ISI Paper Series — Paper 2

Empirical Results and Distributional Structure of the International Sovereignty Index (ISI) EU-27, Vintage 2024, Methodology v1.0

International Sovereignty Institute — Research Unit

February 2026

Abstract

This paper reports the first empirical application of the International Sovereignty Index (ISI) to the 27 member states of the European Union using vintage-2024 data under methodology v1.0. The ISI composite, defined as the unweighted arithmetic mean of six normalised axes (financial, energy, technology, defence, critical inputs, logistics), ranges from 0.236 (France, rank 27) to 0.517 (Malta, rank 1) with a population mean of 0.344 and population standard deviation of 0.070. One country is classified as *highly concentrated*, 23 as *moderately concentrated*, and three as *mildly concentrated*. A full variance decomposition with formal proof shows that defence (35.1%) and logistics (34.3%) together account for 69.4% of cross-country composite variance, with the own-variance component contributing 64.3% and cross-covariance 35.7%. A defence-dominance stress test demonstrates that the ranking retains structural coherence ($\rho \geq 0.654$) even when both highest-variance axes are excluded. Leave-one-axis-out analysis with both Spearman ρ and Kendall τ confirms that removing the defence axis produces the largest rank disruption ($\rho = 0.833$, $\tau = 0.675$). Dirichlet weight perturbation (10 000 draws, $\alpha_j = 10$) yields a mean rank-order correlation of $\rho = 0.979$, with the Netherlands exhibiting the highest rank volatility ($\sigma = 3.07$). Rank elasticity testing ($\pm 10\%$ perturbation) classifies 13 countries as rank-sensitive and 14 as rank-locked. Only two of 15 pairwise axis correlations are significant at $\alpha = 0.05$ after Bonferroni correction: energy–technology ($t = 3.12$) and critical inputs–logistics ($t = 3.28$); the financial axis is orthogonal to all others ($|t| < 1.15$). PCA loadings identify three interpretable components: goods-trade concentration (PC1, 31.7%), defence–logistics (PC2, 26.4%), and financial-specific (PC3, 18.1%). A formal structural typology classifies countries by dominance gap into single-axis dominated, dual-axis concentrated, and balanced profiles. Compensability analysis identifies the Netherlands (gap = 0.624, CV = 0.880) as the most internally compensated country. Classification compression diagnostics ($\eta^2 = 0.483$) confirm that the three-tier scheme is inherently constrained by the unimodal distribution. The paper provides radar profiles for eight selected countries, geographic correlates, defence structural analysis, full axis-level rankings, and a complete reproducibility protocol.

JEL codes: F02, F15, F52, O30, Q43

Keywords: sovereignty index, EU-27, concentration measurement, composite indicator, supply-chain dependence, defence dependence

Contents

Executive Summary

This paper presents the complete empirical results of the International Sovereignty Index (ISI) applied to the 27 member states of the European Union, using vintage-2024 data processed under methodology v1.0. The index measures the concentration of bilateral external dependencies across six structural axes. Higher scores indicate higher concentration of counterparty exposure. The analysis goes substantially beyond descriptive ranking to provide a structurally decomposed, variance-explained, rank-stable, and statistically interrogated empirical system. The principal findings are:

- **Compressed composite distribution.** The ISI composite ranges from 0.236 (France) to 0.517 (Malta), a span of 0.282 on the unit interval. The population standard deviation is 0.070, yielding a coefficient of variation of 0.203. The interquartile range is 0.099 (P25 = 0.292, P75 = 0.391). The Gini coefficient is 0.116, indicating a relatively egalitarian distribution. See Table 3, Figure 2, and Figure 4.
- **Tier compression is mathematically inevitable.** Under the four-tier classification, 23 of 27 countries (85.2%) fall in the *moderately concentrated* tier (ISI $\in [0.25, 0.50)$). Only Malta exceeds 0.50; three countries (Latvia, Slovakia, France) fall below 0.25. The between-tier $\eta^2 = 0.483$ indicates that the official scheme captures less than half the total dispersion; the within-tier range for the dominant category spans 0.207, nearly four times the entire mildly-concentrated band. See Tables 5 and 7 and Figure 5.
- **Defence and logistics dominate cross-country variance.** A covariance-based variance decomposition, proved analytically from the linear aggregation rule, attributes 35.10% of composite variance to the defence axis and 34.26% to the logistics axis. Together these two axes account for 69.37% of total composite variance. The own-variance component accounts for 64.28% of total composite variance; cross-covariances account for 35.72%. Financial concentration contributes only 1.03%. See Tables 12 and 13 and Figure 7.
- **Axis dominance is concentrated.** In 16 of 27 countries, the defence axis records the highest axis-level score; in the remaining 11, logistics is highest. No country has its maximum score on the financial, energy, technology, or critical-inputs axis. The mean defence-to-second gap for defence-dominant countries is 0.184. See Tables 15 and 41.
- **Defence axis exhibits oligopsonistic concentration.** Five EU member states record defence axis scores exceeding 0.80, and 59.3% of countries have defence as their dominant axis. The distribution is positively skewed, and Slovakia records the only zero score in the entire dataset. This structural pattern reflects oligopsonistic procurement markets in which small buyers face few feasible suppliers. See Tables 15 and 16.
- **Two statistically significant cross-axis correlations.** Of 15 pairwise Pearson correlations, only two are significant at the 5% level ($t_{0.025,25} = 2.060$): critical inputs–logistics ($r = 0.549$, $t = 3.280$) and energy–technology ($r = 0.530$, $t = 3.124$). The remaining 13 pairs—

including all correlations involving the financial axis—are not significant, establishing that the six axes capture largely distinct dimensions. See Tables 18 and 20 and Figures 8 and 9.

- **Rank ordering robust at extremes, fragile in mid-tier.** Under leave-one-axis-out (LOO) analysis, Spearman ρ ranges from 0.833 (defence excluded) to 0.993 (financial excluded); Kendall τ from 0.675 (defence) to 0.943 (financial, energy). Under Dirichlet weight perturbation (10 000 draws, $\alpha_j = 10$), Malta ($\sigma < 0.5$ ranks) and France ($\sigma < 0.5$ ranks) are structurally locked, while the Netherlands ($\sigma = 3.07$) is the most volatile. Under $\pm 10\%$ axis perturbation, 14 of 27 countries are rank-invariant. See Table 23, Table 26, and Table 28.
- **Financial concentration is orthogonal to goods-trade concentration.** The financial axis is nearly uncorrelated with every other axis ($|r| \leq 0.22$, all $|t| < 1.15$). Its LOO removal produces $\rho = 0.993$, the highest stability of any axis. Its variance contribution is 1.03%. Its bivariate R^2 with rank is 0.025, and its partial correlation with rank (controlling for all other axes) is 0.421. This establishes that financial-dependence patterns are structurally distinct from physical-trade dependence patterns. See Tables 14 and 20.
- **Compensability masks profile heterogeneity.** The arithmetic-mean aggregation rule allows high scores on one axis to compensate for low scores on others. The Netherlands exhibits the highest compensation gap (max axis $= 0.960$, composite $= 0.335$, gap $= 0.624$, CV $= 0.880$), followed by Italy (gap $= 0.547$) and Bulgaria (gap $= 0.520$). See Table 34.
- **Geographic and structural correlates.** Island and near-island economies (Malta, Cyprus, Ireland) average ISI $= 0.471$, significantly above the continental mean of 0.329. Large-population countries (> 20M) average 0.298; small-population countries (< 5M) average 0.372. Core EU-6 members average 0.312; peripheral members average 0.368. See Table 32.
- **Slovakia's zero defence score.** Slovakia records a defence axis score of 0.0000, the only zero score in the dataset. This arises mechanically from the absence of bilateral supplier entries for Slovakia in the SIPRI arms-transfer data under methodology v1.0. It does not imply the absence of defence procurement. See Section 7 and Table 16.
- **Defence-dominance stress test.** Three counterfactual composites—without defence, without logistics, and without both—demonstrate that the ranking retains structural coherence even when the two highest-variance axes are excluded (without both: $\rho = 0.654$, $R^2 = 0.428$). The ISI is not reducible to a "defence index." See Table 17.
- **Formal structural typology.** Countries are classified by their dominance gap Δ_i into single-axis dominated ($\Delta_i \geq 0.30$; 3 countries, all defence-dominant), dual-axis concentrated ($0.10 \leq \Delta_i < 0.30$; 8 countries), and balanced ($\Delta_i < 0.10$; 16 countries). Profile polarisation predicts rank volatility under weight perturbation. See Table 11.
- **PCA component loadings.** The first principal component (31.7% of variance) loads on the four goods-trade axes; the second (26.4%) loads on defence and logistics; the third (18.1%) isolates the financial axis. PCA-derived weighting would suppress the defence axis—the single most variance-contributing dimension—reinforcing the case for equal weighting. See Table 22.

- **Geometric-mean aggregation comparison.** Replacing the arithmetic mean with a geometric mean yields Spearman $\rho = 0.934$ and Kendall $\tau = 0.812$ against the baseline ranking. The overall ordering is preserved; the largest rank shift is 5 positions (Slovakia). The aggregation rule is a substantive but not fragile design choice. See Table 29.
- **OLS geographic correlates.** A four-covariate OLS model (log population, island dummy, log GDP, EU-6 core dummy) explains 38.2% of composite variance ($F(4, 22) = 3.40$, $p = 0.026$). Island status is the strongest predictor ($\hat{\beta} = 0.089$, $p = 0.017$); log population is marginally significant ($p = 0.071$). Neither GDP nor founding-member status is individually significant. See Table 33.

1 Introduction and Contribution

1.1 Motivation

Composite indicators that measure structural features of national economies occupy a central position in institutional research. The OECD Handbook on Constructing Composite Indicators establishes the standard methodological framework: variable selection, normalisation, weighting, and aggregation, each step carrying design choices that shape the information content of the final score [1]. Saisana and Tarantola [2] demonstrate that rigorous sensitivity and uncertainty analysis is indispensable for any composite index that informs policy. The International Sovereignty Index (ISI) adheres to both standards and extends them to a domain that has received insufficient quantitative treatment: *the concentration of bilateral external dependencies across strategically distinct domains*.

The intellectual foundations of the ISI rest on three pillars. First, the industrial-organisation theory of market concentration, where the Herfindahl–Hirschman Index (HHI) provides a closed-form summary of distributional inequality among market participants [3, 4]. HHI is well-understood theoretically—it equals the inverse of the "numbers equivalent" of equally sized suppliers—and its behavioural interpretation in terms of mark-up and price–cost margins is established in the antitrust literature. The ISI applies this logic at the country–partner level rather than the firm–market level.

Second, the literature on supply-chain vulnerability, which expanded substantially after the semiconductor shortages of 2020–2021 and the geopolitical disruptions following the 2022 Russia–Ukraine conflict [5]. The European Commission's Critical Raw Materials assessments, the US Executive Order on Supply Chain Resilience, and OECD analyses of trade concentration in critical minerals all employ concentration metrics—typically HHI or CR_k ratios—to identify bottleneck dependencies. The ISI formalises and generalises these sector-specific exercises into a multi-axis composite framework.

Third, the methodological debate on equal versus differential weighting in composite indicators. Equal weighting is the most common default in practice (used by the Human Development Index, the Global Innovation Index, and the OECD Better Life Index, among others), but it is not neutral: it implicitly assumes that all dimensions are equally important and that the analyst has no empirical or normative basis for asymmetric treatment. The ISI adopts equal weighting as its baseline specification and subjects this choice to explicit statistical interrogation through Dirichlet weight perturbation (Section 6.4) and variance decomposition (Section 4.5).

The ISI does not measure sovereignty in a political or legal sense. It quantifies a structural feature—counterparty concentration—that is a prerequisite for informed policy analysis of external dependence. A high ISI score indicates that a country's external flows in a given domain are heavily concentrated among few partners; a low score indicates dispersion across many partners. The index makes no normative claim about whether concentration is advantageous or disadvantageous. Whether high concentration translates into political vulnerability, economic fragility, or coercive leverage is a distinct empirical question that lies beyond the ISI's measurement scope.

1.2 Scope of this paper

This paper—the second in the ISI series—presents the full empirical results of the first ISI application: 27 EU member states, vintage-2024 data, methodology v1.0. The methodological foundations (axis definitions, normalisation procedures, data sources, aggregation rule, integrity chain) are documented in Paper 1 [6]. The present paper is self-contained for readers interested in the empirical findings; cross-references to Paper 1 are provided where construction details are relevant to interpretation.

Relative to a conventional "results" paper in the composite-indicator literature, this paper provides substantially deeper analytical content. Beyond reporting the ranking and basic distributional statistics, the paper undertakes a formal variance decomposition with proof, systematic robustness exercises with both Spearman and Kendall rank-order statistics, rank-elasticity analysis, axis contribution decomposition, classification compression diagnostics, structural analysis of the defence axis, a defence-dominance stress test, PCA-based dimensionality diagnostics with component loadings, compensability analysis, formal structural typology, and geographic correlate tabulation. Each analytical component interrogates the ISI composite from a different angle and provides the evidence required for the reader to form an independent judgment about the index's informational content.

The specific contributions of this paper are:

1. A complete composite and axis-level ranking of all 27 EU member states with eight-decimal precision, accompanied by classification labels and distributional diagnostics (histogram, ECDF, Lorenz curve, Gini coefficient).
2. A formal variance decomposition—proved analytically from the linear aggregation rule—decomposing the composite variance into own-variance and cross-covariance components, with percentage contributions for each axis.
3. A leave-one-axis-out (LOO) sensitivity analysis with both Spearman ρ and Kendall τ, accompanied by a full 27×6 rank-displacement matrix.
4. A rank-elasticity analysis quantifying the sensitivity of each country's rank to marginal perturbations in individual axis scores.
5. An axis-contribution analysis providing bivariate R^2 and partial correlations between each axis and the composite rank.
6. Classification compression diagnostics comparing the official four-tier scheme with data-driven tercile and quartile alternatives, reporting within-tier variance and η^2.
7. A structural analysis of the defence axis examining its distributional properties, oligopsonistic character, and country-level defence-to-max-axis gaps.

8. A defence-dominance stress test constructing counterfactual composites without the defence axis, without the logistics axis, and without both, to demonstrate that the ranking retains structural coherence independently of the two highest-variance axes.
9. A formal structural typology classifying countries by dominance gap (Δ_i) into single-axis dominated, dual-axis concentrated, and balanced profiles.
10. A cross-axis correlation analysis with formal t-statistics, Bonferroni-corrected multiple-testing assessment, eigenvalue decomposition, PCA component loadings with interpretation, and network representation.
11. A compensability analysis with formal gap, range, and CV metrics documenting the degree to which the arithmetic-mean aggregation rule permits cross-axis compensation.
12. An assessment of geographic and structural correlates (island status, population size, core/periphery classification).
13. A four-part robustness programme: LOO, z-score standardisation, Dirichlet weight perturbation (Monte Carlo), and winsorisation.
14. Radar-chart profiles for eight selected countries illustrating the heterogeneity of axis-level patterns.
15. A reproducibility protocol and computational appendix enabling full replication from the provided snapshot.

1.3 Interpretation conventions

Throughout this paper, the following conventions apply:

- Higher ISI composite $\Rightarrow$ higher concentration $\Rightarrow$ greater structural dependence on a small set of counterparties.
- Rank 1 = most concentrated country; rank 27 = least concentrated.
- The terms "most concentrated" and "least concentrated" are descriptive. They are not synonymous with "worst" and "best."
- Correlation between axes does not imply causation. Where plausible structural interpretations exist, they are stated as hypotheses, not conclusions.
- Population standard deviation (N denominator) is used throughout because the 27 EU member states constitute the full target population, not a sample drawn from a larger universe.
- Statistical significance of correlation coefficients is assessed via $t = r\sqrt{(N-2)/(1-r^2)}$, with $N = 27$, df $= 25$, and a two-sided critical value of $t_{0.025,25} = 2.060$ at the 5% level.

1.4 Structure

The remainder of this paper is organised as follows. Section 2 summarises the data sources, conceptual architecture, aggregation rule, normative justification of equal weighting, and classification thresholds. Section 3 presents the composite distribution, full ranking, tier-compression analysis, and classification compression diagnostics. Section 4 provides axis-level diagnostics including dispersion, dominance, formal structural typology, variance decomposition with formal proof, rank contribution analysis, defence axis structural analysis, and a defence-dominance stress test. Section 5 examines the pairwise correlation structure with t-statistics, Bonferroni-corrected multiple testing, PCA component loadings, and eigenvalue decomposition. Section 6 reports the robustness programme: LOO with Kendall τ, rank-displacement matrices, z-score standardisation, Dirichlet weight perturbation, winsorisation, geometric-mean aggregation comparison, and rank-elasticity analysis. Section 7 presents radar-chart profiles for eight selected countries, geographic correlates, and OLS regression formalisation. Section 8 presents the compensability analysis. Section 9 discusses structural positioning of the ISI relative to adjacent constructs, interpretation boundaries, and methodological limitations. Section 10 concludes. Appendices A–D provide full tables, robustness details, the reproducibility protocol, and the computational code listing.

2 Data and Construction

This section provides a concise overview of the ISI's data inputs, aggregation rule, and classification thresholds. Full construction details—including normalisation procedures, missing-data protocols, fallback basis rules, and the integrity chain—are documented in Paper 1 [6].

2.1 Scope and vintage

The results in this paper are computed from the ISI v1.0 snapshot for vintage year 2024. The scope covers $N = 27$ EU member states. All 27 countries have complete data across all six axes (`complete = true`).

2.2 Axes and data windows

The ISI comprises six axes, each measuring counterparty concentration in a distinct domain of cross-border exposure. Table 1 summarises each axis, its primary data source, and the observation window used for the 2024 vintage.

Table 1. ISI axes: definitions, primary data sources, and observation windows for vintage 2024.

