading supplier can expose a system even when total projected supply appears sufficient. USGS commodity summaries provide the commodity-specific production and reserve data that an observational implementation would ordinarily require (USGS, 2026). Those sources establish the practical importance of the question; they are not used to retrofit or validate the synthetic values analyzed here.

This paper addresses the primary question: *How has geographic concentration in the supplied critical-mineral production panel evolved from 2015-2026, and to what extent can concentration, volatility, growth, diversification, reserves, and producer-network structure explain and predict vulnerability to supply disruptions?* Eleven secondary questions operationalize that inquiry: the level and direction of concentration; the minerals diversifying or concentrating most; the growth-concentration and volatility-concentration relationships; systemic producer countries; deterministic and stochastic losses; temporal predictability; feature importance; and robustness to alternative vulnerability definitions.

The contribution is methodological and empirical within the supplied panel. First, it subjects the source to a formal audit rather than accepting precomputed shares and risk scores. Second, it distinguishes the number of listed producers from the effective number implied by their shares. Third, it uses a country-mineral bipartite network to identify cross-commodity chokepoints. Fourth, it combines transparent equal weights with PCA, simulation, machine learning, and broad weight sensitivity rather than claiming that any single index is objective. Fifth, it uses information at year t to predict a structural stress target at t+1 with whole-year train, validation, and test blocks, preventing the most common form of panel time leakage. Finally, it uses rank aggregation across method families so the ultimate ordering depends on triangulation rather than a preferred model.

The evidence is decisive within that scope. The panel's average concentration declined modestly, but the change was heterogeneous and did not eliminate extreme single-country dependence. Gallium, Terbium, Dysprosium, Tungsten, and Germanium occupied the five highest consensus positions. Gallium combined a 94.8% leading-country share, HHI near 0.900, network dependence, and the largest simulated tail loss. Terbium and Dysprosium were also consistently extreme across concentration, CMVI, PCA, network, Monte Carlo, and ML ranks. Nickel, Copper, Manganese, Tantalum, and Tin were comparatively diversified and suffered smaller standardized losses. China was not simply the dominant producer for one commodity: it led 15 of 24 minerals in 2026, giving it a cross-mineral systemic role far beyond the next-ranked countries. The top three remain strongly supported under alternative concentration measures, Monte Carlo regimes, network thresholds, and index weights; the fourth and fifth consensus positions are more weight-sensitive.

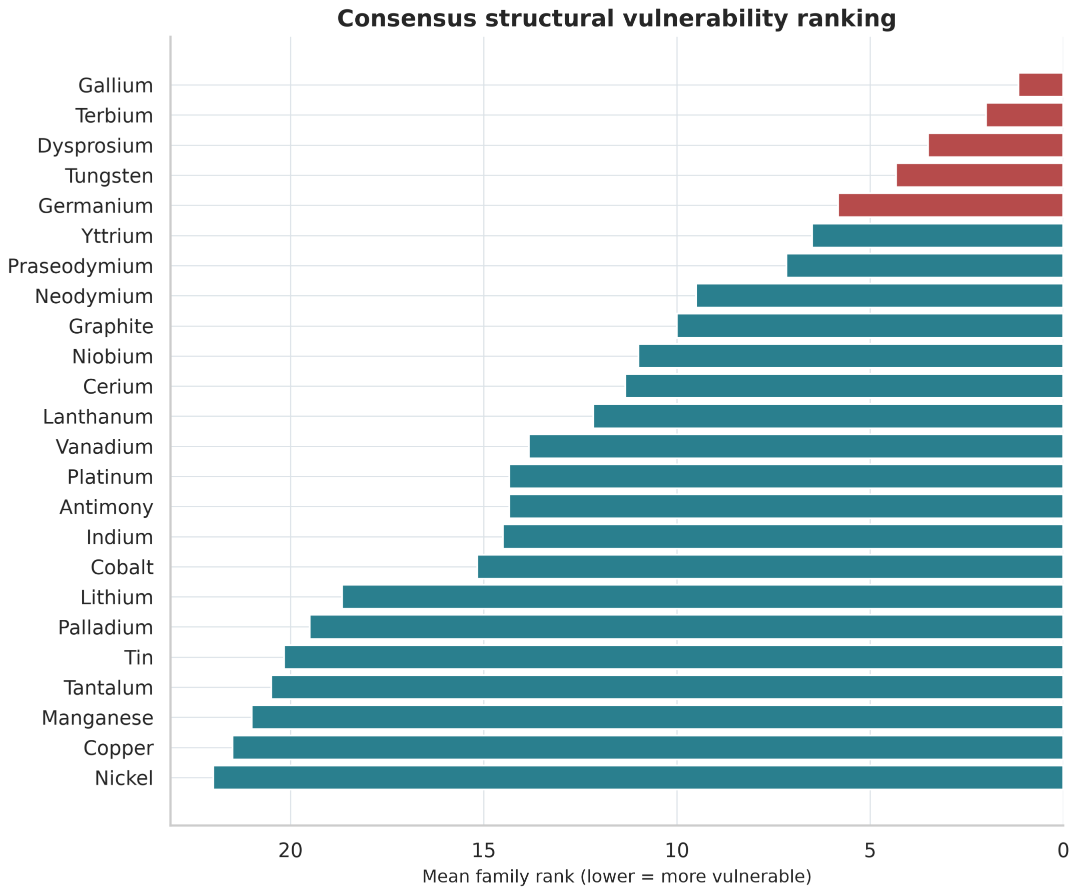

Figure 31. Consensus structural vulnerability ranking. Borda-style mean rank aggregates six method families to avoid reliance on one index or stress assumption.

2. Background and Related Literature

2.1 Criticality and supply-chain resilience

Material criticality is multidimensional. Graedel et al. (2015) formalized criticality using supply risk, environmental implications, and vulnerability to supply restriction, illustrating that no single production statistic fully characterizes exposure. Nassar, Du, and Graedel (2015) applied related logic to rare earth elements and showed that vulnerability can vary with substitution and national context even among elements sharing geological occurrence. Nassar, Graedel, and Harper (2015) added by-product dependence: when a metal is recovered mainly as a companion to another host metal, price signals for the companion may not elicit a proportionate supply response. The supplied panel does not contain substitution elasticities, recycling, co-production, or technology-specific demand, so this study narrows its claim to geographic and network structure while retaining growth, volatility, and reserve proxies.

Resilience is the capacity to maintain function under disruption, not merely the absence of past disruption. In a mineral context, redundancy, geographic diversity, inventories, substitution, recycling, and responsive capacity can all contribute. This project observes only mine-level country production, a normalized refining-share proxy, reserves, annual growth, and a small set of contextual fields. It can measure structural redundancy and stress losses, but not downstream inventory buffering, bilateral trade rerouting, price-mediated demand destruction, or substitution. The term *vulnerability* is therefore used precisely: the share of supply exposed to a standardized structural shock and the concentration/network features associated with that exposure.

The synthetic source contains a precomputed risk score and disruption labels. Those fields are not used in the CMVI or primary predictive model. Doing so would reward a model for reconstructing the source generator's own risk formula and would create a circular appearance of validation. Instead, the study defines a new structural stress outcome from independently recomputed shares and explicit shock assumptions. This choice sacrifices a superficially realistic event classifier in favor of a target whose interpretation and information timing can be audited.

2.2 Concentration measures

The Herfindahl-Hirschman Index is the sum of squared market shares. It increases when share shifts toward large producers and ranges from 1/N under equal shares to one under monopoly. Competition authorities use HHI as a screening measure for market structure, although supply security is not identical to product-market competition (DOJ and FTC, 2023). Here HHI is used descriptively: a high value means production is distributed as if only a small number of equal-size sources existed. The reciprocal, 1/HHI, converts that abstraction into an effective producer count. A mineral may list ten producing countries but have an effective count near two if the remaining producers contribute little.

Concentration ratios preserve a different intuition. CR1 is the leading producer's share; CR3 and CR5 measure the combined share of the three and five largest producers. CR1 maps directly to dominant-producer stress losses, while CR3 captures dependence on a small coalition. HHI is more sensitive to the full distribution. Agreement between them strengthens a vulnerability inference; disagreement is scientifically informative. Niobium illustrates potential disagreement in this panel: it is extremely concentrated in one producer but scores lower on cross-mineral network dependence, reducing its consensus rank relative to its pure concentration rank.

Shannon entropy supplies a complementary diversification measure (Shannon, 1948). Entropy increases as shares spread more evenly and can be normalized by log(N) to range from zero to one for a fixed observed producer count. HHI and inverse normalized entropy are not algebraic duplicates; they weight the distribution differently. Their rank correlation of 0.982 in the latest year is therefore strong evidence that the main concentration ordering is not peculiar to one formula.

2.3 Networks and systemic nodes

Country-mineral production is naturally bipartite: countries connect to minerals but countries do not directly connect to countries in the raw structure. Network analysis preserves this two-mode organization (Newman, 2010). Degree counts how many mineral relationships a country has above a threshold; weighted degree sums its shares across minerals. Betweenness identifies nodes that lie on many shortest paths; PageRank and eigenvector centrality reward connection to important neighbors. None is sufficient alone. A broad low-share producer can have high degree but limited immediate loss impact, while a dominant single-mineral producer can have high share yet little cross-mineral reach.

This project combines centralities with transparent dominance counts and then validates importance through node removal. Thresholds of 1%, 5%, and 10% assess whether conclusions depend on including small edges. The network visualization uses 5% for legibility, but the principal weighted metrics use all edges at or above 1%. Removing the most central nodes and comparing their effects with random removals gives centrality an operational interpretation: a valid systemic ranking should identify countries whose removal creates greater average cross-mineral loss than removal of equally many random countries.

2.4 Stress simulation and predictive modeling

Historical disruptions are rare, heterogeneous, and endogenous to policy and markets. Structural stress tests answer a different question: given the observed share architecture, how much supply would be lost under a standardized outage? Deterministic scenarios are transparent but sparse. Monte Carlo adds a distribution of magnitudes and multi-country events, allowing value-at-risk and conditional-value-at-risk analogues. SupplyVaR95 is the loss exceeded in the worst 5% of simulations; SupplyCVaR95 is the mean loss within that tail. These statistics do not predict the frequency of geopolitical conflict. They compare the loss geometry created by different supply structures under common declared regimes.

Machine learning is used to test whether year-t features contain information about year-(t+1) structural tail exposure. Linear and regularized models establish interpretable baselines; tree ensembles allow interactions and nonlinear thresholds (Breiman, 2001). SHAP decomposes held-out predictions into feature contributions under a game-theoretic framework (Lundberg and Lee, 2017). Because the target is itself a simulated structural outcome, high accuracy demonstrates that future stress exposure is persistent and encoded in lagged structural indicators. It does not demonstrate prediction of real political events. That distinction is central to interpreting the very high out-of-sample R-squared obtained below.

3. Data

3.1 Source and scope

The Kaggle dataset *Critical Minerals and Rare Earths (2015-2026)* was accessed on 8 August 2026 from https://www.kaggle.com/datasets/sergionefedov/critical-minerals-and-rare-earths-20152026. Three files were supplied: a 1,548-row master panel, a 24-row mineral reference table, and a 19-row data dictionary. The master covers 24 minerals, 35 source country labels, and 12 annual periods from 2015 through 2026. Each year has 129 country-mineral rows and all 24 minerals, making mineral-year coverage balanced. Producer counts range from three to nine. There are no duplicate country-mineral-year keys and no missing primary production, reserve, price, or share values. The sole missing field is `disruption_next_year` in 2026, exactly as expected for a forward label.

The source dictionary calls mine production and reserves “simulated” and price a “simulated mineral price.” The reference table reproduces configured producer counts and average top-country shares. The study consequently does not claim that the 2026 panel is a forecast or that the 2015-2025 paths reproduce official statistics. Contextual comparisons to IEA, USGS, World Bank, or EU reports explain why concentration matters; they do not authenticate these numeric values.

3.2 Audit findings and cleaning

The data audit found no negative production or reserve values and no zero production values. It did find internal share inconsistencies. Supplied production shares failed to sum to 100% within 268 of 288 mineral-years using a 0.25 percentage-point tolerance; totals ranged from 93.27% to 106.65%. One supplied Gallium share exceeded 100%. Recomputing HHI from supplied shares differed from the supplied HHI by 0.024 on average and by as much as 0.117. Refining shares failed the same adding-up check in 287 of 288 groups and ranged in total from 75.91% to 129.27%.