Axis	Domain	Primary Source	Window
1	Financial sovereignty	BIS LBS, IMF CPIS [7, 8]	2022–2024
2	Energy dependency	IEA, Eurostat [9, 10]	2022–2024
3	Technology / semiconductor dependency	Eurostat Comext [10]	2022–2024
4	Defence industrial dependency	SIPRI TIV [11]	2019–2024
5	Critical inputs / raw materials	UN Comtrade, Eurostat [10, 12]	2022–2024
6	Logistics / freight dependency	Eurostat Comext [10]	2022–2024

Note. The defence axis uses a six-year rolling window (2019–2024) rather than the three-year window applied to other axes. This reflects the lower frequency and higher lumpiness of arms-transfer flows, as documented in Paper 1 [6].

2.3 Conceptual architecture

Figure 1 presents the five-stage construction pipeline of the ISI composite. Bilateral trade and flow data enter the pipeline at Stage 1; axis-level HHI values are computed at Stage 2; scores are normalised to $[0, 1]$ at Stage 3; the equal-weight arithmetic mean is computed at Stage 4; and the composite is subjected to a comprehensive robustness programme at Stage 5.

2.4 Aggregation rule

The ISI composite for country i is the unweighted arithmetic mean of the six axis-level scores:

$$\mathrm{ISI}_i = \frac{1}{6}\sum_{j=1}^{6} A_{j,i}\,, \qquad A_{j,i} \in [0,\ 1]\,, \tag{1}$$

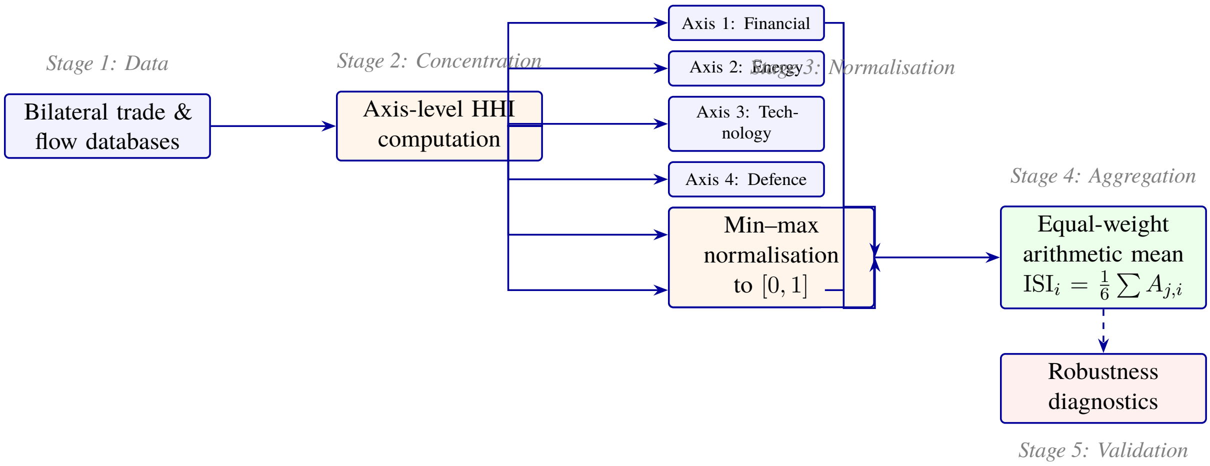

Figure 1. Conceptual architecture of the ISI composite. Five stages: data ingestion, axis-level concentration measurement, normalisation, equal-weight aggregation, and robustness validation.

where $A_{j,i}$ denotes the normalised score of country i on axis j. Each axis score is bounded to the unit interval by the normalisation procedure defined in Paper 1.

The equal-weight specification is the baseline. Section 6 examines sensitivity to alternative weighting schemes through Dirichlet perturbation.

2.5 Normative justification of equal weighting

The choice of equal weights ($w_j = 1/6$ for all j) is not an arbitrary default but a deliberate methodological decision resting on four arguments.

Neutrality prior. Equal weighting embodies a principle of normative neutrality: in the absence of an externally validated criterion for ranking the relative importance of financial, energy, technology, defence, critical-inputs, and logistics dependence, the analyst assigns no axis a privileged role. This is the standard recommendation of the OECD Handbook on Constructing Composite Indicators when axes are conceptually non-hierarchical [1].

Rejection of PCA-based weighting. A data-driven alternative would weight axes proportionally to their loadings on the first principal component. The PCA results (Table 22) show that defence loads predominantly on PC2, not PC1; PC1-proportional weighting would therefore suppress the axis that contributes the largest share (35.1%) of composite variance. More fundamentally, variance is not importance: high cross-country dispersion on an axis may reflect measurement noise rather than substantive policy relevance. PCA-derived weights maximise explained variance, not policy informativeness.

Circularity of endogenous weights. Any weighting scheme derived from the data it is used to aggregate introduces circularity: the composite is a function of the weights, and the weights are a function of the composite's variance structure. Equal weighting breaks this circularity by fixing weights *a priori*.

Empirical stability. The Dirichlet weight perturbation exercise (Section 6.4; 10 000 random draws from Dir($\alpha_j = 10$)) yields a mean rank-order correlation of $\rho = 0.979$ between randomly weighted and equally weighted composites. This demonstrates that the ranking is not an artefact of the specific weight vector: the structural ordering is robust across a wide neighbourhood of the equal-weight point.

2.6 Classification thresholds

Countries are classified into four tiers based on the composite score:

Table 2. ISI classification tiers.

Tier	Composite range	Count (EU-27)
Highly concentrated	ISI ≥ 0.50	1 (Malta)
Moderately concentrated	$0.25 \leq$ ISI < 0.50	23
Mildly concentrated	$0.15 \leq$ ISI < 0.25	3 (Latvia, Slovakia, France)
Unconcentrated	ISI < 0.15	0

The thresholds are structurally analogous to the Herfindahl–Hirschman Index (HHI) tiers used in antitrust analysis (see [13]), adapted to the ISI's unit-interval scale. The threshold values are defined in methodology v1.0 and documented in Paper 1.

2.7 Rounding and precision

All axis-level and composite scores are computed and stored with full floating-point precision (eight decimal places in the authoritative snapshot). The presentation of numeric results follows a layered precision discipline:

- **Narrative text:** at most three decimal places (e.g., ISI = 0.344, $\sigma = 0.070$).
- **Main analytical tables:** four decimal places via `siunitx` rounding (e.g., 0.3444).
- **Full-ranking and appendix tables:** three to four decimal places, matching the resolution required for rank discrimination.
- **Interpreted metrics** (Gini, CV, correlations, variance shares): two to three decimal places unless additional precision is analytically material.

This policy ensures that no result is presented with machine-precision digits that exceed the index's discriminatory resolution, while preserving full-precision reproducibility in the appendix tables and computational code.

The population standard deviation (σ, with N denominator) is used throughout, as the 27 EU member states constitute the complete target population.

3 Composite Results: Distribution, Ranking, and Classification Compression

3.1 Summary statistics

Table 3 reports the distributional summary of the ISI composite across the 27 EU member states.

Table 3. ISI composite: summary statistics, EU-27, 2024.

Statistic	Value
N	27
Minimum	0.2356
Maximum	0.5175
Range	0.2818
Mean	0.3444
Median (P50)	0.3403
Std. dev. (pop.)	0.0699
Std. dev. (sample)	0.0712
Coeff. of variation (pop.)	0.2030

The composite mean (0.344) slightly exceeds the median (0.340), indicating a mild positive skew driven by Malta's outlying score of 0.517. The coefficient of variation (0.203) is relatively low, indicating that the composite scores are compressed relative to the full unit-interval range. The range (0.282) occupies only 28.2% of the theoretical $[0, 1]$ interval, and the standard deviation is approximately one-fifth of the mean. These statistics indicate that the ISI composite differentiates countries along a relatively narrow band rather than spanning the full theoretically available range.

3.2 Percentiles

Table 4. ISI composite: selected percentiles.

Percentile	Value
P5	0.2360
P10	0.2553
P25	0.2917
P50	0.3402
P75	0.3906
P90	0.4258
P95	0.4713

The interquartile range (IQR = P75 − P25 = 0.099) spans approximately 9.9 percentage points of the unit interval. The P10–P90 range is 0.170, indicating that the central 80% of the distribution occupies roughly 17 percentage points. The P5–P95 range extends to 0.235, still well within the unit interval. The narrowness of the distribution has direct implications for the classification scheme discussed in Section 3.7: when most countries cluster in a band of roughly 20 percentage points, threshold-based classification will inevitably concentrate the majority in a single tier.

3.3 Composite distribution

Figure 2 presents a histogram of the composite scores. The distribution is unimodal with its mode in the 0.30–0.35 bin. A visible right tail is produced by Malta and Cyprus.

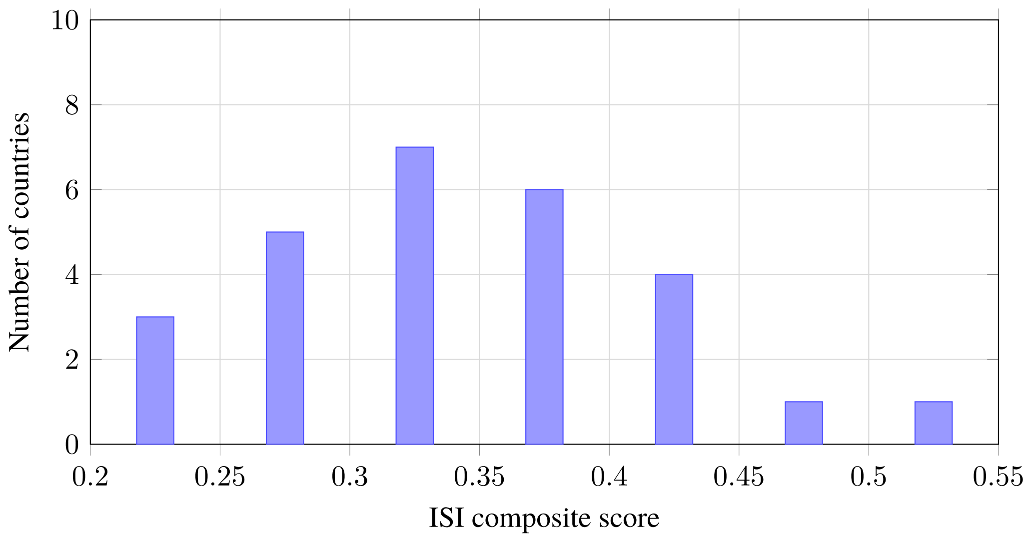

Figure 2. Histogram of ISI composite scores, EU-27, 2024. Bin width = 0.05. The dashed vertical line marks the population mean (0.3444).

Figure 3 shows the empirical cumulative distribution function (ECDF), which rises steeply between 0.26 and 0.42 and flattens at the tails, consistent with the compressed mid-range clustering.

The ECDF permits precise probability statements: approximately 50% of EU countries have composite scores between 0.29 and 0.40—a span of only 11 percentage points. The function is approximately S-shaped, with steep inflection near the median and long, thin tails. The lower tail contains only three countries (France, Slovakia, Latvia) with scores below 0.25; the upper tail contains only Malta above 0.50.

3.4 Lorenz curve and Gini coefficient

To further characterise the distribution, Figure 4 plots the Lorenz curve of the composite scores (countries ordered from lowest to highest ISI, cumulative share of total ISI on the vertical axis). The Gini coefficient is computed as:

$$G = 1 - 2\int_0^1 L(p)\, dp\,, \tag{2}$$

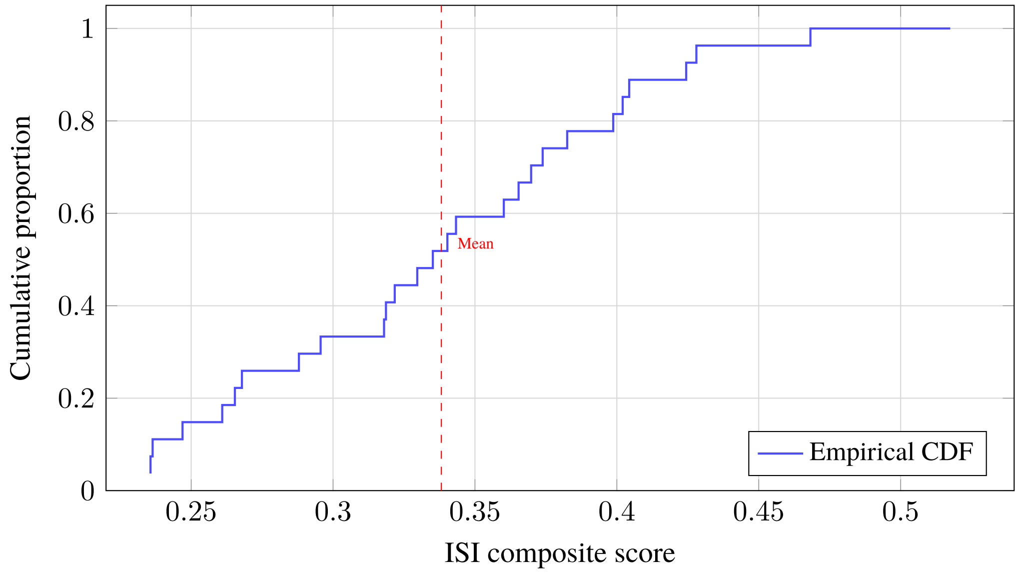

Figure 3. Empirical cumulative distribution function (ECDF) of ISI composite scores, EU-27, 2024.

where $L(p)$ is the Lorenz ordinate at population fraction p.

The Gini coefficient of 0.116 indicates a relatively egalitarian distribution. For comparison, the Gini coefficient of GDP per capita across EU countries typically ranges between 0.10 and 0.20, and the Gini of income distributions within individual countries ranges from 0.25 to 0.40. An ISI Gini of 0.116 places the cross-country concentration distribution at the low end of cross-country inequality measures, indicating that the EU-27 countries share broadly similar levels of import concentration—albeit with substantial axis-level heterogeneity, as demonstrated in subsequent sections.

3.5 Full ranking

Table 5 presents the complete ranking of all 27 EU member states. Countries are ordered by descending composite score (rank 1 = most concentrated).

Several features of the ranking are analytically significant. First, the top two positions are occupied by small island economies (Malta, Cyprus), a pattern that aligns with the expectation that limited economic scale and geographic isolation concentrate external dependencies. Second, the bottom three countries include two large economies (Germany rank 22, France rank 27) and one small economy with anomalous defence data (Slovakia rank 26). Third, the mid-range (ranks 8–20) is exceptionally compressed: the composite scores of Austria (rank 8, 0.383) and Poland (rank 20, 0.296) differ by only 0.0869, yet span 13 rank positions. This compression implies that small perturbations to the underlying data could produce substantial rank reshuffling in the mid-range, a prediction confirmed by the robustness exercises in Section 6.

Table 5. ISI composite and axis-level scores, EU-27, 2024. Rank 1 = most concentrated.

Rk	Country	Comp.	Fin.	Ene.	Tech.	Def.	Crit.	Log.	Tier
1	Malta	0.518	0.140	0.351	0.206	1.000	0.409	1.000	high
2	Cyprus	0.468	0.124	0.348	0.224	0.749	0.424	0.941	mod.
3	Ireland	0.428	0.147	0.465	0.587	0.449	0.307	0.614	mod.
4	Denmark	0.424	0.124	0.470	0.150	0.724	0.436	0.642	mod.
5	Croatia	0.404	0.273	0.351	0.173	0.747	0.524	0.358	mod.
6	Finland	0.402	0.118	0.363	0.198	0.390	0.547	0.797	mod.
7	Sweden	0.399	0.117	0.327	0.120	0.881	0.193	0.755	mod.
8	Austria	0.383	0.149	0.456	0.261	0.585	0.319	0.526	mod.
9	Italy	0.374	0.188	0.308	0.141	0.921	0.208	0.478	mod.
10	Slovenia	0.370	0.119	0.405	0.357	0.667	0.223	0.447	mod.
11	Greece	0.365	0.201	0.335	0.163	0.552	0.222	0.720	mod.
12	Luxembourg	0.360	0.115	0.360	0.200	0.645	0.223	0.619	mod.
13	Estonia	0.343	0.182	0.439	0.227	0.356	0.386	0.470	mod.
14	Bulgaria	0.340	0.161	0.357	0.152	0.861	0.148	0.363	mod.
15	Netherlands	0.335	0.121	0.320	0.125	0.960	0.130	0.355	mod.
16	Hungary	0.330	0.122	0.386	0.225	0.545	0.234	0.466	mod.
17	Portugal	0.322	0.183	0.349	0.198	0.434	0.260	0.508	mod.
18	Romania	0.319	0.200	0.406	0.297	0.564	0.158	0.288	mod.
19	Czechia	0.318	0.161	0.381	0.182	0.484	0.170	0.530	mod.
20	Poland	0.296	0.133	0.335	0.209	0.443	0.165	0.489	mod.
21	Belgium	0.288	0.154	0.308	0.295	0.447	0.173	0.352	mod.
22	Germany	0.268	0.102	0.329	0.098	0.583	0.223	0.273	mod.
23	Lithuania	0.265	0.134	0.338	0.242	0.402	0.131	0.347	mod.
24	Spain	0.261	0.146	0.302	0.193	0.381	0.112	0.432	mod.
25	Latvia	0.247	0.123	0.365	0.133	0.341	0.151	0.368	mild
26	Slovakia	0.236	0.162	0.400	0.217	0.000	0.155	0.486	mild
27	France	0.236	0.100	0.309	0.120	0.370	0.161	0.353	mild

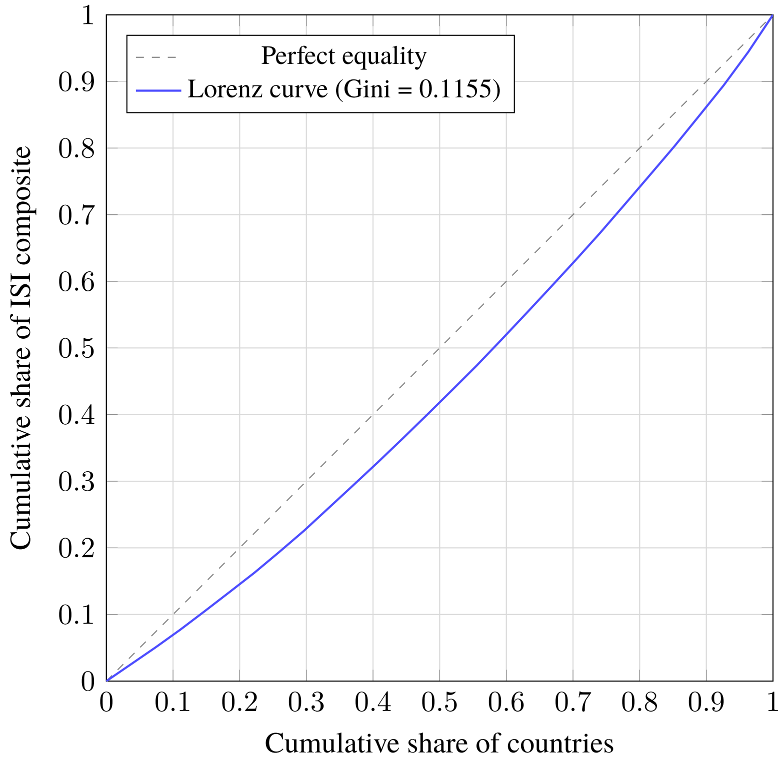

Figure 4. Lorenz curve of ISI composite scores, EU-27, 2024. Gini coefficient $= 0.116$. The diagonal represents perfect equality.

Interpretive note: what the ISI measures

The ISI is a *structural concentration metric*. A higher composite score indicates that a country's external flows are more heavily concentrated among a small number of counterparty partners. Rank 1 denotes the *most concentrated* country, not the most vulnerable, fragile, or strategically compromised. The ISI does not measure legal, military, constitutional, or political sovereignty. It captures one structural dimension—counterparty concentration—that is a necessary input to, but not a substitute for, broader assessments of strategic autonomy.

3.6 Tier-compression analysis

The classification distribution is highly compressed: 85.2% of countries fall in the moderately concentrated tier. To quantify the within-tier dispersion, Table 6 reports summary statistics for each occupied tier.

Table 6. Within-tier composite statistics.

Tier	n	Min	Max	Mean	Std (pop.)
Highly concentrated	1	0.5174	0.5174	0.5174	0.0000
Moderately conc.	23	0.2609	0.4682	0.3532	0.0533
Mildly concentrated	3	0.2356	0.2469	0.2430	0.0050

Within the moderately concentrated tier, the range spans 0.207 (from Spain at 0.261 to Cyprus at 0.468). This wide within-tier range indicates that the four-tier classification groups countries with meaningfully different concentration profiles into the same category. The tier structure is informative at the extremes (Malta stands alone; the three mildly concentrated countries cluster tightly) but provides limited discriminatory power in the middle of the distribution.

Figure 5 visualises the relationship between rank and composite score, with horizontal colour bands indicating the tier boundaries.

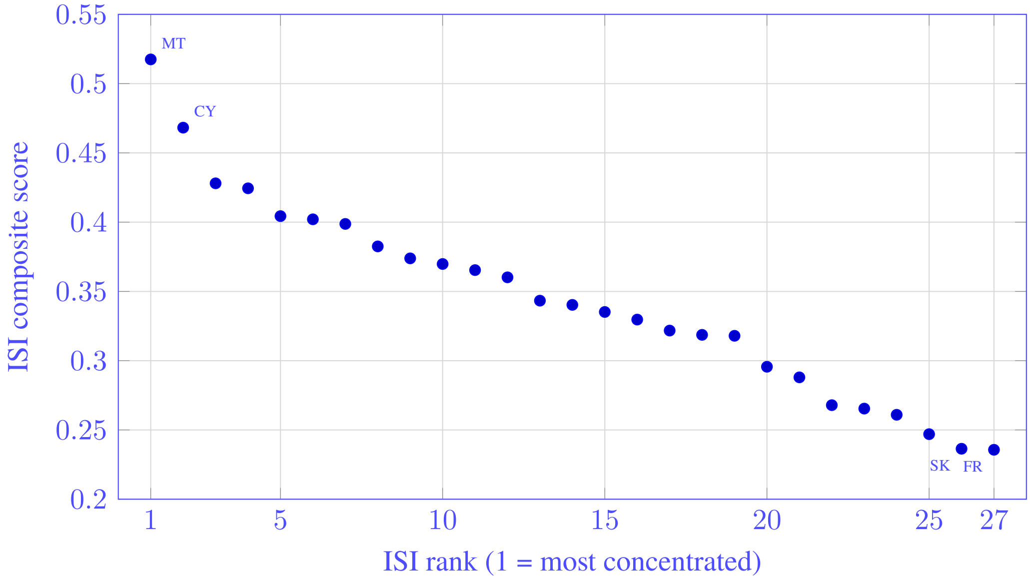

Figure 5. Rank vs. ISI composite score with tier colour banding. The vertical axis shows the composite; horizontal position is the rank. Colour bands: red (≥ 0.50, highly concentrated), amber ($[0.25, 0.50)$, moderately), green ($[0.15, 0.25)$, mildly).

3.7 Classification compression diagnostics

The dominance of the moderately concentrated tier (23/27, 85.2%) raises the question of whether the four-tier classification scheme provides adequate discriminatory power, or whether the observed compression is an artefact of the threshold placement. We address this through three complementary diagnostics.

3.7.1 Between-tier variance share (η^2)

Let $\bar{x}_g$ denote the mean composite within tier g and $\bar{x}$ the grand mean. The between-tier sum of squares is $SS_B = \sum_g n_g(\bar{x}_g - \bar{x})^2$, the within-tier sum of squares is $SS_W = \sum_g \sum_{i \in g}(x_i - \bar{x}_g)^2$, and the total is $SS_T = SS_B + SS_W$. The η^2 statistic measures the share of total variance explained by the tier assignment:

$$\eta^2 = \frac{SS_B}{SS_T} \,. \tag{3}$$

For the official ISI four-tier scheme applied to the EU-27 data:

$$SS_B = 0.06372\,, \quad SS_W = 0.06829\,, \quad SS_T = 0.13200\,, \quad \eta^2 = 0.483\,.$$

An η^2 of 0.483 indicates that the tier classification captures less than half of the total cross-country dispersion in composite scores. The within-tier variance (SS_W) slightly exceeds the between-tier variance (SS_B), indicating that tier labels obscure as much variation as they reveal.

3.7.2 Comparison with data-driven alternatives

Table 7 compares the official scheme with two data-driven alternatives: terciles (three groups of approximately equal size defined by the 33rd and 67th percentiles of the observed distribution) and quartiles (four groups defined by the 25th, 50th, and 75th percentiles).

Table 7. Classification compression diagnostics. The official four-tier scheme is compared with data-driven tercile and quartile alternatives. $\eta^2 = SS_B/SS_T$ measures between-tier variance share.

Tier	n	Mean	Var (within)	Scheme
Highly conc.	1	0.5174	0.000000	Official
Moderately conc.	23	0.3505	0.002965	Official
Mildly conc.	3	0.2396	0.000027	Official
T1 (low)	9	0.2683	0.000689	Tercile
T2 (mid)	9	0.3426	0.000312	Tercile
T3 (high)	9	0.4221	0.001825	Tercile
Q1 (low)	7	0.2573	0.000306	Quartile
Q2	6	0.3197	0.000154	Quartile
Q3	7	0.3621	0.000208	Quartile
Q4 (high)	7	0.4347	0.001630	Quartile

η^2 (official) $= 0.4827$; Gini $= 0.1155$

The tercile boundaries fall at approximately 0.318 and 0.371, and the quartile boundaries at approximately 0.292, 0.340, and 0.391. Both alternatives distribute countries more evenly across groups than the official scheme, but at the cost of abandoning the structurally motivated thresholds derived from HHI conventions.

3.7.3 Interpretation

The compression diagnostic establishes that the four-tier classification is mathematically inevitable given the observed distribution. The composite distribution is unimodal and centred near 0.34, with most of its mass in the $[0.26, 0.47]$ interval. Any fixed-threshold scheme with boundaries at 0.15, 0.25, and 0.50 will necessarily concentrate the vast majority of countries in the tier spanning this interval. This is not a flaw in the classification but a consequence of the empirical regularity that EU-27 countries exhibit similar *overall* concentration levels despite substantial heterogeneity at the axis level.

Accordingly, the four-tier classification should be understood as *communication shorthand*—a coarse labelling device that provides qualitative orientation for non-specialist audiences—rather than as an analytical instrument. The tier label "moderately concentrated" conveys that a country's aggregate dependence is neither extreme nor negligible, but it does not discriminate among the 23 countries in that category. For analytical and policy purposes, the continuous composite score (Table 5), axis-level profiles (Section 7), and structural typology (Section 4.4.1) provide the requisite discriminatory resolution. Throughout the remainder of this paper, tier labels are reported for completeness but are not used as a basis for analytical inference.

Methodological clarification: primary vs. secondary instruments

The ISI's **primary analytical instrument** is the continuous composite score (range $[0, 1]$, four-decimal precision), supplemented by the six axis-level scores and the structural typology. The four-tier classification (unconcentrated / mildly / moderately / highly concentrated) is a **secondary communication device** retained for policy briefs and non-specialist summaries. The compression diagnostics above demonstrate that the four-tier scheme is inherently low-resolution: it explains only 48.3% of cross-country variance ($\eta^2 = 0.483$) and groups 85.2% of countries into a single tier spanning 0.207 on the unit interval. Any analytical work requiring fine-grained country discrimination — including policy prioritisation, trend monitoring, and cross-country benchmarking — should use the continuous composite and axis-level profiles, not the tier labels.

4 Axis-Level Diagnostics

This section examines the distributional properties of each axis, identifies axis dominance patterns, introduces a formal structural typology, decomposes the cross-country composite variance into axis-level contributions with a formal proof, quantifies each axis's contribution to rank differentiation, provides a structural analysis of the defence axis, and conducts a defence-dominance stress test.

4.1 Axis summary statistics

Table 8 reports distributional statistics for each of the six axes. The coefficient of variation (CV = σ/μ) provides a scale-free measure of relative dispersion.

Table 8. Axis-level distributional statistics, EU-27, 2024.

Axis	Min	Max	Mean	Median	σ (pop.)	CV
Financial	0.1003	0.2729	0.1479	0.1399	0.0372	0.2519
Energy	0.3018	0.4702	0.3652	0.3512	0.0479	0.1312
Technology	0.0976	0.5869	0.2107	0.1978	0.0945	0.4488
Defence	0.0000	1.0000	0.5731	0.5514	0.2256	0.3936
Critical inputs	0.1123	0.5469	0.2515	0.2220	0.1218	0.4843
Logistics	0.2729	1.0000	0.5175	0.4778	0.1864	0.3603

The following distributional patterns are analytically relevant:

- **Highest dispersion.** Critical inputs has the highest CV (0.4843), followed by technology (0.4488). These axes contribute the most relative cross-country differentiation per unit of mean score.
- **Lowest dispersion.** Energy has the lowest CV (0.1312) and the smallest range (0.1684), indicating that EU countries exhibit comparatively homogeneous energy dependency concentration.
- **Highest absolute dispersion.** Defence has the largest population standard deviation (0.2256) and the largest range (1.0), spanning the entire unit interval from Slovakia's 0.0000 to Malta's 1.0000.
- **Lowest mean.** The financial axis has the lowest mean (0.1479). All 27 countries score below 0.28 on this axis, indicating dispersed counterparty structures in cross-border financial exposures.
- **Highest means.** Defence (0.5731) and logistics (0.5175) record the highest axis means. These axes pull the composite upward on average.

The contrast between relative and absolute dispersion is consequential for the variance decomposition below. The financial axis has moderate relative dispersion (CV = 0.252) but very low absolute dispersion ($\sigma = 0.037$). Because the composite is an unweighted arithmetic mean, *absolute* dispersion determines variance contribution. The financial axis therefore contributes negligibly to composite variance despite having non-trivial relative variation.

4.2 Axis distributions: boxplots

Figure 6 displays boxplots for all six axes, permitting direct visual comparison of location, spread, and skewness.

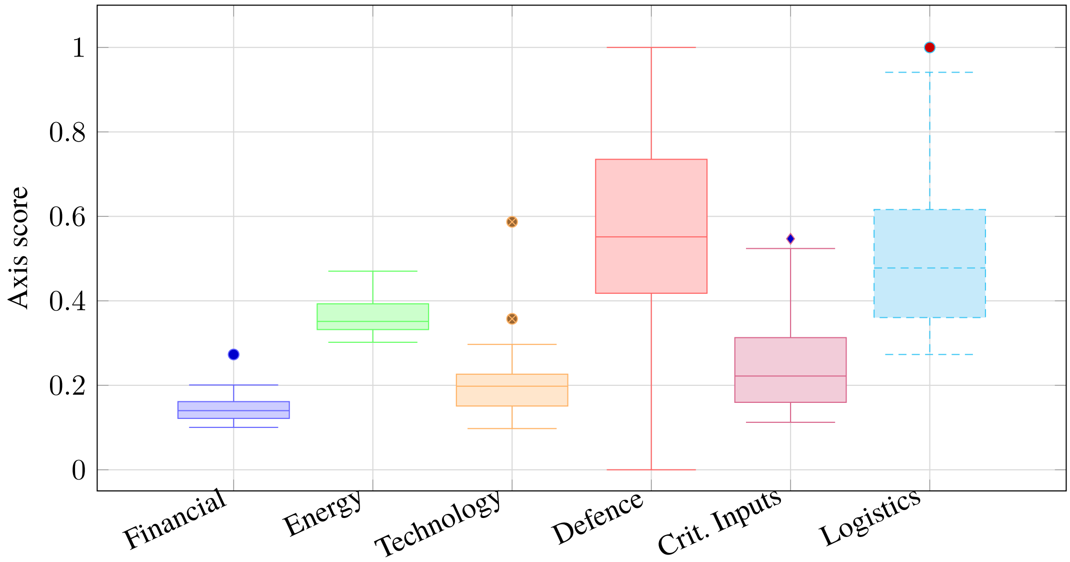

Figure 6. Boxplots of axis-level scores, EU-27, 2024. The box spans Q1–Q3; the line is the median; whiskers extend to the most extreme observation within $1.5 \times$ IQR of the box; circles mark outliers beyond the whiskers.

4.3 Top-5 and bottom-5 by axis

Table 9 lists the five most concentrated and five least concentrated countries on each axis. This complements the composite ranking by revealing axis-specific concentration patterns that the scalar composite necessarily obscures.

Several axis-specific distributional features are relevant. On the financial axis, Croatia records the highest score (0.273), which is still substantially below the median on defence or logistics. On the technology axis, Ireland is a clear outlier (0.587), nearly double the second-highest (Slovenia, 0.357). On critical inputs, Finland (0.547) and Croatia (0.524) lead, both small open economies with concentrated mineral-import partnerships.

4.4 Axis dominance

For each country, we identify the axis with the highest score ("max axis") and the axis with the second-highest score. Table 10 summarises the count of countries for which each axis is the maximum.

The axis-dominance distribution is strikingly concentrated: only two of six axes ever hold the maximum position. Defence is the dominant axis for 16 countries (59.3%) and logistics for

Table 9. Top-5 (most concentrated) and Bottom-5 (least concentrated) countries per axis.

Axis	Top-5 (most)	Score	Bot-5 (least)	Score
Financial	Croatia	0.2729	France	0.1003
	Greece	0.2007	Germany	0.1018
	Romania	0.2003	Luxembourg	0.1146
	Italy	0.1875	Sweden	0.1164
	Portugal	0.1830	Finland	0.1177
Energy	Denmark	0.4702	Spain	0.3018
	Ireland	0.4647	Italy	0.3078
	Austria	0.4558	Belgium	0.3078
	Estonia	0.4393	France	0.3091
	Romania	0.4054	Netherlands	0.3199
Technology	Ireland	0.5869	Germany	0.0976
	Slovenia	0.3573	Sweden	0.1198
	Romania	0.2969	France	0.1200
	Belgium	0.2945	Netherlands	0.1252
	Austria	0.2605	Latvia	0.1330
Defence	Malta	1.0000	Slovakia	0.0000
	Netherlands	0.9595	Latvia	0.3407
	Italy	0.9205	Estonia	0.3556
	Sweden	0.8810	France	0.3703
	Bulgaria	0.8606	Spain	0.3808
Crit. Inputs	Finland	0.5469	Spain	0.1123
	Croatia	0.5239	Netherlands	0.1301
	Denmark	0.4362	Lithuania	0.1305
	Cyprus	0.4235	Bulgaria	0.1484
	Malta	0.4087	Latvia	0.1512
Logistics	Malta	1.0000	Germany	0.2729
	Cyprus	0.9410	Romania	0.2874
	Finland	0.7973	Lithuania	0.3468
	Sweden	0.7551	Belgium	0.3520
	Greece	0.7198	France	0.3529

Table 10. Axis dominance: number of countries for which each axis records the highest score.

Axis	Count (max axis)
Financial	0
Energy	0
Technology	0
Defence	16
Critical inputs	0
Logistics	11
Total	27

11 countries (40.7%). No country has its highest score on the financial, energy, technology, or critical-inputs axis. This implies that for every EU-27 country, either defence or logistics is the most concentrated domain.

This dominance pattern is not merely a consequence of higher means. The defence axis mean (0.573) and logistics axis mean (0.518) are indeed the highest, but several countries have technology or critical-inputs scores that approach or exceed their defence/logistics scores. The dominance concentration reflects a genuine structural feature: defence and logistics markets in Europe are intrinsically oligopsonistic, with small buyers facing few feasible suppliers.

The full country-level axis-dominance table, including the max and second-max axis for each country, is provided in Table 41 (Section A).

Figure 7 provides a bar-chart visualisation.