No row was deleted. Mine production share was recomputed as country tonnage divided by mineral-year total tonnage. Reserve share used the same operation on reserves. Because refining tonnage was unavailable, supplied refining percentages were divided by their mineral-year sum and explicitly labeled a *normalized refining proxy*. Supplied production share, HHI, top-country share, risk score, and disruption fields were retained under `_source` names for audit comparison. Country labels were standardized for display and network joining: DRC became Democratic Republic of the Congo; USA became United States; and concatenated South Africa, South Korea, and New Caledonia names were separated. The original label remains in `country_source`.

Programmatic validation requires production, reserve, and refining-proxy shares to sum to one within numerical tolerance; HHI, entropy, and CMVI to remain within bounds; CR1 <= CR3 <= CR5; effective producers <= actual producers; target year = feature year + 1; and all simulated losses between zero and 100%. These assertions convert methodological decisions into executable contracts.

3.3 Units and transformations

Mine production and reserves are in metric tonnes, but tonnage scales differ radically across minerals. The paper never ranks minerals by raw tonnes. Production trends are indexed to 2015 = 100 and growth is calculated within mineral. Prices are reported as source-simulated U.S. dollars per tonne but are not a primary vulnerability component because cross-mineral price levels reflect fundamentally different products. Reserve-to-production ratios are dimensionless years-of-supply proxies. Country shares, concentration ratios, loss rates, and CMVI values are proportions unless displayed as percentages.

The principal tidy table has one row per country-mineral-year. Derived mineral-year, country-year, and country-mineral tables are generated separately. The mineral-year panel contains global production, global reserves, recomputed shares and concentration, annual and rolling growth, rolling volatility, acceleration, dominant-country change and persistence, reserve ratios, and yearly network dependency. Market entry and exit features were evaluated, but the balanced positive-production source has no producer entry or exit; that absence is reported rather than replaced with manufactured turnover.

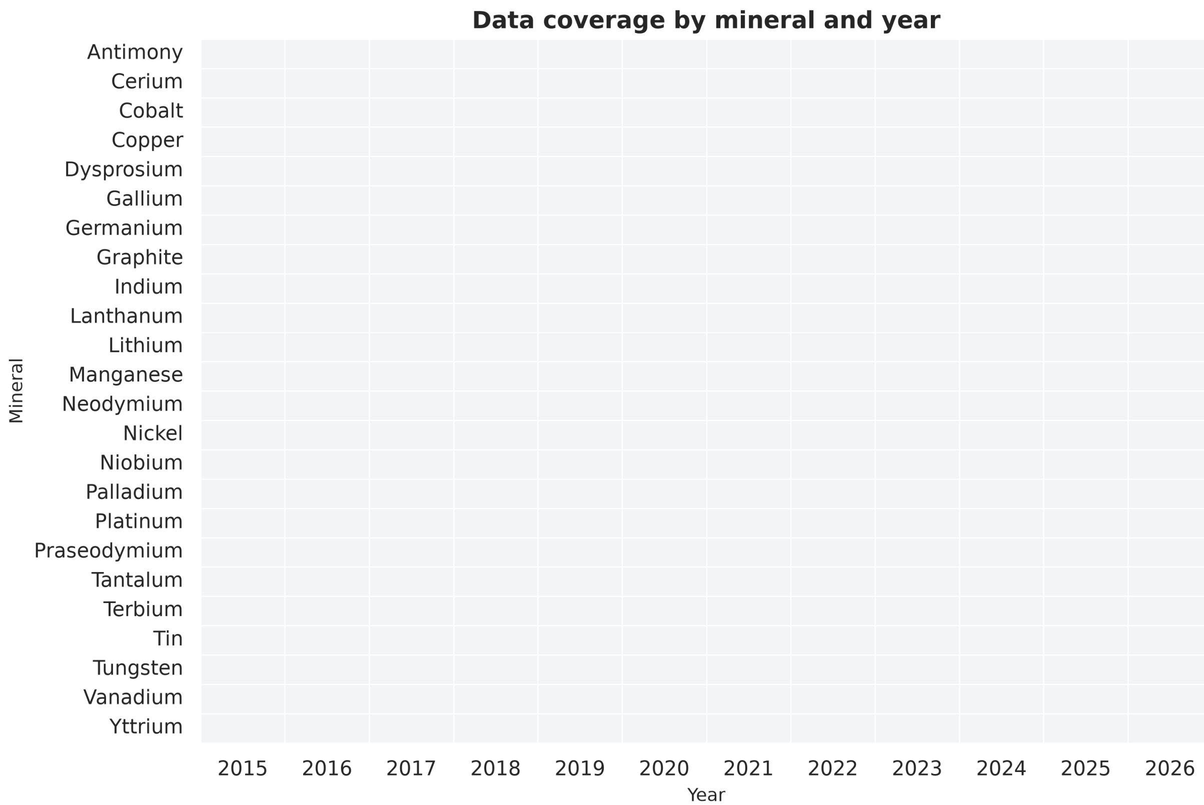

Figure 1. Data coverage by mineral and year. Availability in the balanced synthetic panel. A filled cell indicates at least one producer record.

4. Methods

4.1 Production shares and concentration ratios

For country c, mineral m, and year t, production share is:

`s_cmt = P_cmt / sum_c(P_cmt)`.

CR1, CR3, and CR5 are the sums of the largest one, three, and five shares. These ratios answer RQ1 and map directly to standardized producer-removal scenarios. Because some minerals have fewer than five producers, CR5 equals one when all producers are included.

4.2 HHI and effective producers

HHI is `HHI_mt = sum_c(s_cmt^2)`. Values are reported on 0-1 and 0-10,000 scales. Normalized HHI adjusts the equal-share lower bound: `(HHI - 1/N)/(1 - 1/N)` for N > 1. The effective number of producers is `N_eff = 1/HHI`. This transformation is especially useful for interpretation: Gallium's latest-year effective count is only 1.111, even though four producer countries appear in the panel.

4.3 Entropy and diversification

Shannon entropy is `H_mt = -sum_c(s_cmt ln(s_cmt))`. Normalized entropy divides H by ln(N), producing a diversification score between zero and one for N > 1. HHI emphasizes large shares; entropy gives more relative weight to smaller producers. Rankings are compared with Spearman and Kendall correlations.

4.4 Growth, volatility, reserves, and temporal change

Annual growth is `(P_mt - P_m,t-1)/P_m,t-1`, with nonpositive denominators treated as missing. CAGR uses the eleven intervals between 2015 and 2026. Rolling growth and sample standard deviation use three-year windows. Mineral-level volatility is the sample standard deviation of annual growth; robust dispersion is the median absolute deviation. Downside volatility uses negative-growth years when at least two exist. The coefficient of variation summarizes the production level path within mineral.

Temporal concentration change is 2026 HHI minus 2015 HHI. The five qualitative labels—strongly diversifying through strongly concentrating—are empirical quintiles of that distribution, not subjective policy thresholds. Continuous deltas remain primary. Dominant-country persistence is the expanding fraction of observed years led by the current leading country. Reserve vulnerability combines reserve-share HHI with the inverse reserve-to-production ratio.

4.5 Statistical hypotheses and bootstrap uncertainty

Four hypotheses were specified before inspecting inferential results. Hypothesis A tests association between CAGR and latest HHI using Pearson and Spearman correlations. Hypothesis B tests association between production-growth volatility and mean 2015-2026 concentration. Hypothesis C tests paired 2015-versus-2026 HHI using both a paired t-test and Wilcoxon signed-rank test; Shapiro-Wilk checks the paired-difference distribution. Hypothesis D splits CAGR at the cross-mineral median and compares HHI using Mann-Whitney U and Cliff's delta. All tests are two-sided. Effect sizes, p-values, and 95% intervals are reported whether significant or not. Five thousand nonparametric bootstrap resamples estimate intervals for correlations, mean and median HHI change, and Cliff's delta. The mineral, not the producer row, is the resampling unit for cross-mineral hypotheses.

4.6 Fixed-effects panel regression

The panel model is:

```
HHI_mt = beta1 Growth_mt + beta2 RollingVolatility_m,t-1 + beta3 DemandGrowth_mt + beta4
ReserveHHI_mt + alpha_m + gamma_t + error_mt.
```

Mineral fixed effects alpha_m control time-invariant mineral characteristics; year fixed effects gamma_t absorb common annual shifts. Standard errors are clustered by mineral with a finite-sample correction. The source's producer count is constant within each mineral and is therefore perfectly absorbed by mineral fixed effects; it is deliberately excluded rather than reported from a pseudo-inverse solution. A second fixed-effects specification predicts simulated structural stress at t+1 from HHI, CR1, rolling volatility, network dependency, and reserve HHI measured at t. VIF, Jarque-Bera, Breusch-Pagan, fitted-versus-residual, and mineral/year residual diagnostics are retained. The design estimates conditional associations, not causal effects.

4.7 Country-mineral network and dominance

The weighted bipartite graph links a country to a mineral when its production share exceeds a threshold. Principal network calculations use 1%; 5% and 10% variants test stability. Country degree, weighted degree, betweenness, PageRank, and eigenvector centrality are normalized. The Country Mineral Dominance Score equally averages normalized minerals produced, top-one count, top-three count, weighted share across minerals, and historical leadership persistence. A systemic score averages network centrality, dominance score, and weighted degree. Equal combination is transparent and no claim is made that the weights estimate social welfare.

Node-removal experiments remove the top one, two, and three systemic countries. For each removal, the analysis calculates each mineral's lost share, the mean/median/maximum loss across minerals, remaining mineral coverage, and the count losing at least 20%. One thousand random removals of equal cardinality provide a baseline. Since tonnes are not comparable across minerals, systemic loss is the average of within-mineral percentage losses rather than a sum of unlike tonnages.

4.8 Critical Mineral Vulnerability Index

The equal-weight CMVI has six min-max normalized components: mine concentration (HHI), dominant dependence (CR1), production volatility, positive growth pressure (CAGR position), network dependency, and reserve vulnerability. The baseline is:

```
CMVI_equal = (C + D + V + G + N + R) / 6.
```

Quartiles define LOW, MODERATE, HIGH, and CRITICAL classes. Equal weights are not declared objective; they are a reproducible reference. PCA on standardized components provides a data-driven alternative (Jolliffe and Cadima, 2016). PC1 is sign-oriented to correlate positively with equal-weight vulnerability and rescaled to 0-1. An outcome-calibrated weight set is not treated as confirmatory because fitting weights to an algebraically generated shock outcome would create circular validation.

4.9 Deterministic and stochastic supply shocks

Eight deterministic scenarios reduce the largest producer by 10%, 25%, 50%, 75%, or 100%; reduce the top two or top three by 50%; or remove the top two completely. Loss is the sum of affected producer shares multiplied by shock magnitude. Surviving shares are renormalized to compute post-shock HHI and effective producers.

The Monte Carlo model uses 10,000 simulations for every mineral-year under low, medium, and high structural-stress regimes. The common base event rates are 3%, 8%, and 15%. Producer event propensity is multiplied by 0.5 plus normalized cross-mineral centrality and capped at 75%. Conditional magnitudes follow Beta(2,3). A correlated systemic switch occurs with half the base event rate, selects a country proportional to centrality plus a small floor, and draws severity from Beta(3,2). These parameters are scenario assumptions, not observed disruption probabilities. Their purpose is comparative: identical rules expose how different share and network architectures convert shocks into losses. Metrics are expected loss, median, VaR95, VaR99, CVaR95, and probabilities of loss exceeding 10%, 20%, and 30%.

4.10 Leakage-safe machine learning and explainability

The target is next-year medium-regime SupplyCVaR95. Features at t include HHI, CR1, CR3, producer count, effective producers, entropy, current and rolling growth, rolling volatility, acceleration, concentration change, dominance persistence, reserve HHI, reserve years, and two network-exposure measures. Each row satisfies target_year = feature_year + 1. Source risk and disruption labels are excluded.

Target years through 2022 form training data (n = 168); 2023-2024 form validation (n = 48); and 2025-2026 form untouched test data (n = 48). Expanding-year folds tune tree hyperparameters only in training. Mean, median, and previous-year target baselines, Linear Regression, Ridge, Lasso, Elastic Net, Random Forest, Extra Trees, Gradient Boosting, and HistGradientBoosting are compared using MAE, RMSE, R-squared, median absolute error, and Spearman ranking correlation. The selected validation model is refit on training plus validation and evaluated once on test. Thirty-repeat held-out permutation importance, SHAP, partial dependence, and residual diagnostics explain and audit the best compatible tree model.