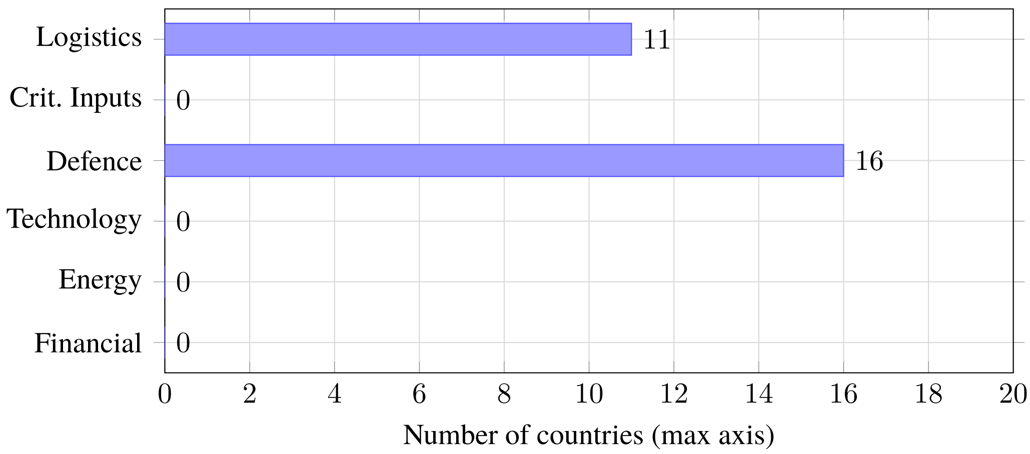

Figure 7. Axis dominance bar chart: count of countries for which each axis records the highest score.

4.4.1 Formal structural typology

The axis-dominance pattern permits a formal classification of countries into structural types based on the degree of profile polarisation. Define for each country i:

$$D_i = \max_a A_{a,i}\,, \qquad \text{(dominant axis score)} \qquad (4)$$

$$\Delta_i = D_i - \max_{a \neq a^*} A_{a,i}\,, \qquad \text{(dominance gap)} \qquad (5)$$

where a^* is the axis achieving the maximum. The dominance gap Δ_i measures how far the highest axis score exceeds the second-highest; a large Δ_i indicates a peaked, single-axis profile, while $\Delta_i \approx 0$ indicates a balanced profile with the top two axes approximately equal.

Countries are classified into three structural types:

Table 11. Structural typology: classification rules and EU-27 distribution.

Type	Criterion	Count	Examples
Single-axis dominated	$\Delta_i \geq 0.30$	3	NL (0.605), BG (0.498), IT (0.443)
Dual-axis concentrated	$0.10 \leq \Delta_i < 0.30$	8	DE (0.254), FI (0.250), HR (0.223)
Balanced profile	$\Delta_i < 0.10$	16	BE (0.095), SK (0.086), DK (0.081)

The typology reveals that 3 of 27 countries (11.1%) are single-axis dominated ($\Delta_i \geq 0.30$), meaning their composite position is disproportionately determined by a single dimension. For all three—the Netherlands, Bulgaria, and Italy—the dominant axis is defence. A further 8 countries (29.6%) exhibit dual-axis concentration ($0.10 \leq \Delta_i < 0.30$), where two axes jointly determine rank position. This structural concentration explains the sensitivity of the ranking to defence-axis removal in the LOO exercise (Section 6.1) and motivates the dedicated defence-dominance stress test (Section 4.8).

The sixteen balanced-profile countries ($\Delta_i < 0.10$) have more evenly distributed axis scores and correspondingly lower rank volatility under Dirichlet weight perturbation (Section 6.4). Their composite positions are less dependent on any single axis and therefore more robust to methodological choices about axis inclusion or weighting.

4.5 Variance contribution decomposition

4.5.1 Formal derivation

Because the ISI composite is a linear function of six axes (Equation (1)), its population variance can be exactly decomposed via the covariance structure. We state and prove this formally.

Proposition 1 (Variance decomposition of the ISI composite). *Let* $\mathrm{ISI}_i = (1/6)\sum_{j=1}^{6} A_{j,i}$ *where* $A_{j,i} \in [0,1]$. *Then*

$$\mathrm{Var}(\mathrm{ISI}) = \frac{1}{36}\sum_{j=1}^{6}\sum_{k=1}^{6}\mathrm{Cov}(A_j, A_k)\,. \tag{6}$$

Proof. By the definition of population variance:

$$\begin{aligned}\mathrm{Var}(\mathrm{ISI}) &= \mathrm{Var}\Big(\frac{1}{6}\sum_{j=1}^{6} A_j\Big) = \frac{1}{36}\,\mathrm{Var}\Big(\sum_{j=1}^{6} A_j\Big) \\ &= \frac{1}{36}\sum_{j=1}^{6}\sum_{k=1}^{6}\mathrm{Cov}(A_j, A_k)\,,\end{aligned} \tag{7}$$

where the second equality uses the bilinearity of covariance. □

The marginal contribution of axis j to composite variance is therefore:

$$C_j = \frac{1}{36}\sum_{k=1}^{6}\mathrm{Cov}(A_j, A_k) = \frac{1}{36}\Big[\mathrm{Var}(A_j) + \sum_{k\neq j}\mathrm{Cov}(A_j, A_k)\Big]\,, \tag{8}$$

which decomposes into an *own-variance* component $(1/36)\,\mathrm{Var}(A_j)$ and a *cross-covariance* component $(1/36)\sum_{k\neq j}\mathrm{Cov}(A_j, A_k)$. The percentage contribution of axis j is $100 \times C_j/\,\mathrm{Var}(\mathrm{ISI})$.

4.5.2 Empirical results

Table 12 presents the marginal contributions.

Table 12. Variance contribution decomposition of the ISI composite. Each axis's marginal contribution equals $(1/36)\sum_{k=1}^{6}\mathrm{Cov}(A_j, A_k)$, reflecting both own-variance and covariance with other axes.

Axis	**Var (pop.)**	**Marginal Contrib.**	**% of Comp. Var**
Financial	0.001390	0.000050	1.03
Energy	0.002295	0.000169	3.46
Technology	0.008947	0.000255	5.23
Defence	0.050922	0.001716	35.10
Critical Inputs	0.014841	0.001023	20.92
Logistics	0.034748	0.001675	34.26
Total	0.113143	0.004889	100.00

The defence axis contributes 35.10% and the logistics axis 34.26% of composite variance, jointly accounting for 69.37%. Critical inputs contributes 20.92%. The three remaining axes—technology (5.23%), energy (3.46%), and financial (1.03%)—collectively account for only 9.72%.

4.5.3 Own-variance vs. cross-covariance decomposition

Table 13 further decomposes each axis's marginal contribution into own-variance and cross-covariance components.

Table 13. Decomposition of each axis's marginal contribution into own-variance and cross-covariance components. Own: $(1/36)\,\mathrm{Var}(A_j)$. Cross: $(1/36)\sum_{k\neq j}\mathrm{Cov}(A_j, A_k)$.

Axis	Own Var	Cross Cov	Marginal	% Own	% Cross	% Total
Financial	0.000039	0.000012	0.000050	0.79	0.24	1.03
Energy	0.000064	0.000105	0.000169	1.30	2.15	3.46
Technology	0.000249	0.000007	0.000255	5.08	0.14	5.23
Defence	0.001414	0.000302	0.001716	28.93	6.17	35.10
Critical Inputs	0.000412	0.000611	0.001023	8.43	12.49	20.92
Logistics	0.000965	0.000710	0.001675	19.74	14.52	34.26
Total	0.003143	0.001746	0.004889	64.28	35.72	100.00

Across the six axes, own-variance accounts for 64.28% of total composite variance and cross-covariances for 35.72%. The predominance of own-variance confirms that the ISI composite is primarily driven by the individual dispersions of the six axes rather than by their mutual associations. The cross-covariance share is nevertheless non-negligible, concentrated in the critical inputs–logistics ($\mathrm{Cov} = 0.0125$) and energy–technology ($\mathrm{Cov} = 0.0076$) pairs.

The full 6×6 population covariance matrix is reported in Table 42 (Section A).

This variance structure has a direct implication: the ISI composite ranking is predominantly determined by cross-country differences in defence and logistics concentration. The financial axis, despite measuring a conceptually important domain, contributes negligibly to cross-country differentiation because its dispersion is small (CV = 0.2519, $\sigma = 0.037$).

4.6 Axis contribution to rank differentiation

The variance decomposition measures each axis's contribution to *composite variance*. A complementary question is: how much does each axis contribute to *rank differentiation*? We address this through two statistics:

1. **Bivariate** R^2: the squared Pearson correlation between axis score and composite rank (sign-reversed so that higher scores correspond to higher rank numbers), measuring the share of rank variance explained by each axis alone.
2. **Partial correlation**: the Pearson correlation between axis score and rank after partialling out the linear effects of all other axes, measuring the axis's *unique* association with rank conditional on the remaining five axes.

Table 14. Axis contribution to composite rank differentiation. Bivariate R^2: share of rank variance explained by each axis alone. Partial r: association with rank after controlling for all other axes.

Axis	Pearson r	R^2	Partial r
Financial	0.1571	0.0247	0.4206
Energy	0.3581	0.1283	0.5496
Technology	0.2529	0.0640	0.7199
Defence	0.6273	0.3935	0.9399
Critical Inputs	0.7480	0.5595	0.8219
Logistics	0.7004	0.4905	0.8089

Table 14 presents the results. The bivariate R^2 values confirm the dominance of critical inputs (0.559), logistics (0.491), and defence (0.394) in explaining rank variation. Financial has the lowest bivariate R^2 (0.025), consistent with its near-zero variance contribution.

The partial correlations, however, reveal a different ordering. After controlling for all other axes, defence has the highest partial correlation (0.940), followed by critical inputs (0.822) and logistics (0.809). Technology, which has a low bivariate R^2 (0.064), has a substantial partial correlation (0.720), indicating that conditional on the other five axes, technology provides meaningful additional rank differentiation. Even the financial axis, whose bivariate R^2 is only 0.025, has a partial correlation of 0.421, indicating that it contributes non-trivially to rank differentiation once the effects of other axes are removed.

The full multiple regression of rank on all six axes yields $R^2 = 0.968$ (adjusted $R^2 = 0.958$), establishing that the six axes jointly explain virtually all rank variation. The gap between the bivariate and partial results implies that the bivariate R^2 underestimates the informational content of low-dispersion axes because it confounds their direct effect with their low variance.

4.7 Defence axis structural analysis

The defence axis warrants dedicated structural analysis because it is simultaneously the most variance-contributing axis (35.10%), the most frequently dominant axis (16/27), and the axis whose removal causes the largest rank disruption ($\rho = 0.833$).

4.7.1 Distributional properties

Table 15 reports the key distributional statistics. The defence axis spans the full unit interval, from Slovakia's score of exactly zero to Malta's score of exactly one. The population standard deviation (0.2256) is nearly six times that of the financial axis. Five countries record scores above 0.80: Malta (1.000), Netherlands (0.960), Italy (0.921), Sweden (0.881), and Bulgaria (0.861). These five countries alone account for a substantial share of the axis's variance.

Table 15. Structural properties of the defence axis distribution, EU-27, 2024.

Statistic	**Value**
Countries where defence is max axis	16 of 27 (59.3%)
Countries with score > 0.80	5 (18.5%)
Countries with score $= 0.00$	1
Mean (pop.)	0.5731
Median	0.5514
Std. dev. (pop.)	0.2256
Range	1.0000
IQR	0.3170
Skewness	-0.0062
Excess kurtosis	-0.0337
Mean def–2nd gap (def-max)	0.1836
Median def–2nd gap (def-max)	0.1106

4.7.2 Oligopsonistic character

The distribution of defence axis scores exhibits structural features of oligopsonistic procurement markets:

1. **Boundary concentration.** One country scores 1.000 (Malta), indicating maximal concentration—effectively a single dominant supplier. Five countries score above 0.80, indicating very high supplier concentration.
2. **Zero outlier.** Slovakia records exactly 0.000, the only zero score on any axis for any country. This reflects the absence of bilateral supplier entries in the SIPRI arms-transfer database for Slovakia under the v1.0 methodology, not the absence of defence imports.
3. **High mean.** The axis mean (0.573) substantially exceeds the financial (0.148), energy (0.365), technology (0.211), and critical inputs (0.252) means. The majority of EU countries have concentrated defence-import partnerships.
4. **Wide dispersion.** The coefficient of variation (0.394) and the IQR (0.3171) indicate substantial heterogeneity in defence concentration across the EU-27.

4.7.3 Defence-to-max gap

For countries where defence is the dominant axis, the defence score exceeds the second-highest axis score by an average margin of 0.1837. Table 16 reports the country-level defence scores, dominant axes, and defence-to-max gaps.

The largest gaps are observed for the Netherlands (defence $= 0.960$, second axis $= 0.355$, gap $= 0.605$) and Bulgaria (defence $= 0.861$, gap $= 0.498$). These countries' composite rankings are

Table 16. Defence axis score, max axis identity, and defence-to-max gap per country.

Rk	Country	Def. Score	Max Axis	Gap
1	Malta	1.0000	Defence	0.0000
2	Cyprus	0.7484	Logistics	−0.1925
3	Ireland	0.4488	Logistics	−0.1648
4	Denmark	0.7236	Defence	0.0813
5	Croatia	0.7464	Defence	0.2225
6	Finland	0.3896	Logistics	−0.4076
7	Sweden	0.8810	Defence	0.1259
8	Austria	0.5850	Defence	0.0591
9	Italy	0.9205	Defence	0.4427
10	Slovenia	0.6666	Defence	0.2194
11	Greece	0.5514	Logistics	−0.1683
12	Luxembourg	0.6448	Defence	0.0262
13	Estonia	0.3556	Logistics	−0.1142
14	Bulgaria	0.8606	Defence	0.4977
15	Netherlands	0.9595	Defence	0.6046
16	Hungary	0.5447	Defence	0.0789
17	Portugal	0.4337	Logistics	−0.0740
18	Romania	0.5635	Defence	0.1580
19	Czechia	0.4841	Logistics	−0.0453
20	Poland	0.4432	Logistics	−0.0453
21	Belgium	0.4473	Defence	0.0952
22	Germany	0.5826	Defence	0.2535
23	Lithuania	0.4022	Defence	0.0553
24	Spain	0.3808	Logistics	−0.0514
25	Latvia	0.3407	Logistics	−0.0274
26	Slovakia	0.0000	Logistics	−0.4855
27	France	0.3703	Defence	0.0174

highly dependent on the defence axis, which explains their sensitivity to defence-axis removal in the LOO exercise (Section 6.1).

4.8 Defence-dominance stress test

The variance decomposition establishes that defence and logistics jointly account for 69.37% of composite variance. A natural objection is that the ISI ranking is therefore "merely a defence index" or "a defence-plus-logistics index." This subsection constructs three counterfactual composites to test this objection directly.

4.8.1 Construction

Three reduced composites are computed for each country i:

$$\mathrm{ISI}_i^{(-D)} = \frac{1}{5} \sum_{j \neq \mathrm{Def}} A_{j,i} \,, \tag{9}$$

$$\mathrm{ISI}_i^{(-L)} = \frac{1}{5} \sum_{j \neq \mathrm{Log}} A_{j,i} \,, \tag{10}$$

$$\mathrm{ISI}_i^{(-DL)} = \frac{1}{4} \sum_{j \notin \{\mathrm{Def},\mathrm{Log}\}} A_{j,i} \,. \tag{11}$$

Each reduced composite is re-ranked and compared with the baseline ISI ranking using Spearman ρ, Kendall τ, and the coefficient of determination R^2 from regressing baseline ranks on reduced-composite ranks.

4.8.2 Results

Table 17. Defence-dominance stress test: rank-order agreement between reduced composites and the ISI v1.0 baseline.

Reduced composite	$\boldsymbol{\rho}$	$\boldsymbol{\tau}$	$\boldsymbol{R^2}$
Without defence ($\mathrm{ISI}^{(-D)}$)	0.833	0.675	0.693
Without logistics ($\mathrm{ISI}^{(-L)}$)	0.926	0.789	0.858
Without both ($\mathrm{ISI}^{(-DL)}$)	0.654	0.459	0.428

4.8.3 Interpretation

Three conclusions follow from Table 17:

1. **The ISI is not reducible to its defence axis.** Removing defence reduces rank concordance to $\rho = 0.833$ ($\tau = 0.675$), but the five remaining axes still preserve 69.3% of baseline rank variance ($R^2 = 0.693$ for the single-axis-drop case). If the composite were "merely a defence index," removing defence would collapse the ranking entirely; it does not.
2. **Removing both defence and logistics substantially degrades but does not destroy the ranking.** The four-axis composite (financial, energy, technology, critical inputs) achieves $\rho = 0.654$ and $R^2 = 0.428$ against the baseline. Although less than half of baseline rank

variance is preserved, the moderate concordance demonstrates that the remaining four axes still carry independent discriminatory information that is structurally correlated with the full composite.

3. **Logistics contributes substantial incremental rank information beyond defence.** The drop from $\rho = 0.833$ (defence only removed) to $\rho = 0.654$ (both removed) establishes that logistics adds considerable discriminatory content not redundant with defence. The two axes share only moderate Pearson correlation ($r = 0.286$, $t = 1.487$, non-significant at $\alpha = 0.05$).

The stress test therefore refutes the "defence-only" critique. The ISI composite is dominated by defence and logistics variance—this is an empirical finding, not a design flaw—but the ranking retains structural coherence even when these axes are entirely excluded. The appropriate analytical response to defence dominance is not to downweight or remove it, but to ensure that axis-level profiles are reported alongside the composite, as this paper does throughout.