4.11 PCA, clustering, weight sensitivity, and consensus

K-Means, Ward agglomerative clustering, and Gaussian mixtures are evaluated for k = 2 through 6 using silhouette, Calinski-Harabasz, and Davies-Bouldin metrics. The highest silhouette candidate is selected before interpreting groups. PCA supplies the common two-dimensional map.

Weight sensitivity draws 10,000 vectors from a uniform Dirichlet distribution over the six CMVI components. Each mineral receives median rank, 5th-95th percentile rank, and probabilities of appearing in the top three, top five, or top quartile. Robustness also compares HHI with entropy, CR1 with CR3, equal CMVI with PCA, three network thresholds, three Monte Carlo regimes, alternative shock magnitudes, raw versus winsorized volatility, min-max versus percentile normalization, and two temporal test-start years.

The final consensus uses six families, each contributing one rank: a concentration composite (HHI, CR1, inverse entropy); equal CMVI; PCA; network dependency; Monte Carlo CVaR95; and ML-predicted next-year CVaR. The Borda-style mean family rank is primary and median rank is retained. Grouping correlated concentration indicators into one family prevents counting essentially the same evidence several times.

5. Results

5.1 Global production structure (RQ1)

The latest-year panel is highly heterogeneous. Gallium is most concentrated: HHI 0.900, CR1 94.8%, and only 1.111 effective producers. Niobium follows on HHI at 0.799, then Terbium (0.750), Dysprosium (0.694), and Tungsten (0.688). These are not marginal distinctions. Gallium's share distribution behaves nearly like a monopoly even though four countries are listed.

At the other end, Tin has HHI 0.184 and 5.437 effective producers. Copper, Nickel, and Manganese have HHI between 0.184 and 0.197; their normalized entropy values are near 0.89-0.91. The distinction between actual and effective producers is still material, but no single producer dominates those structures to the same degree.

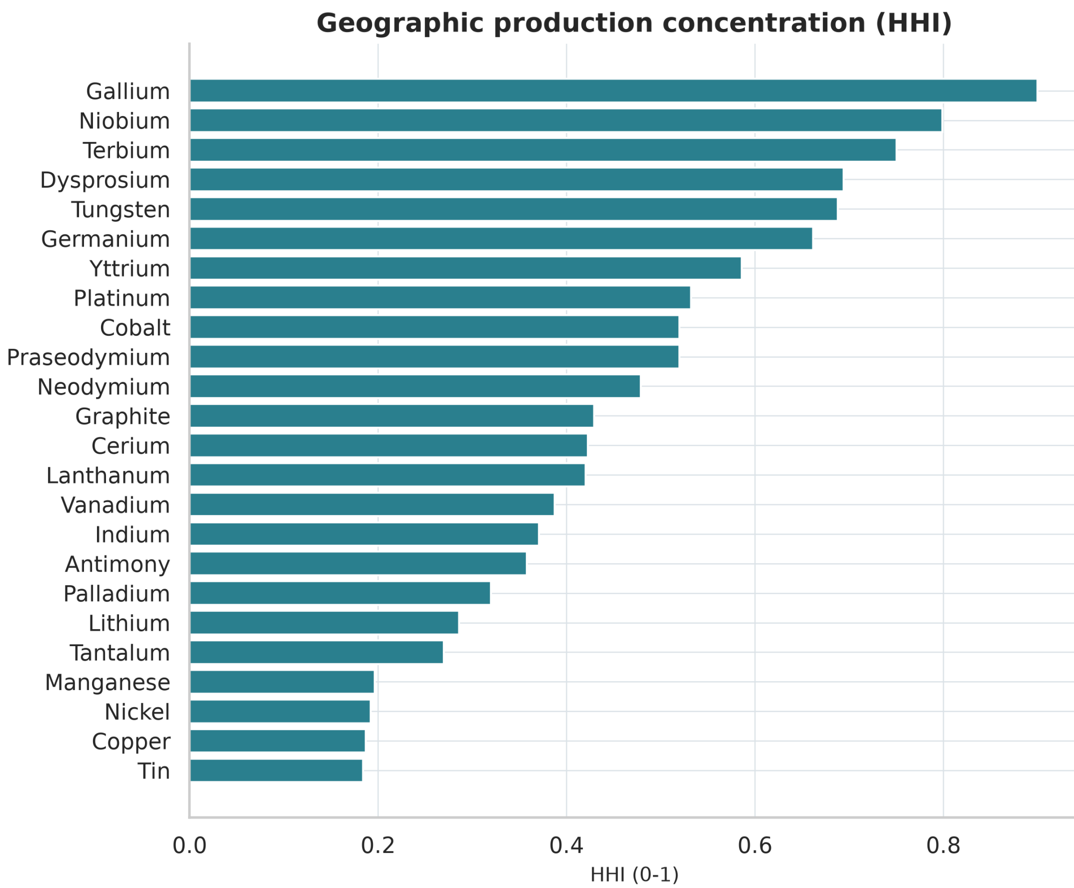

Figure 7. Geographic production concentration (HHI). HHI is computed from normalized mine-production shares; higher values imply fewer effective producers.

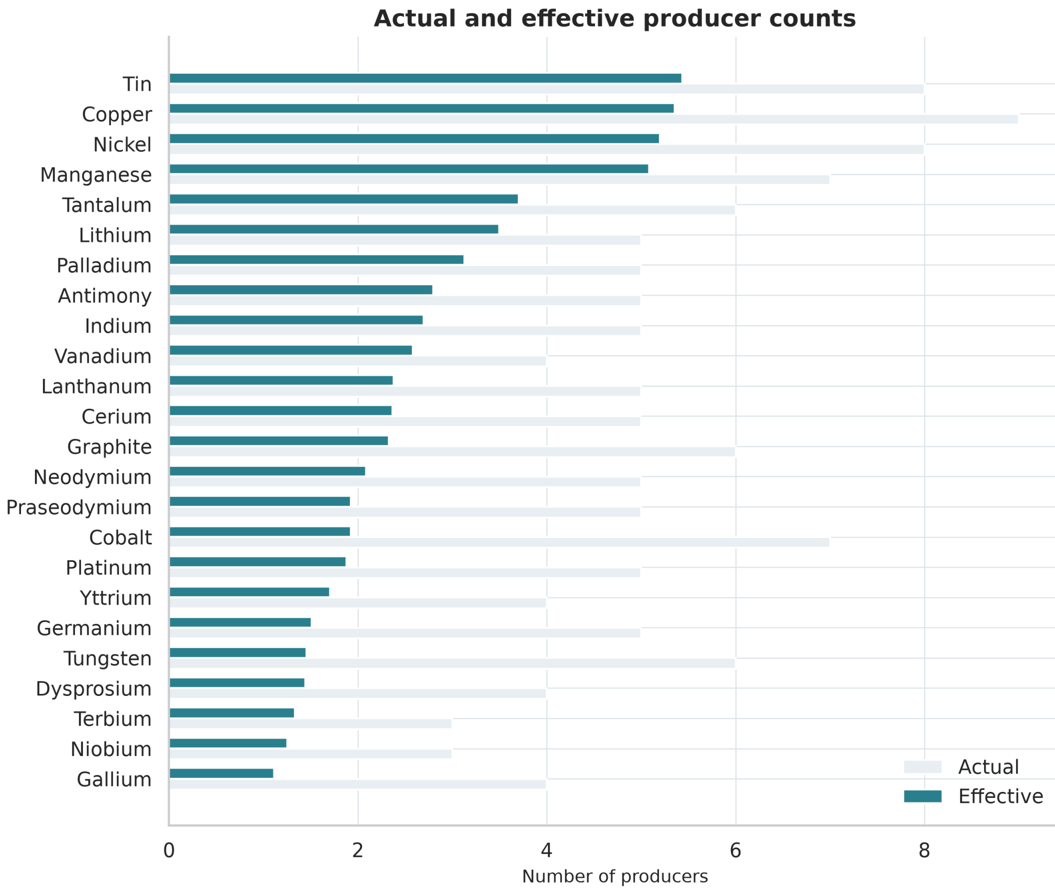

Figure 10. Actual and effective producer counts. Effective producers equals 1/HHI and can be far below the reported producer count.

5.2 Evolution of geographic concentration (RQ2-RQ4)

From 2015 to 2026, mean HHI fell by 0.014 and median HHI fell by 0.010. Sixteen minerals diversified while eight concentrated. The paired t-test rejects no mean change (`t = -2.579`, p = 0.017, mean-change bootstrap 95% CI [-0.024, -0.003]). The Wilcoxon result also rejects a zero-centered change (`W = 72.000`, p = 0.025). The two tests agree that the panel diversified modestly on average.

Average direction conceals mineral-level divergence. Neodymium shows the largest HHI decline (-0.066), followed by Praseodymium, Yttrium, Dysprosium, and Lanthanum. Cobalt shows the largest increase (0.034), followed by Vanadium, Copper, Germanium, and Tungsten. Importantly, diversification does not imply low vulnerability: Dysprosium's HHI declined by 0.046 yet remained 0.694, the fourth-highest latest concentration. A favorable trend can start from such an extreme base that the endpoint remains fragile.

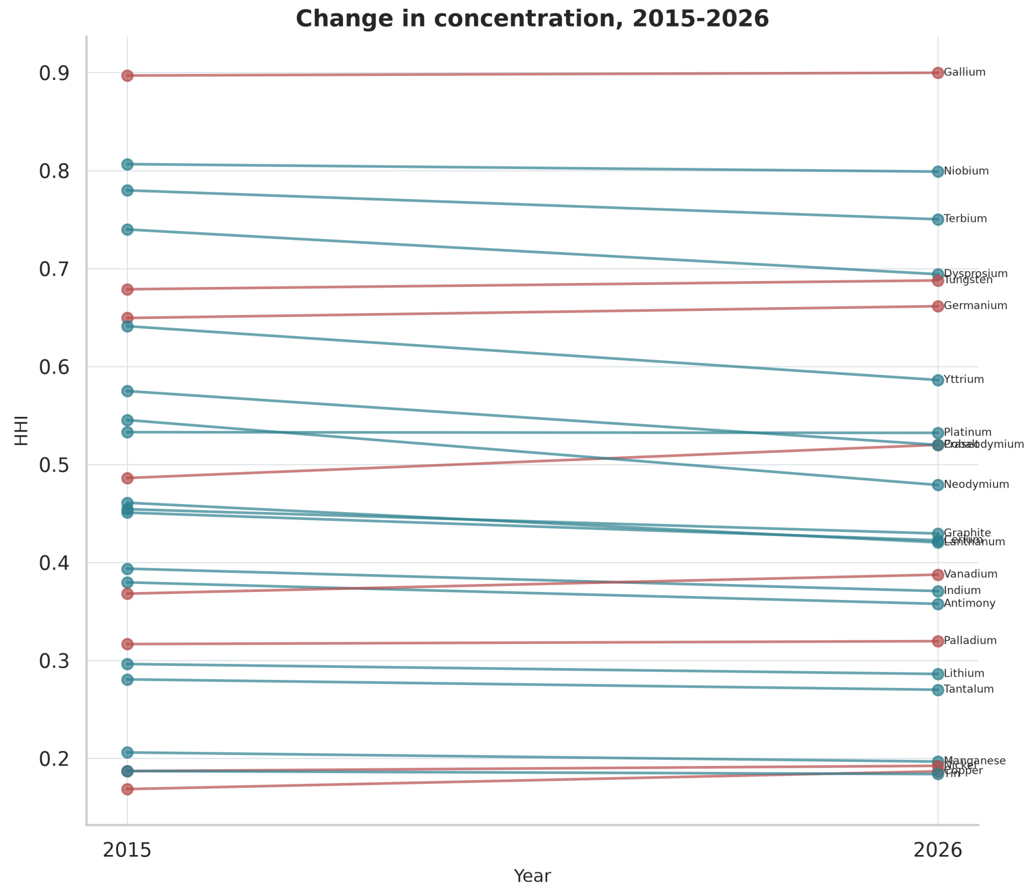

Figure 9. Change in concentration, 2015-2026. Lines connect recomputed 2015 and 2026 HHI; upward movement indicates concentration.