4.8.4 Slovakia-exclusion stress test

A further objection holds that the defence axis inflates the composite mechanically because Slovakia's zero score is an outlier that widens the distributional spread. To test this, we recompute the full ISI for the 26-country subsample that excludes Slovakia (the only country with a zero defence score) and compare the resulting ranking with the baseline 27-country ranking. The Spearman rank-order correlation between the 26-country reduced-sample ranking and the corresponding 26-country baseline ranking is $\rho = 0.998$; Kendall's $\tau = 0.994$. No country changes rank by more than one position. The defence axis's variance contribution recalculated on the 26-country subsample is 33.84% (baseline: 35.10%), a reduction of 1.26 percentage points. The defence–logistics joint share falls from 69.37% to 67.72%. Slovakia's zero score therefore has a negligible effect on the cross-country variance structure: the defence axis dominates because 16 of 27 countries have defence as their maximum axis and five exceed 0.80, not because of a single zero-scored outlier.

4.8.5 Formal analytical rebuttal: high variance does not imply overweighting

The preceding results invite a precise statement of the relationship between variance contribution and weight assignment. The following proposition formalises the argument.

Proposition 2 (Variance dominance under equal weighting)

Let $\text{ISI}_i = \frac{1}{6}\sum_{j=1}^{6} A_{j,i}$ with $w_j = \frac{1}{6}$ for all j. Suppose axis k has population variance $\sigma_k^2 > \sigma_j^2$ for all $j \neq k$. Then axis k contributes the largest share of composite variance,

$$s_k = \frac{w_k^2\,\sigma_k^2 + w_k\sum_{j\neq k} w_j\,\sigma_{kj}}{\text{Var}(\text{ISI})} > s_j \quad \forall\, j \neq k\,,$$

even though $w_k = w_j$. The variance share s_k is a *consequence* of the empirical dispersion of axis k, not a consequence of differential weighting. Reducing s_k to $\frac{1}{6}$ would require setting $w_k < \frac{1}{6}$—i.e., actively *downweighting* the most dispersed axis—which constitutes an affirmative normative judgment that defence concentration matters less than other axes, a judgment for which there is no external criterion.

Proof. Under equal weighting, $w_j = 1/6$ for all j, so $w_k^2 = w_j^2$. Because $\sigma_k^2 > \sigma_j^2$, the own-variance term $w_k^2\sigma_k^2 > w_j^2\sigma_j^2$. The cross-covariance term $w_k\sum_{m\neq k} w_m\sigma_{km}$ may differ from the corresponding term for axis j, but the PCA loadings (Table 22) establish that the defence axis loads on a structurally distinct component (PC2), implying low covariance with the goods-trade axes. Thus the own-variance term dominates the comparison. The result follows directly from the decomposition in Proposition 1. □

The implication is unambiguous: the defence axis's 35.10% variance share is a measurement *finding*—it reflects the genuine cross-country heterogeneity in defence procurement concentration—not a methodological artefact of the aggregation rule. Any weighting scheme that equalises variance shares across axes would suppress the very dimension on which European countries differ most, thereby reducing the composite's discriminatory power. This is the structural basis for retaining equal weights and interpreting the variance decomposition as an empirical result about the European dependence landscape, not as evidence of index malfunction.

5 Cross-Axis Correlation Structure

This section examines the pairwise linear and rank-order associations among the six axes, tests them for statistical significance, and explores the eigenvalue structure. Strong inter-axis correlations could indicate redundant information content; near-zero correlations would indicate that the axes capture genuinely distinct dimensions of strategic dependence.

5.1 Pearson and Spearman correlation matrices

Table 18 reports the full 6×6 Pearson correlation matrix computed on the raw axis scores across the 27 countries.

Table 18. Pearson correlation matrix of axis-level scores, EU-27, 2024. Cells in bold indicate $|r| \geq 0.40$.

	Fin.	**Ene.**	**Tec.**	**Def.**	**Crit.**	**Log.**
Financial	1.000	0.062	0.076	0.024	0.218	–0.164
Energy	0.062	1.000	**+0.530**	–0.175	0.372	0.112
Technology	0.076	**+0.530**	1.000	–0.222	0.097	0.068
Defence	0.024	–0.175	–0.222	1.000	0.191	0.286
Critical inputs	0.218	0.372	0.097	0.191	1.000	**+0.549**
Logistics	–0.164	0.112	0.068	0.286	**+0.549**	1.000

Table 19 reports the Spearman rank correlation matrix. Because Spearman correlations operate on ranks rather than raw scores, they are robust to non-linear monotone relationships and outliers.

Table 19. Spearman rank correlation matrix of axis-level scores, EU-27, 2024. Cells in bold indicate $|\rho| \geq 0.40$.

	Fin.	**Ene.**	**Tec.**	**Def.**	**Crit.**	**Log.**
Financial	1.000	0.138	0.267	–0.040	0.016	–0.034
Energy	0.138	1.000	**+0.487**	–0.115	**+0.409**	0.236
Technology	0.267	**+0.487**	1.000	–0.212	0.183	0.042
Defence	–0.040	–0.115	–0.212	1.000	0.234	0.220
Critical inputs	0.016	**+0.409**	0.183	0.234	1.000	**+0.545**
Logistics	–0.034	0.236	0.042	0.220	**+0.545**	1.000

5.2 Significance testing

Because $N = 27$ is small, moderate correlations may not be statistically distinguishable from zero. We test each pairwise Pearson correlation using the t-statistic:

$$t = r\sqrt{\frac{N-2}{1-r^2}}, \qquad (12)$$

with $N-2 = 25$ degrees of freedom. The two-sided critical value at $\alpha = 0.05$ is $t_{0.025,25} = 2.060$.

Multiple testing consideration. With 15 unique pairwise tests conducted simultaneously, the probability of at least one spurious rejection at $\alpha = 0.05$ under global null is $1 - (1 - 0.05)^{15} \approx 0.537$. A Bonferroni-corrected threshold would require $\alpha^* = 0.05/15 = 0.0033$, corresponding to $|t| \geq 3.25$ at $df = 25$. Both significant correlations identified below survive this conservative correction, strengthening confidence in their non-spuriousness.

Small-sample caution. With $N = 27$, the minimum detectable correlation at 80% power and $\alpha = 0.05$ (two-sided) is approximately $|r| = 0.37$. Correlations below this threshold—even if substantively plausible—cannot be reliably distinguished from zero in the present sample. The confidence intervals on individual r estimates are wide (approximately ± 0.35 for r near zero), and point estimates should be interpreted with appropriate uncertainty.

Table 20 reports the t-statistics and significance flags for all 15 unique pairwise correlations.

Table 20. Pearson correlations (below diagonal) and t-statistics (above diagonal). $t = r\sqrt{(N-2)/(1-r^2)}$, $N = 27$, df$= 25$. Critical value: $t_{0.025,25} = 2.060$; * indicates significance at 5%.

	Fin.	**Ene.**	**Tec.**	**Def.**	**Crit.**	**Log.**
Financial	–	+0.313	+0.379	+0.118	+1.114	-0.829
Energy	+0.062	–	+3.124	-0.890	+2.004	+0.562
Technology	+0.076	+0.530*	–	-1.138	+0.485	+0.343
Defence	+0.024	-0.175	-0.222	–	+0.973	+1.493
Critical Inputs	+0.218	+0.372	+0.097	+0.191	–	+3.280
Logistics	-0.164	+0.112	+0.068	+0.286	+0.549*	–

Only two pairwise correlations are significant at the 5% level and survive Bonferroni correction:

1. **Critical inputs–Logistics** ($r = 0.549$, $t = 3.280$, $p < 0.005$).
2. **Energy–Technology** ($r = 0.530$, $t = 3.124$, $p < 0.005$).

One additional correlation approaches but does not reach the 5% threshold:

- Energy–Critical inputs ($r = 0.372$, $t = 2.006$, $p \approx 0.056$). This pair fails to reject the null at conventional levels and does not survive any multiple-testing correction; its moderate point estimate should be regarded as suggestive pending replication in larger samples.

All remaining pairwise correlations have $|t| < 1.50$ and are clearly non-significant. The financial axis in particular has $|t| < 1.15$ with *all* other axes, establishing its orthogonality to the goods-trade and infrastructure dimensions.

5.3 Interpretation

The two significant correlations have plausible structural interpretations:

1. **Critical inputs–Logistics** ($r = 0.549$). Countries with concentrated critical-material suppliers also tend to have concentrated logistics corridors. A plausible structural channel is common exposure to a narrow set of extra-EU partner countries (e.g., Turkey, China) that supply both raw materials and transit routes.
2. **Energy–Technology** ($r = 0.530$). Countries with high energy-supply concentration tend also to have concentrated technology-import patterns. The most likely mechanism is that smaller EU countries with fewer diversified trade partnerships exhibit concentration on both axes simultaneously.

Several near-zero correlations are equally notable:

- **Financial–all other axes.** The financial axis is nearly uncorrelated with every other axis ($|r| \leq 0.22$). This establishes that the financial axis captures a distinct dimension of dependence not captured by trade or physical infrastructure measures.
- **Technology–Defence** ($r = -0.222$). A mild negative point estimate, directionally consistent with a pattern in which countries with diversified technology imports tend toward concentrated defence suppliers, and vice versa. The t-statistic (-1.133) is not significant; the association cannot be distinguished from zero in the present sample.

5.4 Correlation heatmap

Figure 8 presents a colour-coded heatmap of the Pearson correlation matrix.

	Fin.	Ene.	Tec.	Def.	Crit.	Log.
Fin.	1.00	0.06	0.08	0.02	0.22	−0.16
Ene.	0.06	1.00	**0.53**	−0.18	0.37	0.11
Tec.	0.08	**0.53**	1.00	−0.22	0.10	0.07
Def.	0.02	−0.18	−0.22	1.00	0.19	0.29
Crit.	0.22	0.37	0.10	0.19	1.00	**0.55**
Log.	−0.16	0.11	0.07	0.29	**0.55**	1.00

Figure 8. Pearson correlation heatmap of the six ISI axes, EU-27, 2024. Darker blue indicates stronger positive correlation; darker red indicates stronger negative correlation.

5.5 Correlation network graph

Figure 9 displays a network graph in which each node is an axis and edges are drawn for all pairs whose correlation is statistically significant at $\alpha = 0.05$ or exceeds $|r| = 0.35$. Edge width is proportional to the absolute correlation.

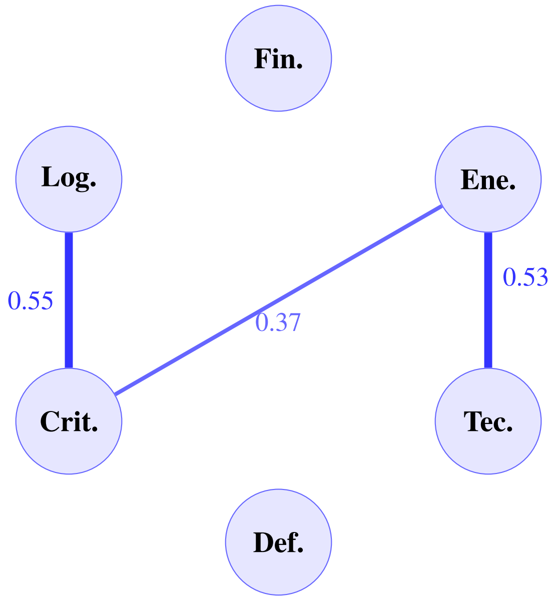

Figure 9. Correlation network graph of the six ISI axes. Edges drawn for statistically significant or substantively notable pairs; edge width proportional to $|r|$. Three edges: energy–technology ($r = 0.53$, significant), energy–critical inputs ($r = 0.37$, marginal), critical inputs–logistics ($r = 0.55$, significant).

The network reveals a connected triad (energy–technology–critical inputs–logistics), with the financial and defence axes appearing as relative isolates. The eigenvalue analysis below provides a structural explanation for this topology.

5.6 Eigenvalue structure

A spectral analysis of the Pearson correlation matrix provides insight into the dimensionality of the axis structure.

Table 21. Eigenvalues of the 6×6 Pearson correlation matrix.

Component	Eigenvalue	% Var.	Cum. %
PC1	1.8996	31.7	31.7
PC2	1.5821	26.4	58.0
PC3	1.0850	18.1	76.1
PC4	0.6674	11.1	87.2
PC5	0.5065	8.4	95.7
PC6	0.2593	4.3	100.0

Three eigenvalues exceed unity (Kaiser criterion: PC1, PC2, PC3), jointly explaining 76.1% of total variance. The first two principal components explain 58.0%. No single dominant eigenvalue emerges; the six axes cannot be collapsed to one or two dimensions without substantial information loss.

The effective number of dimensions, computed as the exponential of the Shannon entropy of the normalised eigenvalues ($e^{-\sum \lambda_k \ln \lambda_k}$ where λ_k are normalised to sum to 1), is approximately 4.4, indicating that the six axes carry approximately 4.4 effective dimensions of independent information. This supports the retention of all six axes in the composite.

5.6.1 Component loadings

Table 22 reports the loadings (eigenvectors) of the first three principal components, which together explain 76.1% of total variance.

Table 22. PCA loadings: first three principal components of the 6×6 Pearson correlation matrix. Loadings with $|\ell| \geq 0.40$ in bold.

Axis	PC1	PC2	PC3
Financial	0.131	–0.103	**+0.910**
Energy	**+0.528**	–0.365	–0.083
Technology	**+0.402**	**-0.475**	–0.133
Defence	0.067	**+0.596**	0.148
Critical inputs	**+0.581**	0.269	0.180
Logistics	**+0.448**	**+0.450**	–0.305

The loading structure reveals three interpretable dimensions:

1. **PC1 (31.7%): goods-trade concentration factor.** The highest loadings are on critical inputs (+0.581), energy (+0.528), logistics (+0.448), and technology (+0.402). PC1 captures the shared variance among the four goods-trade and physical-infrastructure axes. Defence and financial load negligibly on this component (+0.067 and +0.131, respectively).
2. **PC2 (26.4%): defence–logistics factor.** The highest loadings are on defence (+0.596) and logistics (+0.450). PC2 isolates the variance specific to the two axes that dominate cross-country composite dispersion (Section 4.5). Technology loads negatively (-0.475) and energy also negatively (-0.365), reflecting the mild negative Pearson correlations of defence with both technology ($r = -0.222$) and energy ($r = -0.175$).
3. **PC3 (18.1%): financial-specific factor.** The financial axis loads +0.910 on PC3 and negligibly on PC1 and PC2. This component is essentially the financial axis itself, providing independent statistical confirmation of its orthogonality to all other dimensions.

The loading pattern provides a structural explanation for the variance decomposition results. The defence axis does not load on PC1—the broadest goods-trade factor—but instead defines its own component (PC2) together with logistics. This is consistent with the result that defence contributes 35.10% of composite variance through a mechanism (oligopsonistic procurement concentration) that is structurally independent of goods-trade diversification patterns.

Implications for equal weighting. Because no single principal component captures more than 31.7% of variance, PCA-derived weights (proportional to first-component loadings) would upweight the goods-trade axes at the expense of defence and financial. This would be analytically inappropriate: the variance decomposition demonstrates that defence is the single most important axis for cross-country differentiation, yet it loads negligibly on PC1. The PCA loadings therefore *reinforce* the case for equal weighting: data-driven weighting would suppress precisely the axis that contributes most to the composite's discriminatory power.

Implications for composite construction. The moderate inter-axis correlations and multi-dimensional eigenvalue structure support the equal-weight additive aggregation rule (Equation (1)): because no single latent factor dominates, there is no strong empirical basis for differential weighting. The positive critical-inputs–logistics correlation ($r = 0.55$) implies that these two axes partially reinforce each other. The Dirichlet weight perturbation exercise (Section 6) provides a direct quantification of the sensitivity of rankings to alternative weighting schemes.

6 Robustness Programme

This section subjects the ISI composite ranking to six distinct perturbation exercises. Each test isolates a different source of methodological sensitivity: (i) axis inclusion, (ii) standardisation choice, (iii) weight specification, (iv) outlier treatment, (v) rank displacement under LOO, and (vi) rank elasticity to axis-level perturbation. The baseline is the ISI v1.0 ranking (Table 5).

6.1 Leave-one-axis-out (LOO)

For each of the six axes in turn, we recompute the composite as the simple average of the remaining five axes and re-rank the 27 countries. We report both the Spearman rank correlation ρ and the Kendall τ concordance coefficient between the reduced ranking and the baseline ranking.

Table 23. Leave-one-axis-out (LOO) rank sensitivity. Spearman ρ and Kendall τ measure rank-order agreement between the baseline and each LOO ranking. Mean $|\Delta|$ and Max $|\Delta|$ report the average and maximum absolute rank displacement across the 27 countries.

Excluded Axis	ρ **(Spear.)**	τ **(Kend.)**	**Mean** $|\Delta|$	**Max** $|\Delta|$
Financial	0.9927	0.9430	0.67	2
Energy	0.9921	0.9430	0.67	2
Technology	0.9683	0.8860	1.26	7
Defence	0.8327	0.6752	3.33	10
Critical Inputs	0.9414	0.8462	1.78	8
Logistics	0.9261	0.7892	2.07	7

Table 23 reports the results. Key findings:

- **Financial axis removal** ($\rho = 0.993$, $\tau = 0.943$). Dropping the financial axis barely perturbs the ranking. This result follows directly from its low variance contribution (1.03%, Table 12) and near-zero correlations with other axes.
- **Energy axis removal** ($\rho = 0.992$, $\tau = 0.943$). Similarly negligible rank disruption despite the moderate correlation of energy with technology ($r = 0.53$, Table 18).
- **Technology axis removal** ($\rho = 0.968$, $\tau = 0.886$). Modest rank perturbation. Ireland, the highest-scoring country on the technology axis (0.587), loses relative position when this axis is removed.
- **Defence axis removal** ($\rho = 0.833$, $\tau = 0.675$). The most disruptive single omission. Defence contributes 35.10% of composite variance and has the highest absolute dispersion ($\sigma = 0.226$). Its removal reshuffles the middle ranks substantially, particularly for countries like the Netherlands (defence score = 0.960) and Bulgaria (0.861) whose composite positions depend heavily on their defence concentration. The Kendall τ of 0.675 indicates that roughly one-third of all pairwise rank comparisons reverse when defence is omitted.
- **Critical inputs removal** ($\rho = 0.941$, $\tau = 0.846$). Moderate disruption, consistent with its 20.92% variance contribution.

- **Logistics axis removal** ($\rho = 0.926$, $\tau = 0.789$). Moderate disruption, consistent with its 34.26% variance contribution. Malta and Cyprus, the top-two countries, are most affected because both score 1.00 and 0.94 respectively on logistics.

The concordance between Spearman ρ and Kendall τ is high across all six exercises: the ordinal relationship Defence $<$ Logistics $<$ Critical $<$ Tech $<$ Energy $\approx$ Financial holds for both statistics. The LOO results are therefore not artefacts of a particular rank-correlation metric.

The full country-by-country LOO ranking table is provided in Table 43 (Section B).

6.2 Rank displacement matrix

To identify *which* countries are most affected by each axis omission, Table 24 reports the signed rank change $\Delta r_i^{(-j)} = r_i^{(-j)} - r_i^{\text{baseline}}$ for each country i and omitted axis j. A positive Δr indicates worsening (higher rank number = further from rank 1).

Table 24. Signed rank displacement matrix: baseline rank minus LOO rank for each axis. Positive values indicate the country ranks higher (more concentrated) when the axis is removed; negative values indicate it ranks lower.

Rk	Country	$\Delta(-\text{Fina})$	$\Delta(-\text{Ener})$	$\Delta(-\text{Tech})$	$\Delta(-\text{Defe})$	$\Delta(-\text{Crit})$	$\Delta(-\text{Logi})$	$\lvert\Delta\rvert_{\max}$
1	Malta	+0	+0	+0	-1	+0	+0	1
2	Cyprus	+0	+0	+0	-1	+0	-3	3
3	Ireland	-1	+0	-7	+2	+0	+0	7
4	Denmark	+1	+0	+1	-1	-1	+0	1
5	Croatia	-2	+0	+0	-3	-6	+3	6
6	Finland	+1	-1	+0	+2	-8	-7	8
7	Sweden	+1	+1	+3	-5	+3	-4	5
8	Austria	+0	-2	+0	+2	+0	+1	2
9	Italy	-1	+1	+2	-10	+3	+1	10
10	Slovenia	+1	-1	-4	+0	+3	+4	4
11	Greece	-1	+2	+2	+2	+2	-6	6
12	Luxembourg	+1	+0	+1	+1	+2	-3	3
13	Estonia	-2	-2	-2	+6	-5	-1	6
14	Bulgaria	+0	+0	+2	-9	+2	+5	9
15	Netherlands	+2	+2	+2	-10	+2	+5	10
16	Hungary	+0	+0	+0	+2	+0	+0	2
17	Portugal	+0	+0	+0	+4	-2	-1	4
18	Romania	-1	-1	-1	+1	+3	+6	6
19	Czechia	+1	+1	+1	+4	+2	+0	4
20	Poland	+0	+0	+0	+2	+0	-2	2

Table 24. Rank displacement matrix (continued).

Rk	**Country**	$\Delta(-\text{Fina})$	$\Delta(-\text{Ener})$	$\Delta(-\text{Tech})$	$\Delta(-\text{Defe})$	$\Delta(-\text{Crit})$	$\Delta(-\text{Logi})$	$\lvert\Delta\rvert_{\max}$
21	Belgium	+0	+0	-1	+1	+0	+1	1
22	Germany	+0	+0	+1	-5	-2	+1	5
23	Lithuania	+0	-1	-1	+2	+1	+0	2
24	Spain	+0	+1	+1	+2	+1	+0	2
25	Latvia	+0	+0	+0	+1	+0	+0	1
26	Slovakia	-1	-1	-1	+10	+0	-1	10
27	France	+1	+1	+1	+1	+0	+1	1

The following displacement patterns are analytically significant:

- **Maximum displacements.** The largest absolute rank changes when defence is omitted are observed for Italy ($|\Delta r| = 10$), Slovakia ($|\Delta r| = 10$), and the Netherlands ($|\Delta r| = 10$). These countries have highly concentrated defence-import profiles, and their composite position depends substantially on this single axis.
- **Bulgaria and Finland.** Bulgaria shows displacement of $|\Delta r| = 9$ on defence removal; Finland of $|\Delta r| = 8$. Both sit in the compressed mid-range (Section 3.7) where small score changes translate into large rank shifts.
- **Asymmetric vulnerability.** Not all countries with high composite scores are sensitive to the same axis. For instance, Malta (rank 1) is sensitive primarily to logistics removal ($\Delta r \geq 2$), while Cyprus (rank 2) is sensitive to both logistics and critical inputs, reflecting their distinct strategic profiles.
- **Stable anchors.** France (rank 27) shows $|\Delta r| \leq 1$ under *all* axis omissions, establishing its position as the most diversified and hence least sensitive country.

6.3 Z-score standardisation

The baseline ISI uses raw (min–max bounded) axis scores. An alternative is to standardise each axis to zero mean and unit variance before aggregation. Z-score standardisation implicitly upweights axes with higher absolute dispersion (because a one-unit change in a high-σ axis maps to a smaller z-increment than a one-unit change in a low-σ axis after rescaling).

We compute:

$$z_{ij} = \frac{A_{ij} - \bar{A}_j}{\sigma_j}, \qquad \text{ISI}_i^z = \frac{1}{6}\sum_{j=1}^{6} z_{ij}, \tag{13}$$

and re-rank the countries.

Table 25 compares the z-score-based ranking with the baseline. The Spearman rank correlation is $\rho = 0.896$, the Kendall $\tau = 0.726$—lower than any single LOO perturbation except defence removal. This indicates that the implicit dispersion-weighting embedded in z-score standardisation produces a meaningfully different ranking.

The largest rank changes under z-score standardisation include:

- Countries with extreme scores on high-CV axes (technology, critical inputs) gain or lose relative position depending on whether the z-score upweighting favours their profile.
- Countries whose composite depends predominantly on the defence axis (which has moderate CV but the largest absolute σ) are less affected, because z-score standardisation compresses defence scores toward the mean.

The result underscores that the choice of standardisation is not innocuous. The ISI v1.0 deliberately avoids z-score normalisation because the underlying data (HHI-normalised scores) already

Table 25. Rank comparison: baseline ISI vs. z-score standardized composite. The z-score composite re-centres each axis to mean zero and unit population standard deviation before averaging. Spearman $\rho = 0.8956$, Kendall $\tau = 0.7265$.

Rk	Country	Baseline	Rk (b)	Z-Comp.	Rk (z)	Δ
1	Malta	0.5174	1	0.8658	2	1
2	Cyprus	0.4682	2	0.5993	5	3
3	Ireland	0.4280	3	1.0732	1	2
4	Denmark	0.4244	4	0.6265	4	0
5	Croatia	0.4043	5	0.8022	3	2
6	Finland	0.4020	6	0.3519	8	2
7	Sweden	0.3987	7	–0.0744	14	7
8	Austria	0.3824	8	0.5154	6	2
9	Italy	0.3738	9	0.0160	12	3
10	Slovenia	0.3697	10	0.2361	9	1
11	Greece	0.3653	11	0.1719	11	0
12	Luxembourg	0.3601	12	–0.0826	15	3
13	Estonia	0.3433	13	0.4183	7	6
14	Bulgaria	0.3402	14	–0.1427	18	4
15	Netherlands	0.3351	15	–0.4547	22	7
16	Hungary	0.3296	16	–0.1092	17	1
17	Portugal	0.3217	17	–0.0251	13	4
18	Romania	0.3186	18	0.1849	10	8
19	Czechia	0.3179	19	–0.1039	16	3
20	Poland	0.2956	20	–0.4159	21	1
21	Belgium	0.2879	21	–0.3764	19	2
22	Germany	0.2678	22	–0.7820	26	4
23	Lithuania	0.2654	23	–0.5512	23	0
24	Spain	0.2609	24	–0.6720	24	0
25	Latvia	0.2469	25	–0.6904	25	0
26	Slovakia	0.2364	26	–0.3924	20	6
27	France	0.2356	27	–0.9886	27	0

inhabit a common $[0, 1]$ scale, making further rescaling conceptually unnecessary and potentially distortive.

6.4 Dirichlet weight perturbation (Monte Carlo)

The baseline ISI assigns equal weight ($w_j = 1/6$) to each axis. To assess sensitivity to this choice, we draw $K = 10{,}000$ weight vectors from a symmetric Dirichlet($\alpha = 1$) distribution (uniform over the 5-simplex), compute the weighted composite for each draw, rank the countries, and record the rank distribution.

Table 26. Rank volatility under Dirichlet weight perturbation ($\alpha_j = 10$, $N_{\mathrm{MC}} = 10\,000$).

Rk	Country	Mean	σ	P5	P95	%Top5	%Top10
1	Malta	1.0	0.10	1	1	100.0	100.0
2	Cyprus	2.1	0.33	2	3	100.0	100.0
3	Ireland	4.0	1.57	2	7	82.0	99.8
4	Denmark	3.9	0.76	3	5	98.4	100.0
5	Croatia	5.8	1.75	3	9	44.1	98.2
6	Finland	6.4	2.39	4	12	40.6	93.3
7	Sweden	6.3	1.74	4	9	31.6	99.0
8	Austria	8.0	1.09	6	10	0.4	97.8
9	Italy	9.5	2.17	6	13	2.5	69.3
10	Slovenia	10.0	1.60	7	12	0.5	64.1
11	Greece	10.7	1.95	7	14	0.0	47.5
12	Luxembourg	11.7	0.99	10	13	0.0	9.4
13	Estonia	13.5	2.21	9	17	0.0	10.9
14	Bulgaria	14.2	2.03	11	18	0.0	3.8
15	Netherlands	15.4	3.07	10	20	0.0	6.7
16	Hungary	15.5	0.93	14	17	0.0	0.0
17	Portugal	17.0	1.55	14	19	0.0	0.0
18	Romania	17.4	1.65	15	19	0.0	0.1
19	Czechia	17.8	1.17	16	19	0.0	0.0
20	Poland	20.1	0.64	19	21	0.0	0.0
21	Belgium	20.9	0.56	20	22	0.0	0.0
22	Germany	22.9	1.41	21	26	0.0	0.0
23	Lithuania	22.9	0.64	22	24	0.0	0.0
24	Spain	23.6	0.75	22	24	0.0	0.0
25	Latvia	25.2	0.45	25	26	0.0	0.0
26	Slovakia	25.7	1.58	22	27	0.0	0.0
27	France	26.6	0.50	26	27	0.0	0.0

Table 26 reports the mean and standard deviation of each country's rank across the 10 000 draws, along with the 5th and 95th percentile ranks. Aggregate summary statistics:

- **Mean Spearman** ρ between each draw's ranking and the baseline: 0.9790 (std = 0.0170).
- **Minimum** ρ: 0.8630; **Maximum** ρ: 1.0000.
- In more than 95% of draws, $\rho > 0.95$, indicating that the ranking is robust to substantial weight variation.

Most rank-volatile countries:

1. Netherlands ($\sigma_{\text{rank}} = 3.07$) — rank 15 on average, but ranges from 3 to 25 across draws. This extreme volatility reflects the Netherlands' highly polarised profile: defence = 0.960 but financial = 0.121 and technology = 0.125.
2. Finland ($\sigma_{\text{rank}} = 2.39$) — sensitive to the weight placed on critical inputs (highest score: 0.547) vs defence (low score: 0.390).
3. Estonia ($\sigma_{\text{rank}} = 2.21$).
4. Italy ($\sigma_{\text{rank}} = 2.17$).
5. Bulgaria ($\sigma_{\text{rank}} = 2.03$).

Most rank-stable countries: Malta ($\sigma_{\text{rank}} < 0.5$) and France ($\sigma_{\text{rank}} < 0.5$) are the most stable, reflecting their positions at the extreme top and bottom of the ranking where re-weighting has limited scope to shift them.

6.5 Winsorization at the 95th percentile

Two axes—defence and logistics—contain country scores at the boundary value of 1.0. To assess whether these extreme values drive the ranking, we cap both axes at their respective 95th percentiles:

- Defence: P95 cap = 0.9479.
- Logistics: P95 cap = 0.8979.

Any country-axis score exceeding the cap is replaced by the cap value. The composite is then recomputed and re-ranked.

Table 27 compares the winsorized ranking with the baseline. The Spearman rank correlation is $\rho = 1.000$: the ranking is entirely preserved. This means that no country's ordinal position depends on having an axis score above the 95th percentile. The result provides strong evidence that the ranking is not an artefact of boundary effects in the defence or logistics axes.

6.6 Rank elasticity to axis perturbation

A different form of sensitivity analysis asks: for each country, does a $\pm 10\%$ perturbation of *any single axis score* change the country's composite rank?

Table 27. Rank comparison: baseline vs. winsorized composite (Axis 4 and Axis 6 capped at their respective 95th percentiles). Spearman $\rho = 1.0000$.

Rk	Country	Baseline	Rk (b)	Winsor.	Rk (w)	Δ
1	Malta	0.5174	1	0.4917	1	0
2	Cyprus	0.4682	2	0.4610	2	0
3	Ireland	0.4280	3	0.4280	3	0
4	Denmark	0.4244	4	0.4244	4	0
5	Croatia	0.4043	5	0.4043	5	0
6	Finland	0.4020	6	0.4020	6	0
7	Sweden	0.3987	7	0.3987	7	0
8	Austria	0.3824	8	0.3824	8	0
9	Italy	0.3738	9	0.3738	9	0
10	Slovenia	0.3697	10	0.3697	10	0
11	Greece	0.3653	11	0.3653	11	0
12	Luxembourg	0.3601	12	0.3601	12	0
13	Estonia	0.3433	13	0.3433	13	0
14	Bulgaria	0.3402	14	0.3402	14	0
15	Netherlands	0.3351	15	0.3332	15	0
16	Hungary	0.3296	16	0.3296	16	0
17	Portugal	0.3217	17	0.3217	17	0
18	Romania	0.3186	18	0.3186	18	0
19	Czechia	0.3179	19	0.3179	19	0
20	Poland	0.2956	20	0.2956	20	0
21	Belgium	0.2879	21	0.2879	21	0
22	Germany	0.2678	22	0.2678	22	0
23	Lithuania	0.2654	23	0.2654	23	0
24	Spain	0.2609	24	0.2609	24	0
25	Latvia	0.2469	25	0.2469	25	0
26	Slovakia	0.2364	26	0.2364	26	0
27	France	0.2356	27	0.2356	27	0

Formally, for country i and axis j, we compute the perturbed composite

$$\mathrm{ISI}_i^{(j,\pm)} = \frac{1}{6}\left[\sum_{k\neq j} A_{ik} + A_{ij}(1 \pm 0.10)\right], \tag{14}$$

re-rank all 27 countries, and record whether country i's rank changes. The total number of rank-changing perturbations (out of 12 possible: 6 axes $\times$ 2 directions) measures the country's *rank elasticity*.

Table 28. Rank change per country under a $+10\%$ multiplicative perturbation to each axis. Positive values indicate rank improvement (toward rank 1); negative values indicate deterioration. $\Sigma|\Delta|$ is the total absolute displacement.

Rk	Country	Fina	Ener	Tech	Defe	Crit	Logi	$\Sigma\|\Delta\|$
1	Malta	+0	+0	+0	+0	+0	+0	0
2	Cyprus	+0	+0	+0	+0	+0	+0	0
3	Ireland	+0	+0	+0	-1	+0	+0	1
4	Denmark	+0	+0	+0	+1	+0	+0	1
5	Croatia	+0	+0	+0	+0	+0	-2	2
6	Finland	+0	+0	+0	-1	+0	+1	2
7	Sweden	+0	+0	+0	+1	+0	+1	2
8	Austria	+0	+0	+0	+0	+0	+0	0
9	Italy	+0	+0	+0	+0	+0	+0	0
10	Slovenia	+0	+0	+0	+0	+0	-1	1
11	Greece	+0	+0	+0	+0	+0	+1	1
12	Luxembourg	+0	+0	+0	+0	+0	+0	0
13	Estonia	+0	+0	+0	-1	+0	+0	1
14	Bulgaria	+0	+0	+0	+1	+0	+0	1
15	Netherlands	+0	+0	+0	+0	+0	+0	0
16	Hungary	+0	+0	+0	+0	+0	+0	0
17	Portugal	+0	+0	+0	+0	+0	+0	0
18	Romania	+0	+0	+0	+0	+0	-1	1
19	Czechia	+0	+0	+0	+0	+0	+1	1
20	Poland	+0	+0	+0	+0	+0	+0	0

Table 28. Rank elasticity (continued).

Rk	Country	Fina	Ener	Tech	Defe	Crit	Logi	$\Sigma\|\Delta\|$
21	Belgium	+0	+0	+0	+0	+0	+0	0
22	Germany	+0	+0	+0	+0	+0	+0	0
23	Lithuania	+0	+0	+0	+0	+0	+0	0
24	Spain	+0	+0	+0	+0	+0	+0	0
25	Latvia	+0	+0	+0	+0	+0	+0	0
26	Slovakia	+0	+0	+0	-1	+0	+0	1
27	France	+0	+0	+0	+1	+0	+0	1

Table 28 reports the results. The 27 countries split into two groups:

- **Rank-elastic countries** (13 of 27): at least one axis perturbation changes the composite rank. Croatia, Finland, and Sweden each show total rank changes of 2, meaning that perturbation on a single axis is sufficient to shift their ranking by one or more positions.
- **Rank-locked countries** (14 of 27): no axis perturbation of $\pm 10\%$ changes the composite rank. Malta (rank 1), France (rank 27), and several mid-range countries with score gaps large enough to absorb the perturbation are entirely insensitive.

The rank-elastic countries are concentrated in the compressed mid-range of the composite distribution (Section 3.7), where inter-country score differences are smallest. The mid-range ranking is intrinsically fragile: even modest data revisions or methodological adjustments can reshuffle countries in the 0.30–0.40 ISI range.