5.3 Growth, volatility, and concentration (RQ5-RQ6)

The cross-mineral data do not support a simple claim that rapidly growing markets are more concentrated. CAGR and HHI have Pearson r = -0.315 (p = 0.134, bootstrap 95% CI [-0.593, 0.048]) and Spearman rho = -0.162 (p = 0.450). The signs are negative, not positive, and intervals include zero. Median-split high-growth and low-growth groups likewise differ little: Mann-Whitney U = 60.000, p = 0.507, Cliff's delta = -0.167. Heterogeneity dominates any monotonic growth-concentration pattern.

Volatility has a moderate positive association with mean-period HHI. Pearson r = 0.385 (p = 0.063) does not reach the 5% threshold, while Spearman rho = 0.410 (p = 0.047) does. The Pearson bootstrap interval is [0.103, 0.633] and the rank interval is [0.040, 0.679]. Both intervals are positive, but the split between linear and rank-test p-values makes the evidence moderate and specification-sensitive rather than uniformly strong.

The fixed-effects model separates within-mineral changes from cross-sectional differences. Growth is positively associated with HHI conditional on mineral and year effects (`beta = 0.066`, clustered SE 0.010, p < 0.001). Lagged rolling volatility is `beta = -0.030` (p = 0.413); demand growth is `beta = 0.000` (p = 0.440); reserve HHI is `beta = 0.007` (p = 0.575). A separate temporally aligned stress regression relates predictors at t to CVaR at t+1; its HHI coefficient is -0.034 (p = 0.845) and its network-dependency coefficient is 0.339 (p = 0.031). These are within-panel associations in a synthetic generator, not evidence that growth or concentration causes observed disruptions.

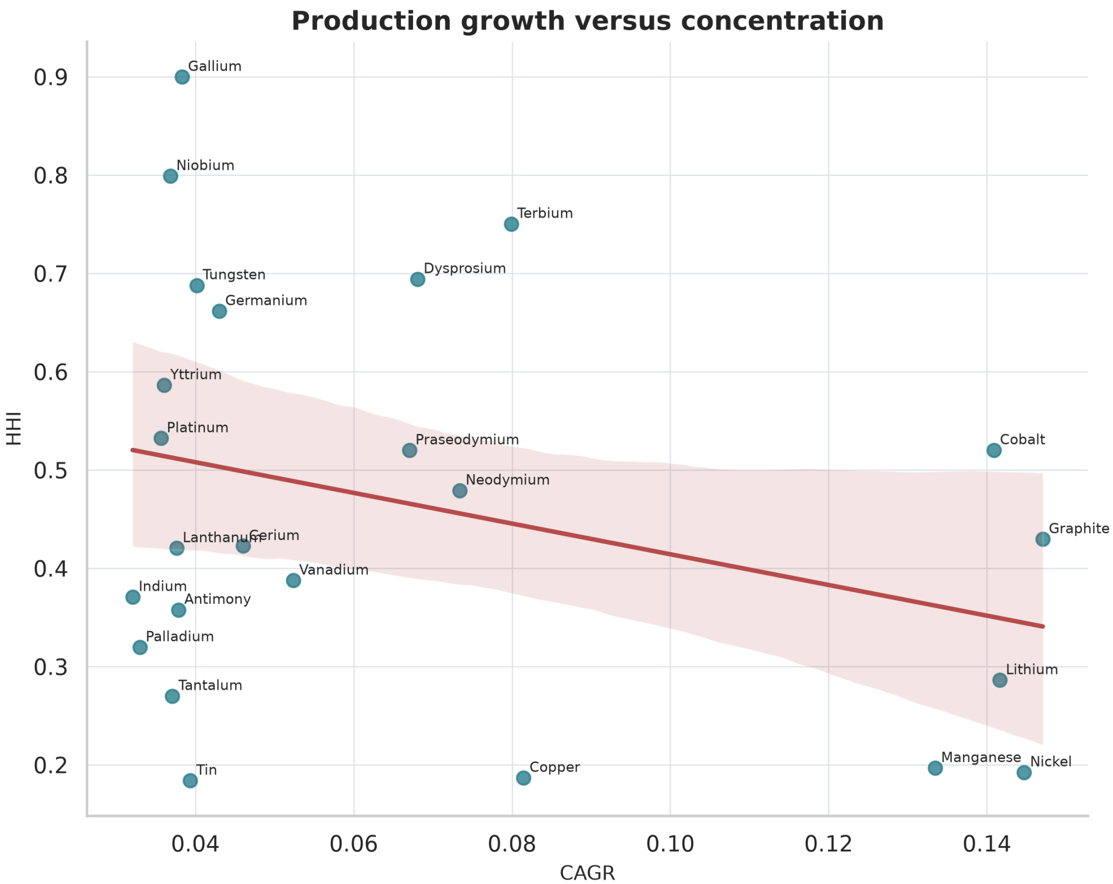

Figure 11. Production growth versus concentration. The line is an OLS descriptive fit; inference is reported separately and is not causal.

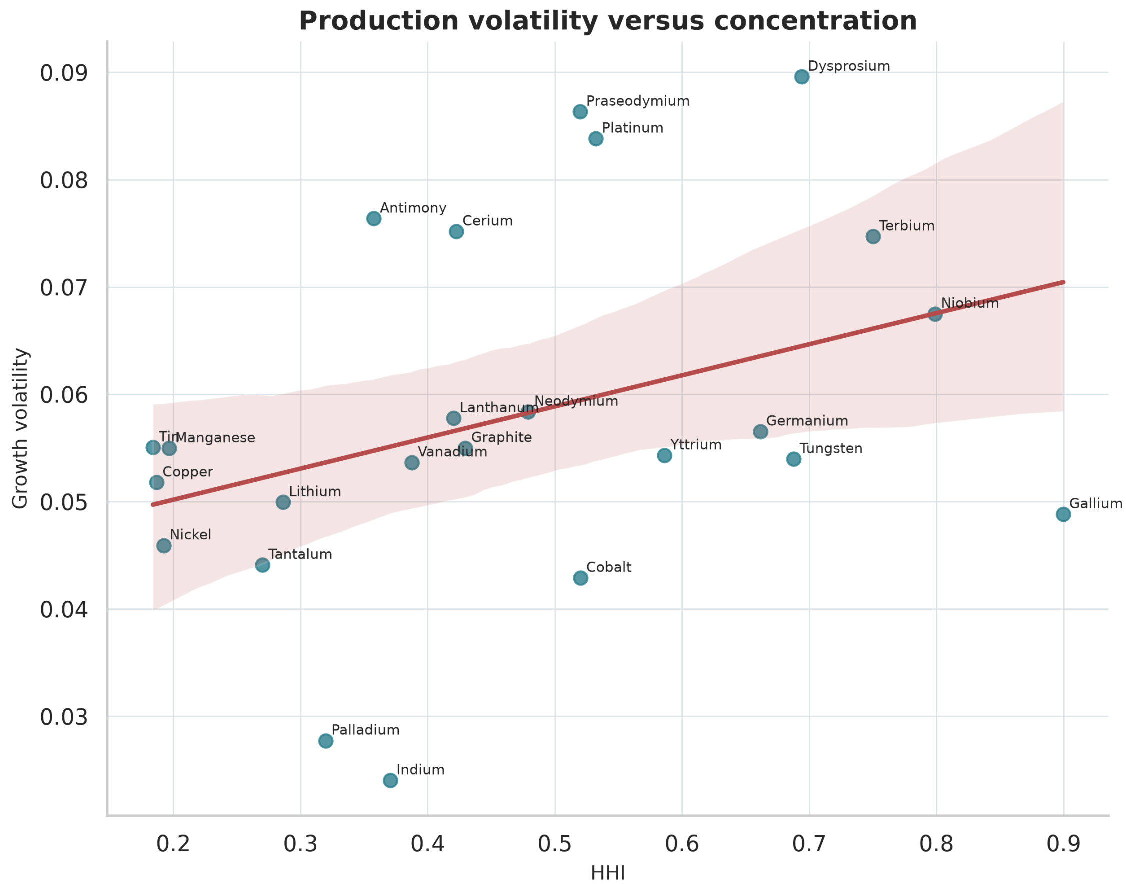

Figure 12. Production volatility versus concentration. Volatility is the standard deviation of annual within-mineral production growth.

5.4 Systemic producer countries (RQ7)

China is the dominant systemic chokepoint by a wide margin. Its normalized systemic score is 1.000, weighted degree is 10.655, and it is the leading producer for 15 of 24 minerals. Russia ranks second at 0.216, followed by Australia (0.197), South Africa (0.174), and the Democratic Republic of the Congo (0.150). The large China-to-second gap is present because China combines high shares with breadth; other countries may dominate important individual minerals but connect to fewer markets.

Threshold checks preserve the principal conclusion. China ranks first at 1%, 5%, and 10% edge thresholds. South Africa, Russia, Australia, Brazil, and the Democratic Republic of the Congo rotate within the remaining top positions as smaller edges are removed. This is moderate evidence about their exact order but robust evidence about China's unique systemic position.

Node removal gives the ranking a supply interpretation. Removing China alone produces mean loss of 44.4% across minerals, versus 2.8% for a random-country removal. 15 minerals lose at least 20%. Targeted top-two and top-three removals widen the gap relative to equal-cardinality random baselines. The result is not driven merely by network aesthetics; high centrality corresponds to large cross-mineral production exposure.

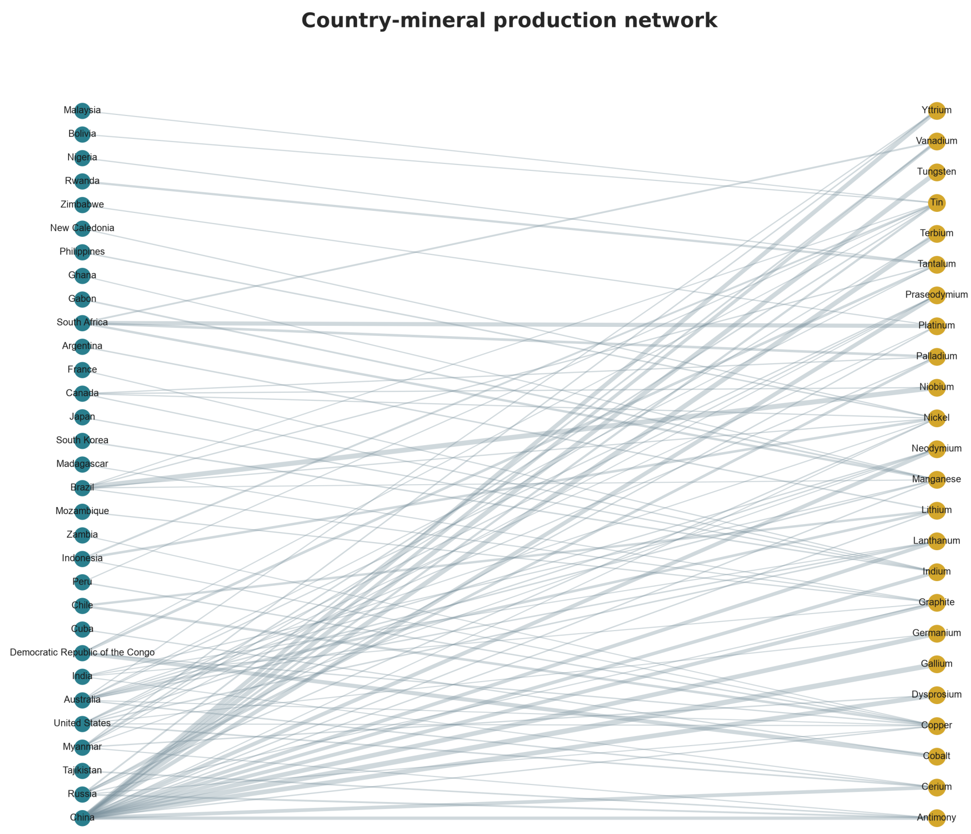

Figure 15. Country-mineral production network. Edges retain producers with at least 5% share to avoid a hairball; width encodes production share.