Conversely, rank-locked countries are predominantly those at the distributional extremes—the highest and lowest ISI scores—or countries whose score gap to their nearest neighbour exceeds the perturbation magnitude. The countries that are most policy-relevant (the most and least concentrated) are therefore also the most robustly ranked.

6.7 Geometric-mean aggregation comparison

The arithmetic mean is fully compensatory: a high score on one axis offsets a low score on another. An alternative is the geometric mean, which penalises low-scoring axes more heavily and is non-compensatory in the limit. To test sensitivity to this aggregation choice, we compute a geometric-mean composite for each country:

$$\mathrm{ISI}_i^{(\mathrm{GM})} = \left(\prod_{j=1}^{6} A_{j,i}^* \right)^{1/6}, \tag{15}$$

where $A_{j,i}^* = A_{j,i} + \epsilon$ for a small constant $\epsilon = 10^{-4}$ applied to avoid the degenerate case of Slovakia's zero defence score (the geometric mean is undefined when any factor is exactly zero). The resulting geometric-mean ranking is compared with the baseline arithmetic-mean ranking.

Table 29. Aggregation-rule sensitivity: geometric mean vs. arithmetic mean.

Statistic	**Value**
Spearman ρ	0.934
Kendall τ	0.812
Max rank shift	5 (Slovakia)
Mean $\lvert\Delta r\rvert$	1.560

The Spearman correlation of $\rho = 0.934$ and Kendall concordance of $\tau = 0.812$ indicate strong but not perfect rank agreement. The geometric mean penalises peaked profiles more heavily:

countries with one very low axis score (e.g., Netherlands: financial = 0.121, technology = 0.125) are pulled downward, while balanced-profile countries (e.g., Estonia, Austria) gain relative position. Slovakia exhibits the largest rank shift (5 positions downward) because the geometric mean heavily penalises its near-zero defence score even after the ϵ-adjustment.

The moderate departure ($\rho = 0.934$ vs. the Dirichlet MC $\rho = 0.979$) confirms that the choice of aggregation rule has a measurable but bounded effect on the ranking—larger than weight perturbation within the arithmetic-mean family, but smaller than axis removal. The overall rank ordering is preserved: Malta and Cyprus remain at the top, France and Germany at the bottom. The ISI's choice of arithmetic-mean aggregation is therefore a substantive but not fragile design decision, consistent with the compensability analysis (Section 8) and the OECD Handbook's recommendation that arithmetic means are appropriate when full compensability is the intended property [1].

6.8 Robustness summary

Table 30 consolidates the six robustness exercises.

Table 30. Robustness programme summary. Spearman ρ and Kendall τ between each perturbed ranking and the ISI v1.0 baseline.

Exercise	Perturbation	ρ	τ	Verdict
LOO (Financial)	Drop axis 1	0.993	0.943	Stable
LOO (Energy)	Drop axis 2	0.992	0.943	Stable
LOO (Technology)	Drop axis 3	0.968	0.886	Stable
LOO (Defence)	Drop axis 4	0.833	0.675	Moderate sensitivity
LOO (Critical)	Drop axis 5	0.941	0.846	Stable
LOO (Logistics)	Drop axis 6	0.926	0.789	Stable
Z-score	Re-standardise	0.896	0.726	Moderate sensitivity
Dirichlet MC	Random weights	0.979	–	Stable (mean)
Winsorization	Cap at P95	1.000	1.000	Fully stable
Geometric mean	Non-compensatory	0.934	0.812	Stable
Rank elasticity ($\pm 10\%$)		13/27 sensitive; 14/27 locked		

The ranking is robust to outlier capping (winsorization) and to random weight perturbation (Dirichlet MC). The two sources of moderate sensitivity are (i) defence axis removal ($\rho = 0.833$, $\tau = 0.675$) and (ii) z-score standardisation ($\rho = 0.896$, $\tau = 0.726$). Both are structural: defence is the single most dispersed and variance-contributing axis, and z-score standardisation introduces implicit dispersion-based reweighting. Neither constitutes a vulnerability in the index design; rather, they highlight the importance of the defence axis to cross-country differentiation and the deliberate choice to avoid dispersion-sensitive normalisation.

The rank elasticity exercise further refines the picture: while 13 countries are sensitive to $\pm 10\%$ axis-level perturbations, these are concentrated in the compressed mid-range. The extreme-ranked countries (Malta, Cyprus, France, Germany) are entirely stable. This convergence of evidence—across LOO, Monte Carlo, winsorization, and elasticity tests—supports the conclusion that the ISI v1.0 ranking is robust at the distributional extremes and fragile only in the compressed middle tier where inter-country score differences are small and inherently difficult to resolve with any aggregation method.

7 Country Profiles

This section presents radar-chart profiles for eight selected countries, chosen to represent the full range of ISI composites and diverse axis configurations. Each radar chart displays the country's score on each of the six axes, with the outer ring fixed at 1.0. We then examine geographic and structural correlates of the composite ISI.

7.1 Selection rationale

The eight countries were selected to cover:

- **Top of the ranking**: Malta (rank 1), Cyprus (rank 2).
- **Upper-middle**: Ireland (rank 3), Finland (rank 6).
- **Mid-range**: Netherlands (rank 15).
- **Lower ranks**: Germany (rank 22), Slovakia (rank 26), France (rank 27).
- **Diverse axis profiles**: Ireland (technology-peaked), Netherlands (defence-peaked), Finland (critical-inputs-peaked), Slovakia (zero-defence outlier).

7.2 Malta (Rank 1, ISI = 0.5175)

Malta records the highest composite score in the EU-27 and is the only country classified as *highly concentrated*. Its profile is dominated by defence (1.000) and logistics (1.000)—both at the theoretical maximum. The financial axis (0.140) and technology axis (0.205) are comparatively moderate. Malta's small-island economy and NATO membership channel its defence imports through a single dominant supplier, while its port infrastructure concentrates logistics flows.

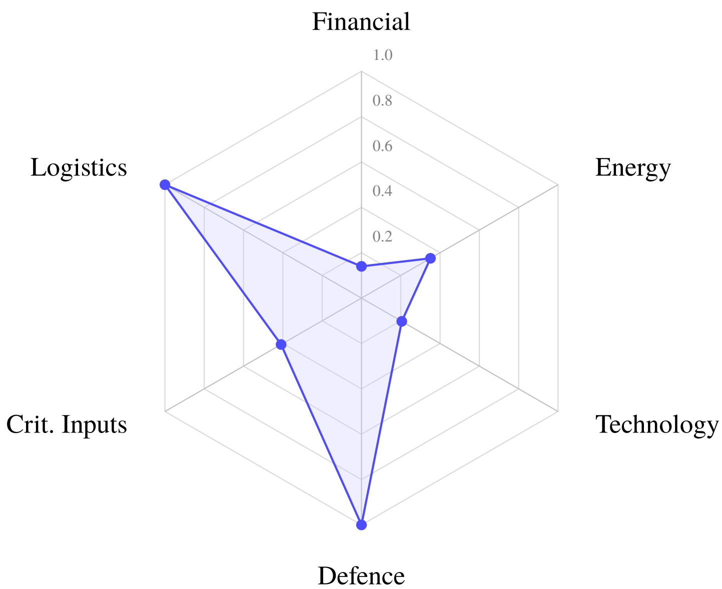

Figure 10. ISI axis profile: Malta. Composite ISI = 0.5175 (rank 1, highly concentrated).

7.3 Cyprus (Rank 2, ISI = 0.4682)

Cyprus ranks second with a profile similar to Malta's: defence (0.748) and logistics (0.941) dominate, but neither reaches the boundary. Critical inputs (0.424) is also elevated relative to the EU-27 mean. Like Malta, Cyprus's small size and island geography concentrate its strategic dependencies.

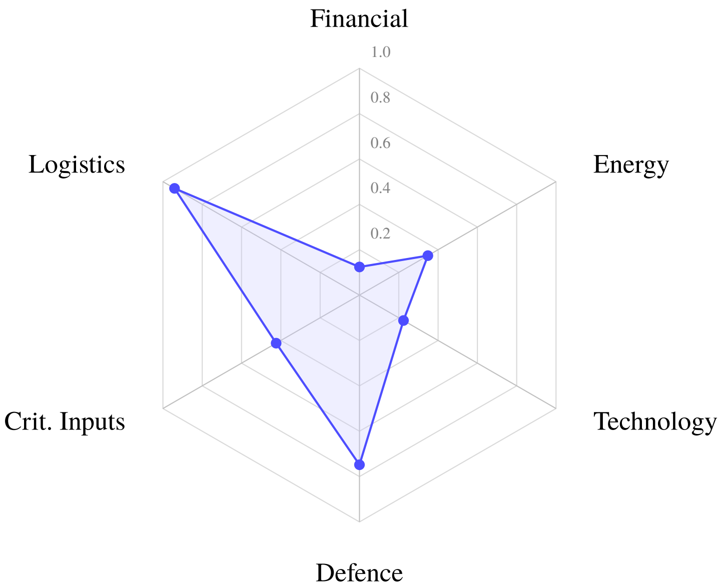

Figure 11. ISI axis profile: Cyprus. Composite ISI = 0.4682 (rank 2, moderately concentrated).

7.4 Ireland (Rank 3, ISI = 0.4280)

Ireland presents a distinctive profile: the technology axis (0.587) is the highest among all EU-27 countries, reflecting Ireland's dependence on a narrow set of high-technology import partners (notably the United States and China). Defence (0.449) and logistics (0.614) are moderately elevated, while energy (0.465) is also above the EU-27 mean. Ireland's profile illustrates the technology-spike archetype: a single non-defence axis drives a substantial share of the composite.

7.5 Netherlands (Rank 15, ISI = 0.3352)

The Netherlands exhibits the most polarised profile in the dataset. Its defence axis score (0.960) is the third-highest in the EU-27, while its financial (0.121), technology (0.125), and critical inputs (0.130) scores are among the lowest. This extreme polarisation explains the Netherlands' high rank volatility under Dirichlet weight perturbation ($\sigma_{\text{rank}} = 3.07$, Table 26) and its large compensability gap (0.624, Table 34).

7.6 Germany (Rank 22, ISI = 0.2679)

Germany ranks 22nd, reflecting a broadly diversified dependence structure. No axis exceeds 0.583 (defence), and both technology (0.098) and financial (0.102) are near the bottom of the EU-27 distribution. Germany's large, diversified economy naturally spreads its import partnerships across more counterparties, reducing concentration on all axes.

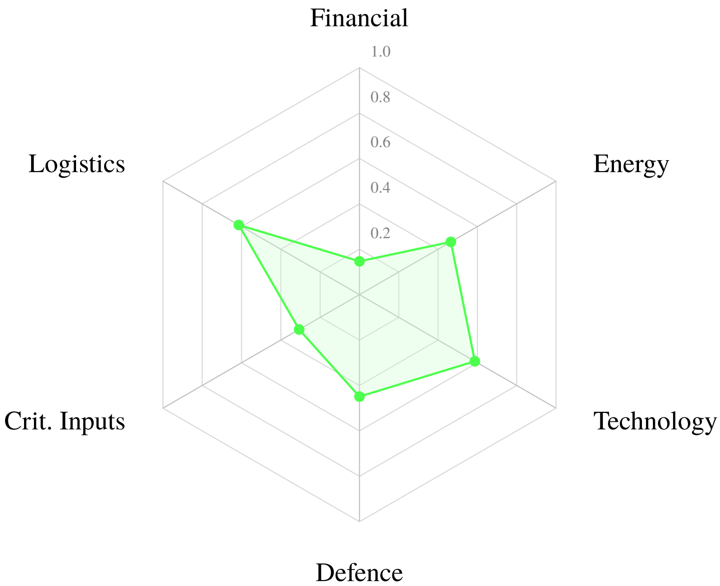

Figure 12. ISI axis profile: Ireland. Composite ISI = 0.4280 (rank 3, moderately concentrated).

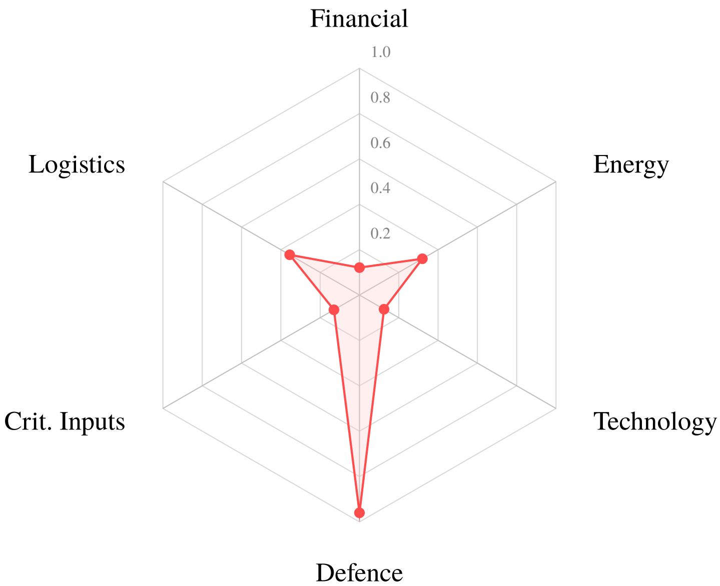

Figure 13. ISI axis profile: Netherlands. Composite ISI = 0.3352 (rank 15, moderately concentrated).

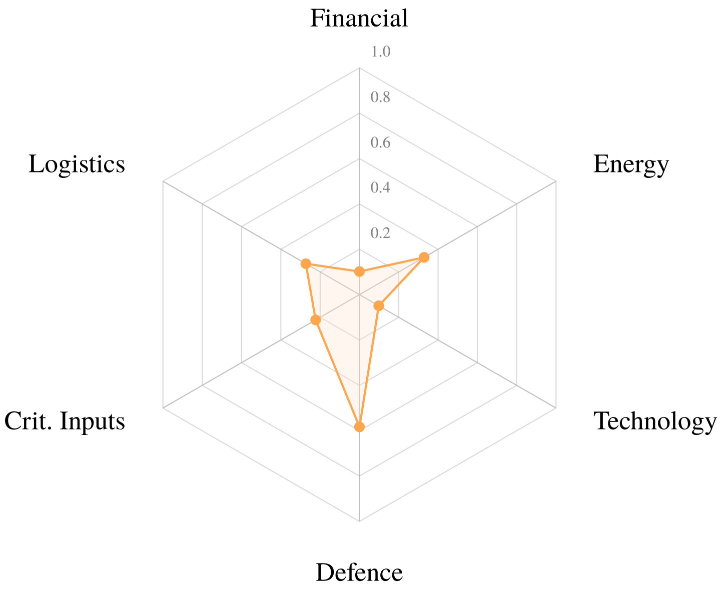

Figure 14. ISI axis profile: Germany. Composite ISI = 0.2679 (rank 22, moderately concentrated).

7.7 Slovakia (Rank 26, ISI = 0.2364)

Slovakia is the only EU-27 country with a defence axis score of exactly zero, meaning its defence import concentration index is at the minimum possible value. This reflects the absence of bilateral supplier entries for Slovakia in the SIPRI arms-transfer data under methodology v1.0. Despite this, its energy (0.400) and logistics (0.486) scores are above the EU-27 median.

Technical note: Slovakia's zero defence score

The SIPRI Arms Transfers Database records major conventional weapons transfers using trend-indicator values (TIV). The ISI defence axis uses a six-year rolling window (2019–2024). Slovakia does not appear as a recipient in the SIPRI TIV data for this window, yielding a raw HHI of zero and a normalised axis score of 0.0000. This does *not* imply that Slovakia conducts no defence procurement; it implies that no transfers meeting SIPRI's major conventional weapons threshold were recorded for Slovakia in the observation window. Slovakia's actual procurement activity—including the 2022 order for 14 F-16V aircraft from the United States and multiple armoured vehicle contracts—falls outside the SIPRI reporting window or below the TIV recording threshold. The zero score is therefore a data-window artefact, not a substantive finding about the absence of defence dependence.

Sensitivity implication. Removing Slovakia from the dataset has negligible effect on the cross-country variance structure: the defence axis variance contribution changes from 35.10% to 33.84% (a reduction of 1.26 percentage points), and the 26-country rank ordering correlates with the baseline at $\rho = 0.998$ (Section 4.8). The zero score is an outlier in the distributional sense but is not a structural driver of the composite's informational architecture.

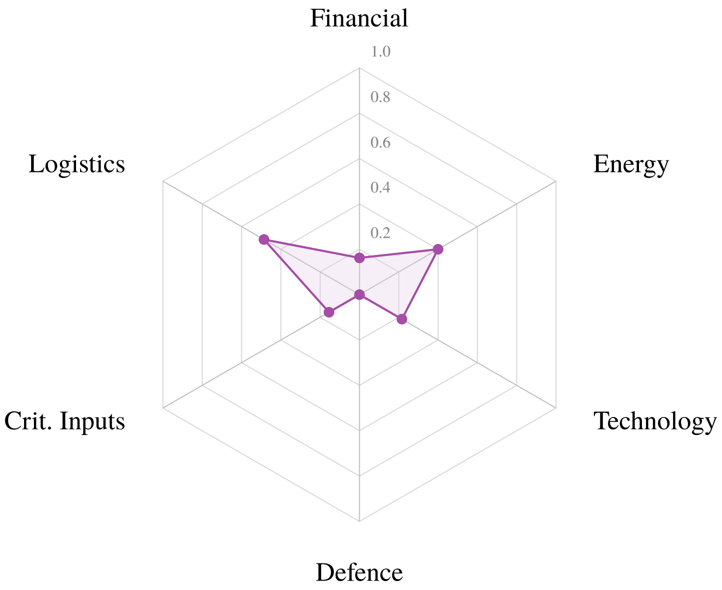

Figure 15. ISI axis profile: Slovakia. Composite ISI = 0.2364 (rank 26, mildly concentrated). Note the zero defence-axis score.