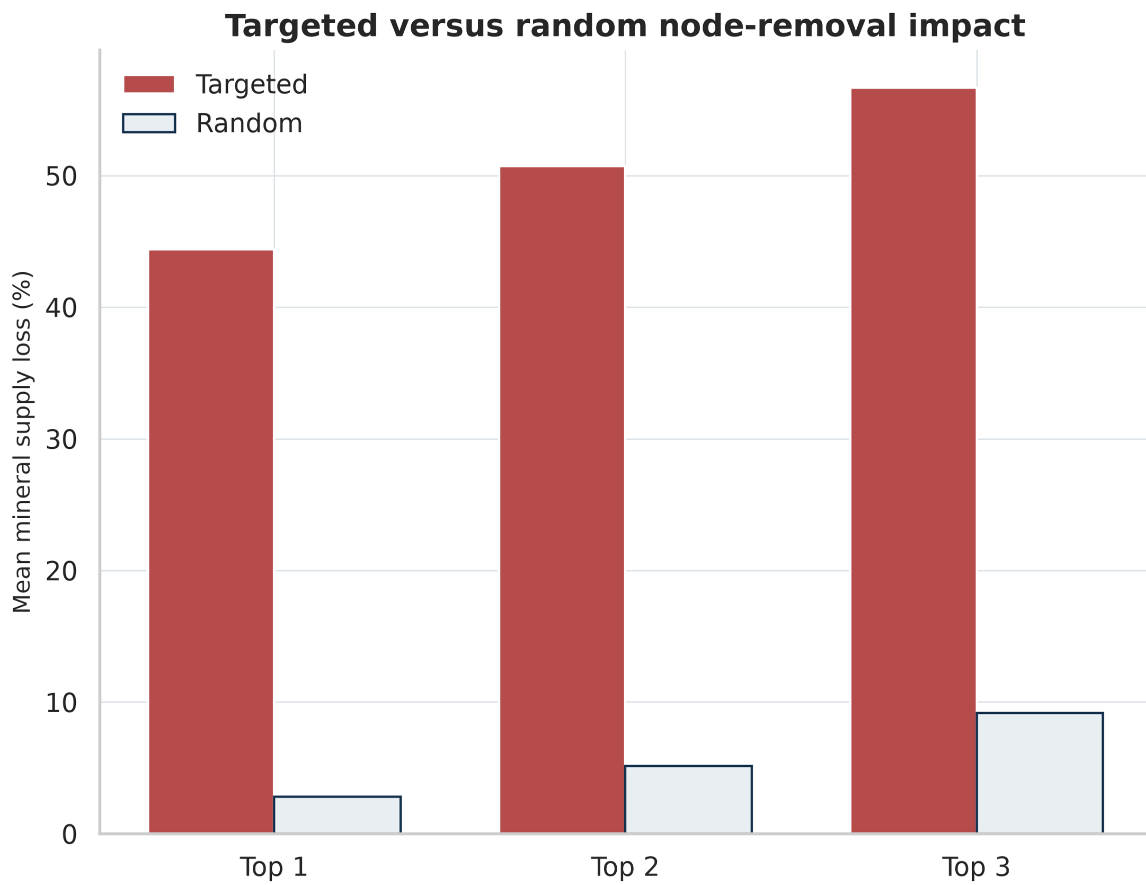

Figure 17. Targeted versus random node-removal impact. Mean mineral supply loss from removing the most systemic countries versus equal-size random removals.

5.5 Composite vulnerability and structural archetypes (RQ10-RQ11)

The equal-weight CMVI and PCA ranking correlate at rho = 0.854 and Kendall tau = 0.667. PC1 explains 59.5% of standardized component variance; PC2 adds 15.3%. Agreement is strong enough to identify a common vulnerability axis but not perfect, which is expected because equal weighting preserves all six components while PCA maximizes variance.

The best clustering solution is kmeans with k = 2. With this selected number of groups, the partition separates relatively high combined concentration/network exposure from lower-exposure structures; it should not be read as evidence for more archetypes than the fitted k permits. Labels were assigned only after fitting, so the analysis does not force preconceived “concentrated and volatile” categories. Cluster membership is descriptive and secondary to continuous scores.

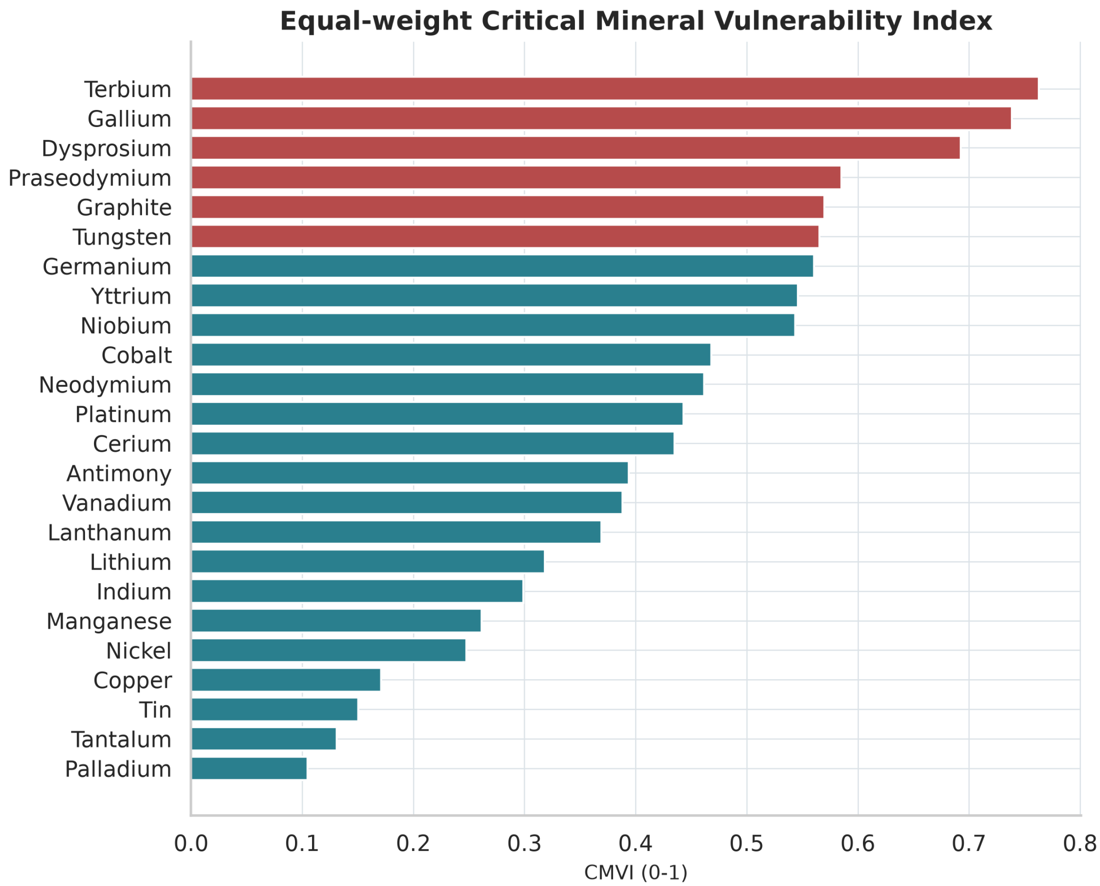

Figure 18. Equal-weight Critical Mineral Vulnerability Index. The equal-weight CMVI averages six normalized structural components; it is a transparent index, not an event probability.

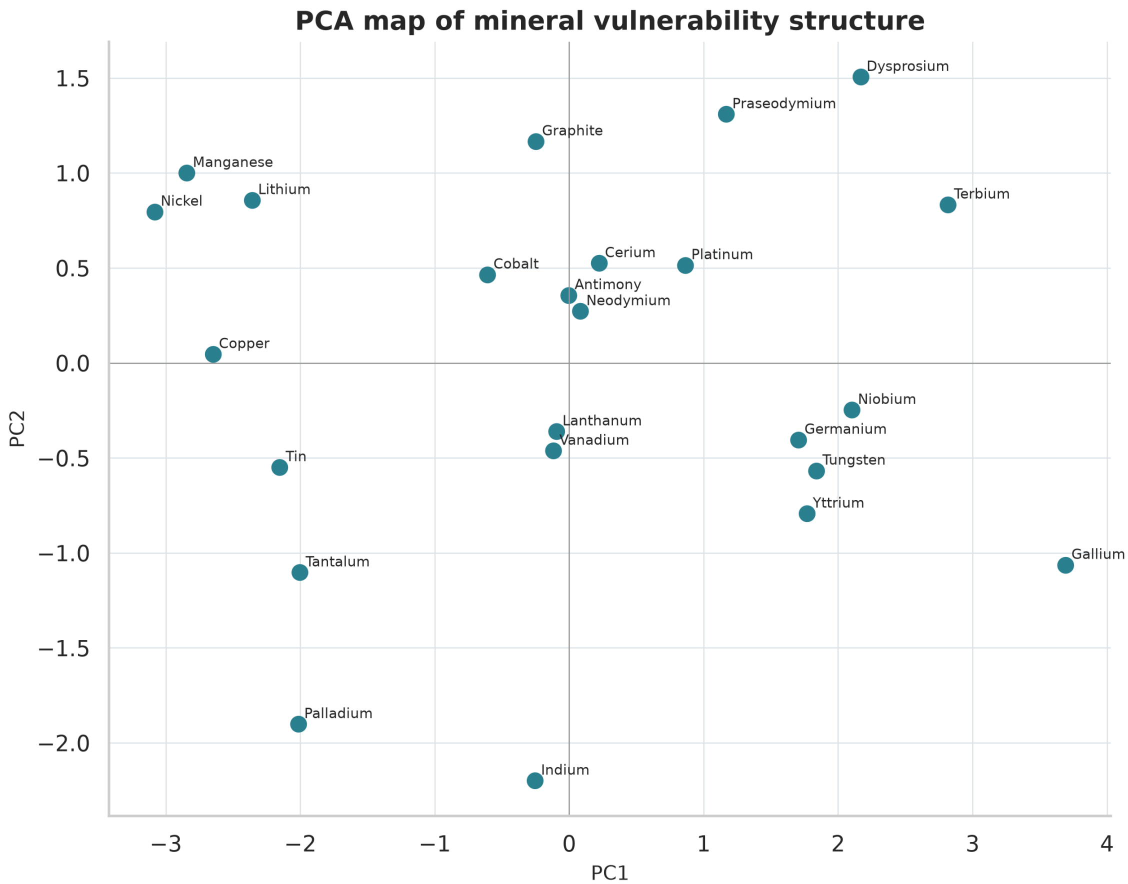

Figure 20. PCA map of mineral vulnerability structure. PCA summarizes standardized CMVI components; labels reveal latent structural similarity.

5.6 Supply shock simulations (RQ8)

The deterministic 50% dominant-producer scenario removes 31.2% of supply for the median mineral. The maximum loss is 47.4%, corresponding to Gallium. Complete dominant-producer loss therefore ranges directly with CR1; top-two and top-three scenarios add coalition exposure and can leave surviving production even more concentrated.

In the medium Monte Carlo regime, Gallium has the highest expected loss (5.4%) and a 11.1% probability of loss above 20% under the declared stress process. Terbium, Dysprosium, Germanium, and Tungsten follow on expected loss. Tail metrics sharpen the ordering: Gallium SupplyVaR95 is 46.1%, while the five highest SupplyCVaR95 values are Gallium 61.6%; Terbium 55.8%; Tungsten 53.8%; Dysprosium 53.1%; Germanium 53.1%. Average-loss and tail-loss views therefore identify the same broad group but can change its internal order.

Low-medium and medium-high CVaR rankings correlate at 0.922 and 0.978, respectively. Stress severity changes levels much more than order. These outcomes are simulation-based conclusions: they establish which structures transform common shocks into the greatest losses, not which countries are most likely to experience real disruptions.

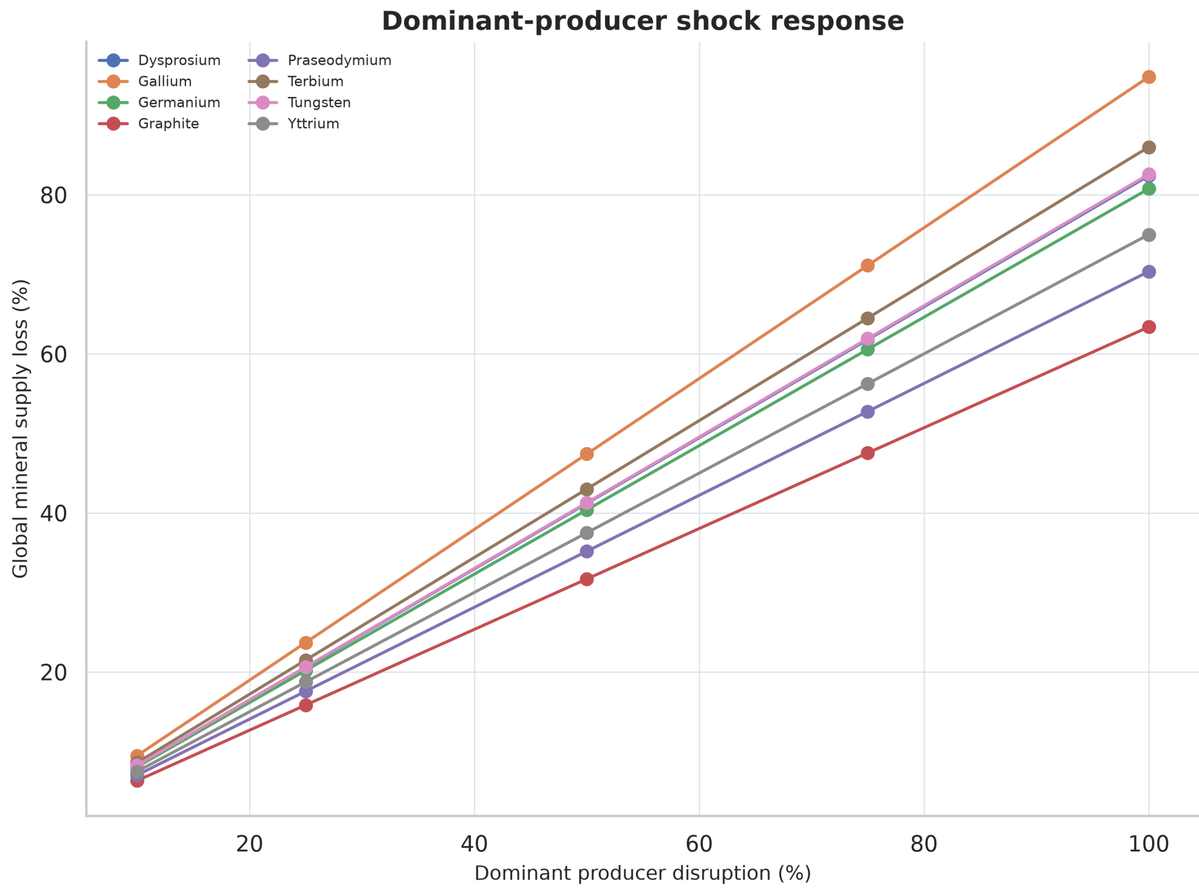

Figure 22. Dominant-producer shock response. Loss equals disrupted fraction multiplied by the dominant producer's recomputed share.

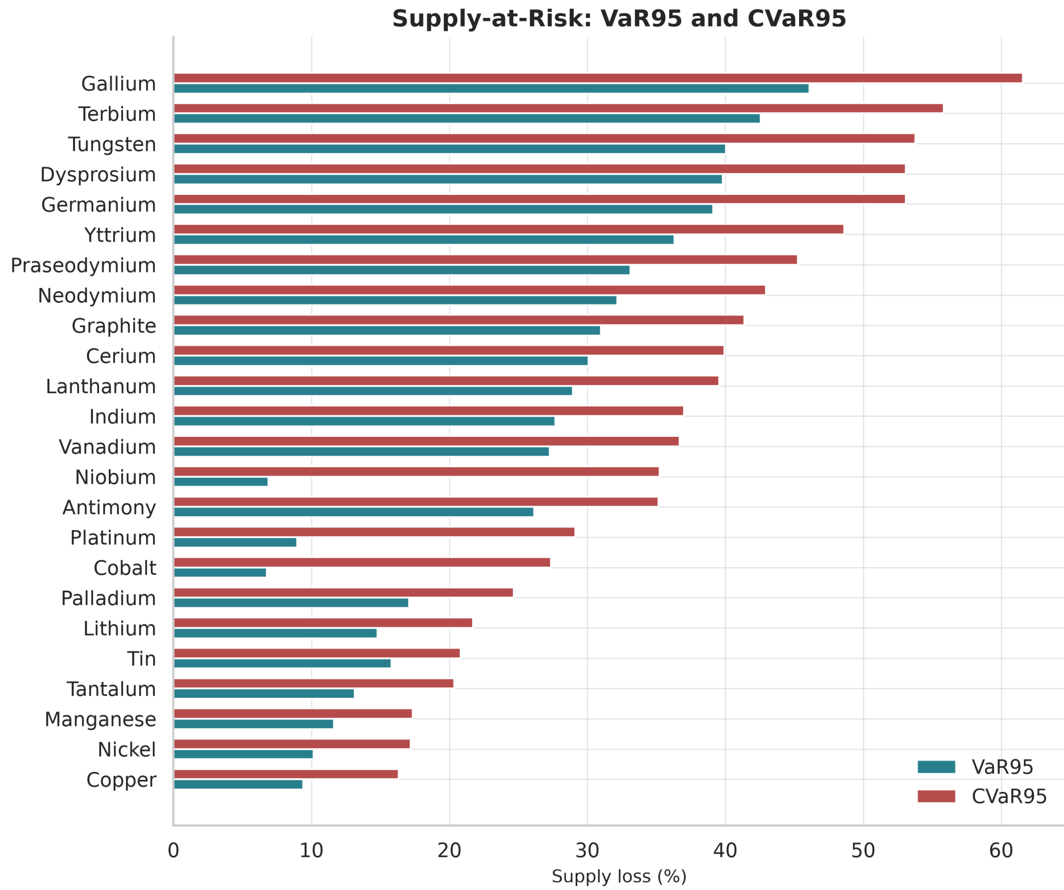

Figure 24. Supply-at-Risk: VaR95 and CVaR95. VaR95 is the 95th percentile loss; CVaR95 averages losses at or above VaR95.

5.7 Machine-learning performance and drivers (RQ9-RQ10)

Linear Regression is the best validation model. On untouched 2025-2026 target observations it achieves MAE 0.008, RMSE 0.010, R-squared 0.994, and Spearman rank correlation 0.995. The mean baseline has test MAE 0.112, while the previous-year target baseline has test MAE 0.010; the selected model therefore adds predictive information beyond constant and persistence benchmarks. Linear Regression provides an interpretable comparator (test R-squared 0.994). Partial-dependence and residual diagnostics are provided as supplementary figures.

The very high R-squared requires careful interpretation. The target is next-year simulated CVaR produced by a persistent share/network system. High accuracy means lagged structure predicts future *structural stress exposure* in the synthetic panel. It would be misleading to translate this roughly 99% structural fit into an explanation of actual geopolitical disruptions. The correct predictive conclusion is narrower and still useful: concentration and network exposure are stable enough that one-year-ahead risk prioritization is highly accurate under the common stress design.

SHAP ranks maximum systemic-country exposure first, followed by overall network dependency, CR1, HHI, effective producers, inverse diversification, CR3, and producer count. The ordering is economically coherent. Exposure rises when a mineral's largest shares belong to countries that are also important elsewhere, not only when its own HHI is high. This explains why Niobium—second in pure concentration—falls to consensus rank 10: its leading producer is less central across the multi-mineral network. Permutation importance produces the same broad family of top features, supporting interpretive stability.

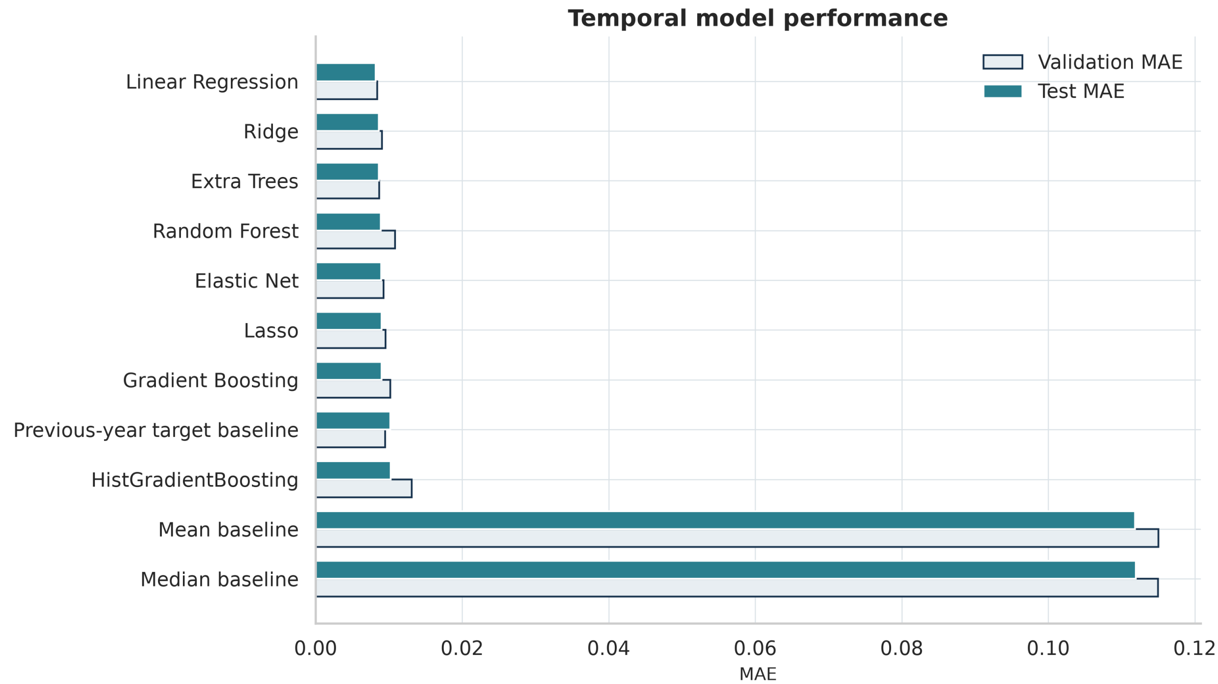

Figure 26. Temporal model performance. Validation and test MAE use whole-year temporal blocks; the final test contains target years 2025-2026.

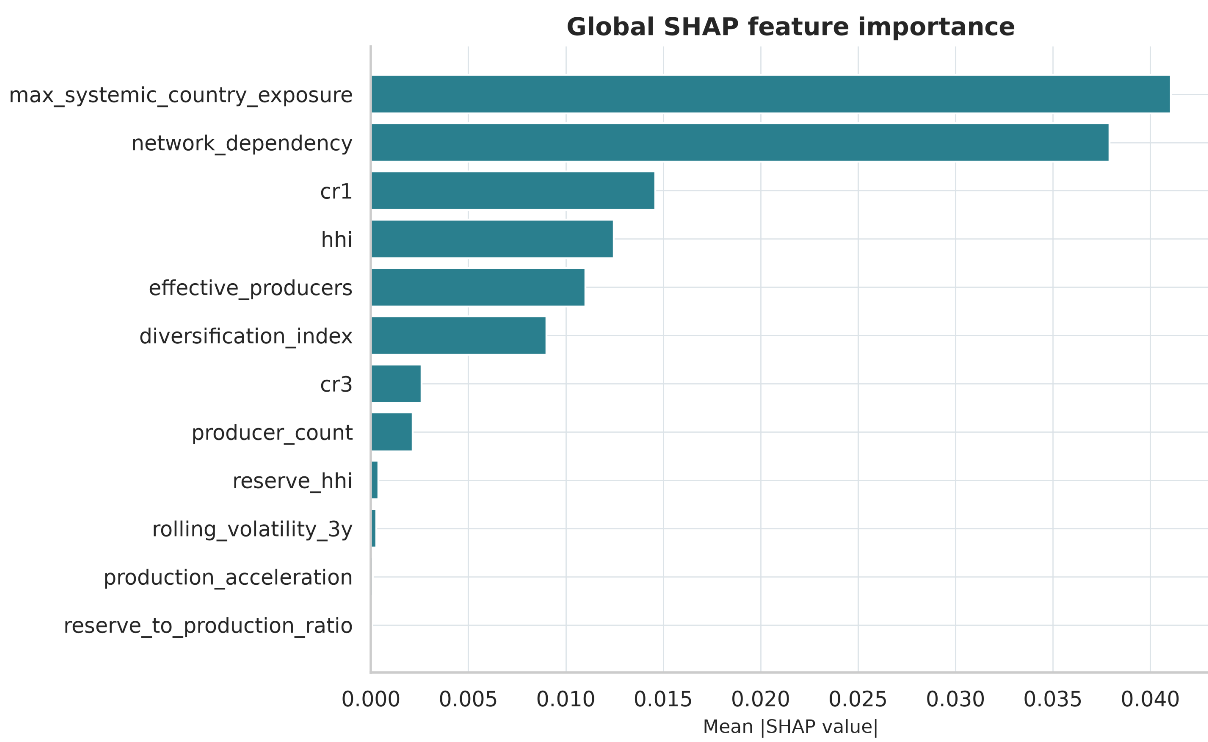

Figure 28. Global SHAP feature importance. Mean absolute SHAP values summarize the best compatible tree model on held-out observations.

5.8 Sensitivity and consensus (RQ11)

Alternative weights do not eliminate rank uncertainty, but they identify a stable extreme group. Gallium appears in the top five under 87.7% of 10,000 Dirichlet weight vectors; Terbium under 98.9%; and Dysprosium under 94.2%. The figure reports full 5th-95th rank intervals rather than only point ranks. Minerals whose risk comes from one component have wider intervals; minerals combining concentration, network dependence, volatility/growth pressure, and reserve vulnerability remain high across most weights.

The final six-family consensus is Gallium, Terbium, Dysprosium, Tungsten, and Germanium at the top and Nickel, Copper, Manganese, Tantalum, and Tin at the bottom. Gallium's mean family rank is 1.167; Terbium's is 2.000; and Dysprosium's is 3.500. Tungsten and Germanium complete the top five. Niobium is the most instructive disagreement: concentration rank 2.000 but network rank 23.000, producing consensus rank 10. Platinum similarly has concentration rank 9.000, network rank 24.000, and consensus rank 14. These disagreements validate the multi-method design: pure HHI detects commodity-specific dependence, while the network family distinguishes cross-mineral systemic importance.

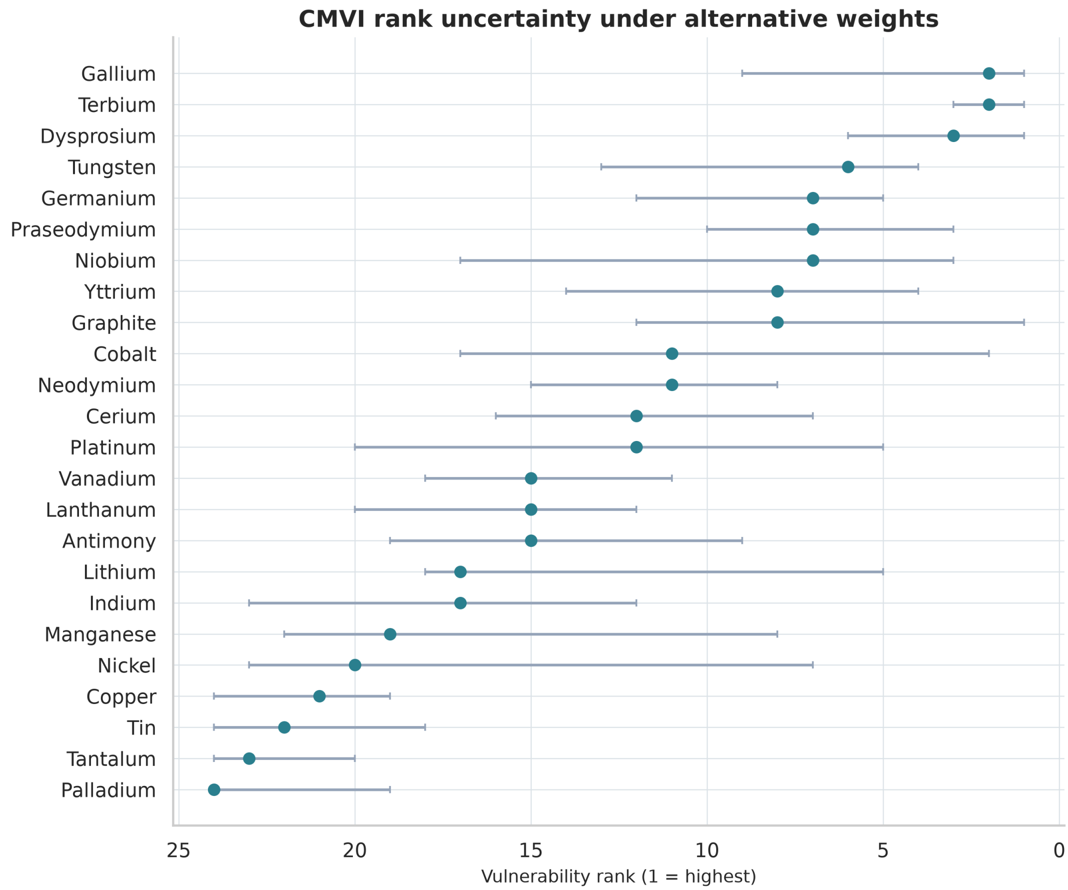

Figure 30. CMVI rank uncertainty under alternative weights. Intervals show 5th-95th percentile rank under 10,000 Dirichlet-distributed component weight vectors.

Table 1. Preregistered hypothesis tests. Statistics, p-values, effect sizes, bootstrap intervals, and mineral sample sizes.

test	statistic	p_value	effect_size	ci_low	ci_high	n
A_Pearson	-0.315	0.134	-0.315	-0.593	0.048	24
A_Spearman	-0.162	0.450	-0.162	-0.517	0.262	24
B_Pearson	0.385	0.063	0.385	0.103	0.633	24
B_Spearman	0.410	0.047	0.410	0.040	0.679	24
C_Paired_t	-2.579	0.017	-0.527	-0.024	-0.003	24
C_Wilcoxon	72.000	0.025	-0.333	-0.025	0.003	24
D_Mann_Whitney	60.000	0.507	-0.167	-0.625	0.306	24

Table 2. Top consensus vulnerability ranks. Lower family ranks indicate greater exposure.

mineral	consensus_rank	mean_family_rank	rank_concentration	rank_cmvi	rank_pca	rank_network	rank_tail_risk	rank_ml
Gallium	1	1.167	1.000	2.000	1.000	1.000	1.000	1.000
Terbium	2	2.000	3.000	1.000	2.000	2.000	2.000	2.000
Dysprosium	3	3.500	5.000	3.000	3.000	3.000	4.000	3.000
Tungsten	4	4.333	4.000	6.000	5.000	4.000	3.000	4.000
Germanium	5	5.833	6.000	7.000	7.000	5.000	5.000	5.000
Yttrium	6	6.500	7.000	8.000	6.000	6.000	6.000	6.000
Praseodymium	7	7.167	10.000	4.000	8.000	7.000	7.000	7.000
Neodymium	8	9.500	11.000	11.000	11.000	8.000	8.000	8.000
Graphite	9	10.000	12.000	5.000	15.000	9.000	9.000	10.000
Niobium	10	11.000	2.000	9.000	4.000	23.000	14.000	14.000
Cerium	11	11.333	13.000	13.000	10.000	11.000	10.000	11.000
Lanthanum	12	12.167	14.000	16.000	13.000	10.000	11.000	9.000

Table 3. Systemic producer countries. The systemic score combines network centrality and dominance.

country	systemic_score	weighted_degree	top1_count	top3_count
China	1.000	10.655	15	15
Russia	0.216	1.516	1	8
Australia	0.197	1.430	1	9
South Africa	0.174	1.618	2	4
Democratic Republic of the Congo	0.150	1.266	2	3
Brazil	0.134	1.328	1	3
Chile	0.101	0.685	1	2
United States	0.100	0.799	0	6
Indonesia	0.087	0.694	1	2
Myanmar	0.079	0.670	0	4

Table 4. Temporal model comparison. Test values use target years 2025-2026.

model	validation_mae	test_mae	test_rmse	test_r2	test_spearman
Linear Regression	0.008	0.008	0.010	0.994	0.995
Extra Trees	0.009	0.009	0.011	0.993	0.991
Ridge	0.009	0.009	0.011	0.994	0.993
Elastic Net	0.009	0.009	0.011	0.993	0.993
Previous-year target baseline	0.009	0.010	0.014	0.990	0.989
Lasso	0.010	0.009	0.011	0.993	0.993
Gradient Boosting	0.010	0.009	0.011	0.993	0.991
Random Forest	0.011	0.009	0.012	0.992	0.991
HistGradientBoosting	0.013	0.010	0.012	0.991	0.990
Median baseline	0.115	0.112	0.133	-0.004	nan
Mean baseline	0.115	0.112	0.133	-0.002	nan

6. Discussion

The supplied mineral system is neither uniformly concentrating nor uniformly resilient. The broad temporal signal is diversification: mean and median HHI decline, two-thirds of minerals move downward, and both paired tests support a distributional shift. But resilience cannot be inferred from that average. The endpoint still contains near-monopoly structures, and several diversifying minerals remain among the most concentrated. The central substantive distinction is between *direction* and *level*. Dysprosium improved materially but remained fragile; Copper concentrated modestly but remained among the most diversified. Policy or portfolio analysis that reports only changes would misclassify both.

The most vulnerable minerals share a structural signature. Gallium, Terbium, Dysprosium, Tungsten, and Germanium have high dominant-producer shares, low effective producer counts, and exposure to a country node that is itself important across many minerals. Their consensus position does not depend on a single method, but the weight experiment distinguishes a stable top trio from more weight-sensitive fourth and fifth positions. Deterministic loss, stochastic tail risk, PCA, equal CMVI, network exposure, and temporal prediction converge most strongly for Gallium, Terbium, and Dysprosium. This triangulation is stronger than claiming that Gallium is “number one” because one index assigns it the largest score.

The lowest-ranked group illustrates the opposite architecture. Nickel, Copper, Manganese, Tantalum, and Tin distribute production across more effective sources and experience smaller standardized tail losses. Lower structural vulnerability does not mean these minerals are unimportant, abundant, environmentally benign, or immune to shortages. Copper could face demand-supply imbalance or long project lead times despite diversified geographic shares; Nickel could remain exposed at the refining stage. This study's resilience statement is conditional on mine-production geography and included proxies.

The country results sharpen the meaning of “chokepoint.” China is structurally different from a country that dominates one mineral. Its weighted degree of 10.655 means that large shares recur across much of the mineral set, and it leads 15 minerals. Targeted removal greatly exceeds a random removal baseline. Russia, Australia, South Africa, the Democratic Republic of the Congo, and Brazil remain important, but their scores are far lower and their precise ordering is threshold-sensitive. The robust conclusion is a single dominant systemic hub plus a second tier of regional/commodity specialists.

Growth produces a nuanced result. Cross-sectionally, fast-growing minerals are not more concentrated; within minerals, annual growth is positively associated with HHI after fixed effects. These results answer different questions and are not contradictory. Cross-sectional rankings reflect stable commodity-specific production architectures. The panel coefficient reflects whether above-normal growth years for a given mineral coincide with above-normal concentration after common year effects. The finding is statistical within the synthetic system, not causal evidence that demand growth inevitably consolidates production.

Machine learning adds two layers. First, it quantifies persistence: one-year-ahead stress exposure is highly predictable, so the structural ordering is not transient noise. Second, SHAP reveals why the best tree model makes its predictions. Cross-mineral systemic-country exposure leads, followed by network dependency and familiar concentration indicators. ML therefore adds interaction and ranking evidence but does not overturn the transparent measures. Its strong performance is partly a property of predicting a structural simulation generated from persistent shares. That limitation is not a failure; it defines what the model successfully learns.

The stress tests translate abstract concentration into operational scale. A median 50% dominant-producer disruption removes 31.2% of mineral supply; Gallium loses 47.4%. In medium Monte Carlo stress, Gallium's worst-five-percent mean loss is 61.6%. Those quantities are easier to use in scenario planning than HHI alone. They identify where diversification, strategic inventory, substitution, recovery, or procurement redundancy would have the largest structural benefit, while stopping short of assigning real-world event probabilities.

7. Limitations

The most important limitation is also the clearest: this is a synthetic dataset. The source dictionary labels production, reserves, and prices as simulated, supplied shares fail accounting identities, and 2026 cannot be treated as a completed historical observation. Results demonstrate an analytical framework and describe the generated panel. They must not be cited as official country production shares. A real-world replication should rebuild the panel from USGS commodity chapters, IEA mining/refining datasets, national geological surveys, or another documented observational source.

The panel models mine production more directly than refining and processing. Refining percentages are normalized proxies because refining tonnage is absent and original shares do not sum to 100. Mine diversity can coexist with processing concentration, a distinction emphasized by the IEA. Trade flows are also absent, preventing bilateral dependency, import reliance, rerouting, sanctions propagation, and logistics analysis. Reserves are simulated and do not capture resource quality, extraction cost, permitting, project maturity, or reserve-definition revisions.

Annual frequency smooths short disruptions. Price is available but simulated, and the project does not estimate price elasticities, substitution, inventory response, recycling, or demand destruction. Production tonnages cannot be added across minerals as if one tonne of Copper were commensurate with one tonne of Gallium; systemic results therefore average relative losses. Environmental, social, governance, and conflict indicators are absent. The network is static within each annual slice and represents production relations, not contractual supply chains.

Monte Carlo propensities and beta magnitudes are stress assumptions. They are intentionally transparent but not empirically estimated geopolitical probabilities. Expected loss and CVaR are comparative structural metrics. Their rank stability across low, medium, and high regimes supports prioritization, but absolute loss probabilities should not be used for insurance pricing or forecasts. Correlated shocks are stylized and do not model alliances, common infrastructure, weather basins, or policy contagion.

The mineral-level sample is small for cross-sectional inference ($n = 24$), limiting power. Clustered panel inference has 24 mineral clusters. Multiple tests are reported as a planned family and interpreted by effect size and uncertainty rather than p-value counting. Fixed-effects models cannot estimate time-invariant producer count. The ML sample has 264 lagged mineral-year observations and a mechanically defined structural target; high predictive performance does not generalize automatically to observed disruptions.

These limitations narrow rather than erase the conclusion. The project can establish exactly how concentration, entropy, network position, and common shocks interact in the supplied system; whether rankings are stable across methods; and whether structural exposure is predictable one year ahead. It cannot establish real-world frequencies, causal effects, or official market shares. That boundary is maintained throughout the paper.

8. Conclusion

Finding 1 - Concentration. The panel is substantially but heterogeneously concentrated. Mean HHI decreased by 0.014 from 2015 to 2026; 16 of 24 minerals diversified and eight concentrated. The paired t-test and Wilcoxon test both support a negative shift. Global production therefore generally diversified within the synthetic panel, but the modest average improvement coexists with extreme endpoints.

Finding 2 - Highest-risk minerals. Across concentration, equal-weight CMVI, PCA, network dependency, Monte Carlo tail loss, and ML-predicted exposure, Gallium, Terbium, Dysprosium, Tungsten, and Germanium are the most structurally vulnerable. Gallium is the strongest case: HHI 0.900, CR1 94.8%, effective producer count 1.111, consensus rank 1, and CVaR95 61.6%. Terbium and Dysprosium are nearly as stable across methods. Tungsten and Germanium complete the high-risk group because concentration and central-country exposure reinforce one another.

Finding 3 - Most resilient minerals. Nickel, Copper, Manganese, Tantalum, and Tin occupy the five lowest consensus positions. Tin, Copper, Nickel, and Manganese have the lowest latest HHI values and roughly five effective producers. Their relative resilience is geographic and structural; it does not imply low economic importance or adequate future capacity.

Finding 4 - Systemic countries. China is the largest chokepoint by a wide margin: systemic score 1.000, weighted degree 10.655, and first place in 15 minerals. Removing China produces 44.4% average mineral loss versus 2.8% for a random country. China, Russia, Australia, South Africa, and Democratic Republic of the Congo form the leading country set, but only China's first position is invariant and dominant across all network thresholds.

Finding 5 - Shock exposure. A 50% disruption to the dominant producer removes 31.2% of supply for the median mineral and 47.4% at the maximum. Gallium, Terbium, Dysprosium, Germanium, and Tungsten have the highest conditional tail losses. These are structural stress outcomes, not geopolitical forecasts.

Finding 6 - Statistical relationships. Cross-sectional growth-concentration tests do not reach 5% two-sided significance. For volatility versus mean HHI, Spearman reaches the 5% threshold ($p = 0.047$) while Pearson does not ($p = 0.063$); the positive relationship is therefore supported by rank evidence but sensitive to functional form. Concentration decreased significantly over time. The fixed-effects regression finds annual growth positively associated with within-mineral HHI after mineral and year effects, while other coefficients must be interpreted as conditional associations. The statistics reject a universal claim that fast-growing minerals are necessarily more concentrated across commodities.

Finding 7 - Machine learning. Linear Regression is the best model, with test MAE 0.008, R-squared 0.994, and rank correlation 0.995. Maximum systemic-country exposure, network dependency, CR1, HHI, and effective producer count are the leading SHAP signals. ML adds evidence that structural exposure persists and that network context matters beyond single-mineral concentration. Its accuracy applies to next-year simulated CVaR, not real events.

Finding 8 - Robustness. HHI and inverse entropy correlate at rho 0.982; equal CMVI and PCA at rho 0.854; and medium versus high Monte Carlo tails at rho 0.978. Gallium, Terbium, and Dysprosium have top-five probabilities of 87.7%, 98.9%, and 94.2% under alternative weights. The principal high-risk conclusion is therefore robust, while middle ranks such as Niobium and Platinum depend meaningfully on whether commodity-specific or systemic network risk is emphasized.

The project discovers that vulnerability in the supplied critical-minerals system is a persistent structural property of concentrated, cross-connected production networks. Modest diversification over time did not dissolve the tail: a small group of minerals remained dependent on one effective producer, and a single country connected those fragile structures across much of the mineral set. Concentration measures identify commodity bottlenecks; network removal identifies the systemic country; Monte Carlo converts both into supply-at-risk; and temporal ML confirms persistence. Within the explicit synthetic-data boundary, the evidence is decisive: resilience is greatest where production shares are both geographically distributed and connected to alternative country nodes, while the greatest vulnerability arises when dominant-producer dependence and systemic network centrality coincide.

References

Breiman, L. (2001). Random Forests. *Machine Learning*, 45, 5-32. https://doi.org/10.1023/A:1010933404324

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1252 establishing a framework for ensuring a secure and sustainable supply of critical raw materials. https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX:32024R1252

Graedel, T. E., Harper, E. M., Nassar, N. T., Nuss, P., & Reck, B. K. (2015). Criticality of metals and metalloids. *PNAS*, 112(14), 4257-4262. https://doi.org/10.1073/pnas.1500415112

Hund, K., La Porta, D., Fabregas, T., Laing, T., & Drexhage, J. (2020). *Minerals for Climate Action*. World Bank.

International Energy Agency. (2025). *Global Critical Minerals Outlook 2025*. Paris: IEA. https://www.iea.org/reports/global-critical-minerals-outlook-2025

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: a review and recent developments. *Philosophical Transactions A*, 374, 20150202. https://doi.org/10.1098/rsta.2015.0202

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *NeurIPS 30*.

Nassar, N. T., Du, X., & Graedel, T. E. (2015). Criticality of the rare earth elements. *Journal of Industrial Ecology*, 19(6), 1044-1054. https://doi.org/10.1111/jiec.12237

Nassar, N. T., Graedel, T. E., & Harper, E. M. (2015). By-product metals are technologically essential but have problematic supply. *Science Advances*, 1(3), e1400180. https://doi.org/10.1126/sciadv.1400180

Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379-423, 623-656.

U.S. Department of Justice & Federal Trade Commission. (2023). *2023 Merger Guidelines*. https://www.justice.gov/atr/merger-guidelines

U.S. Geological Survey. (2026). *Mineral Commodity Summaries 2026*. https://doi.org/10.3133/mcs2026

Appendix

Appendix A. Full consensus triangulation

mineral	consensus_rank	mean_family_rank	rank_concentration	rank_cmvi	rank_pca	rank_network	rank_tail_risk	rank_ml
Gallium	1	1.167	1.000	2.000	1.000	1.000	1.000	1.000
Terbium	2	2.000	3.000	1.000	2.000	2.000	2.000	2.000
Dysprosium	3	3.500	5.000	3.000	3.000	3.000	4.000	3.000
Tungsten	4	4.333	4.000	6.000	5.000	4.000	3.000	4.000
Germanium	5	5.833	6.000	7.000	7.000	5.000	5.000	5.000
Yttrium	6	6.500	7.000	8.000	6.000	6.000	6.000	6.000
Praseodymium	7	7.167	10.000	4.000	8.000	7.000	7.000	7.000
Neodymium	8	9.500	11.000	11.000	11.000	8.000	8.000	8.000
Graphite	9	10.000	12.000	5.000	15.000	9.000	9.000	10.000
Niobium	10	11.000	2.000	9.000	4.000	23.000	14.000	14.000
Cerium	11	11.333	13.000	13.000	10.000	11.000	10.000	11.000
Lanthanum	12	12.167	14.000	16.000	13.000	10.000	11.000	9.000
Vanadium	13	13.833	16.000	15.000	14.000	12.000	13.000	13.000
Antimony	14	14.333	17.000	14.000	12.000	13.000	15.000	15.000
Platinum	14	14.333	9.000	12.000	9.000	24.000	16.000	16.000
Indium	16	14.500	15.000	18.000	16.000	14.000	12.000	12.000
Cobalt	17	15.167	8.000	10.000	17.000	22.000	17.000	17.000
Lithium	18	18.667	20.000	17.000	21.000	16.000	19.000	19.000
Palladium	19	19.500	18.000	24.000	19.000	20.000	18.000	18.000
Tin	20	20.167	24.000	22.000	20.000	15.000	20.000	20.000
Tantalum	21	20.500	19.000	23.000	18.000	21.000	21.000	21.000
Manganese	22	21.000	23.000	19.000	23.000	17.000	22.000	22.000
Copper	23	21.500	21.000	21.000	22.000	18.000	24.000	23.000
Nickel	24	22.000	22.000	20.000	24.000	19.000	23.000	24.000

Appendix B. Monte Carlo assumptions

The low, medium, and high regimes use common base event probabilities 0.03, 0.08, and 0.15. Country-specific probability equals the base multiplied by `0.5 + normalized systemic centrality`, capped at 0.75. Idiosyncratic shock magnitude is Beta(2,3). A correlated event occurs with probability equal to half the base rate, selects one country for the shared year-regime draw with weight proportional to centrality plus 0.05, and draws Beta(3,2) severity. The same event is applied wherever that country produces a mineral. Overlapping idiosyncratic and systemic losses combine sequentially as `1 - (1-a)(1-b)`, so no producer loses more than its physical share. Every mineral-year-regime uses 10,000 simulations and seed 42 plus year. These values define stress environments and are not fitted event probabilities.

Appendix C. CMVI and sensitivity details

All six components are min-max normalized in the latest year and receive baseline weight 1/6. Dirichlet sensitivity draws 10,000 nonnegative vectors summing to one. PCA uses standardized raw components and orients PC1 toward equal-weight risk. Outcome calibration is excluded from confirmatory rankings because the simulated target is constructed from related structural shares. The final consensus weights six method families equally, not each raw metric.

Appendix D. Model search and leakage controls

Tree search uses eight randomized candidates per family over bounded estimator counts, depth/leaf parameters, learning rates, feature fractions, and regularization. Expanding-year folds keep entire years together. Hyperparameters are selected without test data. The final test contains target years 2025 and 2026. The feature list and trained preprocessing/model pipelines are persisted. Source `supply_risk_score`, `high_supply_risk`, `disruption`, and `disruption_next_year` never enter the primary features or target.

Appendix E. Robustness matrix

comparison	spearman	kendall_tau
HHI vs inverse entropy	0.982	0.899
CR1 vs CR3	0.864	0.696
Equal CMVI vs PCA	0.854	0.667
Monte Carlo low vs medium	0.922	0.804
Monte Carlo medium vs high	0.978	0.928
Dominant shock 25% vs 75%	1.000	1.000
Raw vs winsorized-volatility CMVI	0.998	0.986
Min-max vs percentile-normalized CMVI	0.981	0.909

Appendix F. Reproducibility

The repository records Python and package versions, source checksums, random seed, simulation and bootstrap draw counts, feature names, model hyperparameters, output manifests, and QA checks. Run `python run_pipeline.py` in the pinned virtual environment. The pipeline reconstructs raw copies, cleaned data, all tables and figures, models, notebooks, Markdown, PDF, and final JSON. No core conclusion is hardcoded independently of the computed result bundle.