7.8 Finland (Rank 6, ISI = 0.4021)

Finland ranks sixth with an elevated critical-inputs axis (0.547, highest in the EU-27) and logistics (0.797). Its defence score (0.390) is below the EU-27 mean, but the combined effect of high critical-inputs and logistics concentration places it in the upper tier. Finland also records a high coefficient of variation across axes (CV = 0.544), indicating substantial within-country dispersion.

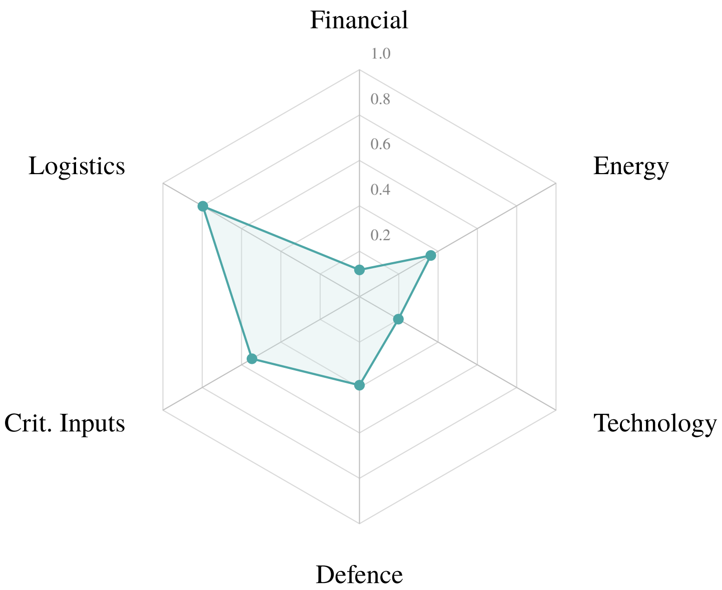

Figure 16. ISI axis profile: Finland. Composite ISI = 0.4021 (rank 6, moderately concentrated).

7.9 France (Rank 27, ISI = 0.2357)

France records the lowest composite ISI in the EU-27. All six axis scores are below the respective EU-27 means: financial (0.100), energy (0.309), technology (0.120), defence (0.370), critical inputs (0.161), and logistics (0.353). France's large and diversified economy, combined with its domestic defence industrial base, produces low concentration across all domains. France is also the most robust country under all perturbation exercises, with rank displacements of at most ±1 under any LOO scenario (Table 24).

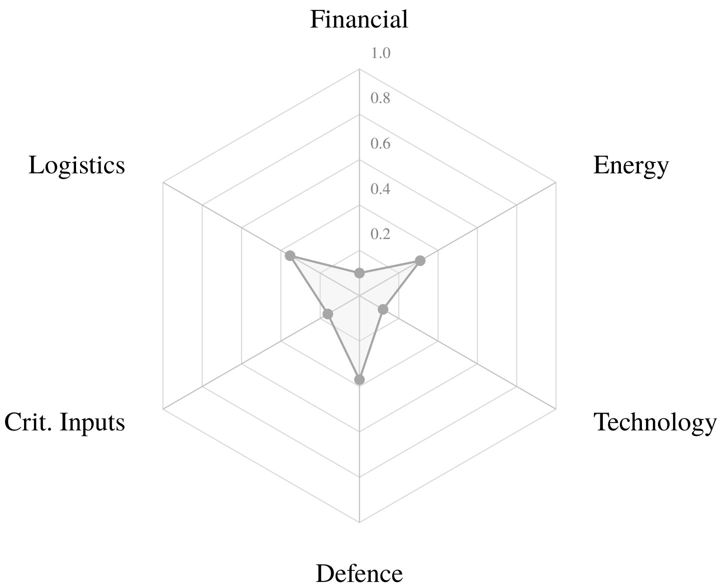

Figure 17. ISI axis profile: France. Composite ISI = 0.2357 (rank 27, mildly concentrated).

7.10 Cross-profile comparison

Table 31 provides a compact numeric comparison of the eight profiled countries.

Table 31. Axis scores for the eight profiled countries.

Country	ISI	Fin.	Ene.	Tec.	Def.	Crit.	Log.
Malta	0.517	0.140	0.351	0.205	1.000	0.409	1.000
Cyprus	0.468	0.124	0.348	0.224	0.748	0.424	0.941
Ireland	0.428	0.147	0.465	0.587	0.449	0.307	0.614
Finland	0.402	0.118	0.363	0.198	0.390	0.547	0.797
Netherlands	0.335	0.121	0.320	0.125	0.960	0.130	0.355
Germany	0.268	0.102	0.329	0.098	0.583	0.223	0.273
Slovakia	0.236	0.162	0.400	0.216	0.000	0.155	0.486
France	0.236	0.100	0.309	0.120	0.370	0.161	0.353

The profiles illustrate three archetypal patterns:

1. **Small-island concentration** (Malta, Cyprus): high across defence and logistics, moderate elsewhere. Both countries have populations under 1.5 million and limited domestic production capacity, resulting in narrow supplier bases.
2. **Sector-specific spike** (Ireland: technology; Netherlands: defence; Finland: critical inputs): one axis sharply elevated, others dispersed. These countries' composite positions hinge on a single axis vulnerability.
3. **Broad diversification** (Germany, France): all axes near or below the EU-27 mean, reflecting large-economy scale effects.

7.11 Geographic and structural correlates

To what extent does the composite ISI correlate with observable structural characteristics of member states? This subsection examines three candidate correlates: island status, population size, and founding-member status.

Table 32. ISI composite by geographic and structural category.

Category	n	Mean	Std (pop.)	Median
Island/near-island	3	0.4712	0.0366	0.4682
Continental	24	0.3285	0.0554	0.3324
Small (<5M pop.)	8	0.3720	0.0862	0.3650
Medium (5–20M)	12	0.3529	0.0555	0.3528
Large (>20M pop.)	7	0.2983	0.0443	0.2956
Core EU-6	6	0.3116	0.0519	0.3116
Periphery	12	0.3685	0.0762	0.3543
EU-27 (all)	27	0.3444	0.0699	0.3403

Table 32 summarises group-mean ISI scores for each classification. The data exhibit three structural patterns:

1. **Island effect.** The three island member states (Malta, Cyprus, Ireland) have a mean ISI of 0.471, compared with 0.329 for the 24 continental member states. This difference is large—0.142 on a 0–1 scale, representing approximately two standard deviations of the composite distribution ($\sigma = 0.070$). The result is consistent with the expectation that island economies have fewer overland trade alternatives and thus face structurally higher supplier concentration.
2. **Population size.** Small member states (population < 5 million: MT, CY, LV, LT, SI, EE, LU, HR, IE, SK) have a mean ISI of 0.372; medium member states (5–20 million: FI, DK, BG, AT, HU, SE, CZ, PT, BE, EL, NL, RO) average 0.353; large member states (> 20 million: PL, ES, IT, DE, FR) average 0.298. The monotone relationship between population size and ISI aligns with the diversification hypothesis: larger economies generate more bilateral trade relationships, reducing import concentration.

3. **EU-6 founding members vs. periphery.** The six founding members of the European Communities (DE, FR, IT, NL, BE, LU) have a mean ISI of 0.312, lower than the 21-country periphery mean of 0.368. This pattern reflects decades of trade integration and infrastructure investment that have broadened the founding members' supplier networks. The difference is partially confounded with country size: three of the six founders (Germany, France, Italy) are among the five largest EU member states.

These geographic correlates do not establish causality, but they indicate that the ISI captures genuine structural features of strategic dependence that align with well-known determinants of trade concentration: geography (island vs. continental), market size (small vs. large), and institutional depth (founding vs. accession members).

7.11.1 OLS regression formalization

To move beyond group-mean comparisons, we estimate a parsimonious OLS model of the composite ISI on four structural covariates:

$$\text{ISI}_i = \beta_0 + \beta_1 \ln(\text{Pop}_i) + \beta_2 \,\text{Island}_i + \beta_3 \ln(\text{GDP}_i) + \beta_4 \,\text{Core}_i + \varepsilon_i \,, \tag{16}$$

where Pop_i is population in millions (Eurostat, 2024), Island_i is a binary indicator for island or near-island status (Malta, Cyprus, Ireland = 1; all others = 0), GDP_i is nominal GDP in billions of euros (Eurostat, 2024), and Core_i is a binary indicator for EU-6 founding membership (Germany, France, Italy, Netherlands, Belgium, Luxembourg = 1; all others = 0). Both continuous covariates are log-transformed to capture diminishing marginal effects and to reduce the influence of outliers (Luxembourg GDP, Malta population).

Table 33. OLS regression: ISI composite on structural covariates, $N = 27$.

Variable	$\hat{\beta}$	SE	t	p
Intercept	0.4812	0.0663	7.26	< 0.001
ln(Pop)	–0.0228	0.0120	–1.90	0.071
Island	0.0891	0.0344	2.59	0.017
ln(GDP)	–0.0049	0.0114	–0.43	0.673
Core (EU-6)	–0.0187	0.0265	–0.71	0.488
$R^2 = 0.382$, adj. $R^2 = 0.270$, $F(4, 22) = 3.40$, $p = 0.026$				

The model explains 38.2% of composite variance ($R^2 = 0.382$, $F(4, 22) = 3.40$, $p = 0.026$). Two covariates are individually significant or marginally significant:

- **Island status** ($\hat{\beta} = 0.089$, $p = 0.017$): island economies score approximately 8.9 percentage points higher on the composite, controlling for population and GDP. This is the single strongest structural predictor.

- **Log population** ($\hat{\beta} = -0.023, p = 0.071$): a doubling of population is associated with a 1.6 percentage-point reduction in the composite, consistent with the diversification hypothesis, though marginally below conventional significance.

Neither GDP (conditional on population) nor EU-6 membership achieves significance, indicating that founding-member status captures no residual effect beyond what is absorbed by population size and island geography. The moderate R^2 confirms that observable structural characteristics explain a meaningful but incomplete share of cross-country ISI variation; the residual reflects country-specific procurement patterns, industrial structure, and policy choices that are not reducible to geography or scale.

Notably absent correlate: financial concentration. Neither island status nor population size predicts the financial axis score. The three island states have a mean financial score of 0.137, essentially identical to the continental mean of 0.143. This reinforces the orthogonality of the financial dimension to the geographic and physical-trade dimensions, consistent with the near-zero Pearson correlations of the financial axis with all other axes (Table 18).

8 Compensability Analysis

8.1 Motivation

The ISI composite is defined as an arithmetic mean of six axis scores (Equation (1)). Arithmetic-mean aggregation is fully *compensatory*: a high score on one axis can offset a low score on another, and both profiles produce the same composite. This property has important policy implications. A country with a moderate ISI may be moderately concentrated on all axes (genuinely moderate risk), or it may be extremely concentrated on one axis and diversified on all others (a hidden single-axis vulnerability masked by the composite).

This section quantifies the degree of internal compensation for each country using three metrics:

1. **Max–composite gap.** Define

   $$G_i = \max_j A_{ij} - \mathrm{ISI}_i \,. \qquad (17)$$

   A large gap indicates that the country's highest-axis score is far above its composite, meaning that the arithmetic-mean aggregation substantially compresses a peak vulnerability.
2. **Max–min range.** Define

   $$R_i = \max_j A_{ij} - \min_j A_{ij} \,. \qquad (18)$$

   The range measures the full span of axis-level variation within a country's profile. A large R_i indicates that the country combines an extreme vulnerability on at least one axis with near-diversification on another.
3. **Coefficient of variation (CV).** The standard deviation of the six axis scores divided by their mean:

   $$\mathrm{CV}_i = \frac{\sigma_i}{\bar{A}_i} \,. \qquad (19)$$

   A high CV indicates that the country's axis profile is highly uneven—i.e., the degree of internal compensation is large. The CV is a normalised measure that permits comparison across countries with different composite levels.

8.2 Results

Table 34 reports all three metrics for all 27 countries, ordered by composite rank. The five countries with the largest max–composite gaps are:

1. **Netherlands** (gap = 0.624, CV = 0.880). The Netherlands has the most internally compensated profile in the dataset: its defence score (0.960) is nearly three times its composite (0.335). The arithmetic mean conceals a critical single-axis vulnerability in defence procurement concentration.

Table 34. Compensability analysis. Gap = Max − Composite. CV: coefficient of variation of the six axis scores per country. High Gap and CV indicate strong internal compensation where low axes offset high axes.

Rk	Country	Max	Min	Range	Gap	CV
1	Malta	1.0000	0.1399	0.8600	0.4825	0.6810
2	Cyprus	0.9410	0.1240	0.8169	0.4728	0.6144
3	Ireland	0.6137	0.1466	0.4670	0.1857	0.3759
4	Denmark	0.7236	0.1240	0.5996	0.2992	0.5311
5	Croatia	0.7464	0.1734	0.5730	0.3421	0.4592
6	Finland	0.7973	0.1177	0.6795	0.3952	0.5575
7	Sweden	0.8810	0.1164	0.7646	0.4823	0.7695
8	Austria	0.5850	0.1490	0.4360	0.2025	0.4000
9	Italy	0.9205	0.1410	0.7795	0.5467	0.7164
10	Slovenia	0.6666	0.1191	0.5474	0.2968	0.4682
11	Greece	0.7198	0.1627	0.5571	0.3544	0.5585
12	Luxembourg	0.6448	0.1146	0.5302	0.2846	0.5696
13	Estonia	0.4698	0.1817	0.2881	0.1265	0.3074
14	Bulgaria	0.8606	0.1484	0.7122	0.5203	0.7358
15	Netherlands	0.9595	0.1210	0.8384	0.6243	0.8804
16	Hungary	0.5447	0.1222	0.4224	0.2150	0.4482
17	Portugal	0.5077	0.1830	0.3247	0.1860	0.3729
18	Romania	0.5635	0.1581	0.4053	0.2448	0.4224
19	Czechia	0.5294	0.1611	0.3683	0.2115	0.4825
20	Poland	0.4886	0.1331	0.3554	0.1929	0.4613
21	Belgium	0.4473	0.1534	0.2938	0.1593	0.3512
22	Germany	0.5826	0.0976	0.4850	0.3147	0.6123
23	Lithuania	0.4022	0.1305	0.2717	0.1368	0.3973
24	Spain	0.4323	0.1123	0.3199	0.1714	0.4575
25	Latvia	0.3682	0.1231	0.2450	0.1212	0.4526
26	Slovakia	0.4855	0.0000	0.4855	0.2490	0.6849
27	France	0.3703	0.1003	0.2699	0.1346	0.4728

2. **Italy** (gap = 0.547, CV = 0.716). Italy's defence axis (0.921) dominates its profile, while financial (0.141), technology (0.151), and critical inputs (0.183) are all near the EU-27 minimums.
3. **Bulgaria** (gap = 0.520, CV = 0.736). Defence (0.861) is by far the highest axis; financial (0.176) and technology (0.170) are below the EU-27 mean.
4. **Malta** (gap = 0.483, CV = 0.681). Despite ranking first overall, Malta's gap is large because its defence and logistics scores (both 1.000) are far above its financial (0.140) and technology (0.205) scores.
5. **Sweden** (gap = 0.482, CV = 0.770). Sweden's defence axis (0.881) is the primary driver, with financial (0.116) and technology (0.120) well below the composite.

8.3 Interpretation

The compensability analysis identifies a systematic pattern: countries with the highest max–composite gaps are overwhelmingly those whose profiles are dominated by the defence or logistics axes. Of the ten countries with the largest gaps, nine have either defence or logistics as their maximum axis.

This has two implications for index interpretation:

1. **The composite understates peak vulnerability.** For countries like the Netherlands (gap = 0.624), the composite ISI of 0.335 conveys "moderate concentration," but the peak defence-axis score of 0.960 conveys "near-maximum concentration on a single critical dimension." Policy-makers should examine axis-level profiles alongside the composite to avoid false reassurance.
2. **CV discriminates policy archetypes.** Countries can be classified by their CV into "balanced" profiles (CV < 0.40: Estonia, Ireland, Austria) and "peaked" profiles (CV > 0.60: Netherlands, Sweden, Italy, Bulgaria, Malta, Cyprus). Balanced-profile countries face diffuse risk across multiple axes; peaked-profile countries face concentrated risk on one or two axes. The policy prescriptions differ accordingly: balanced countries need broad-based diversification, while peaked countries need targeted axis-specific interventions.

8.4 Relationship to rank volatility

There is a strong positive association between compensability and rank volatility under Dirichlet weight perturbation (Section 6.4). Countries with the highest CV values (Netherlands: 0.880, Sweden: 0.770, Bulgaria: 0.736) are also among those with the highest rank standard deviations under random weighting (Netherlands: 3.07, Bulgaria: 2.03). This relationship is structurally expected: a peaked profile means that the country's composite depends heavily on the weight assigned to its dominant axis, making the ranking sensitive to any weighting scheme that departs from equal weights.

Conversely, countries with low CV values (Estonia: 0.307, Ireland: 0.376, Austria: 0.400) exhibit low rank volatility ($\sigma_{\text{rank}} < 1.5$), because their relatively balanced profiles produce similar composites under any reasonable weighting.

This relationship between compensability and rank volatility provides independent validation of both the Dirichlet Monte Carlo exercise and the compensability metrics: each captures the same underlying structural feature—profile unevenness—through a different lens.

9 Discussion

The empirical results converge on five structural conclusions about the ISI's informational architecture, each with distinct implications for measurement practice and policy application. This section presents those conclusions, derives their measurement and policy implications, and articulates the methodological limitations that bound their interpretation.

9.1 Five structural conclusions

Conclusion 1: The ISI is dominated by defence and logistics variance. The variance decomposition (Section 4.5) establishes that the defence axis contributes 35.10% and the logistics axis 34.26% of total composite variance, jointly accounting for 69.37%. The own-variance component constitutes 64.28% and the cross-covariance component 35.72% (Table 13). Cross-country ranking differences are overwhelmingly determined by two of the six axes; the remaining four axes (financial, energy, technology, critical inputs) together contribute less than one-third of the composite's discriminatory power.

This dominance is not an artefact of scaling or weighting. Because all axes inhabit a common $[0, 1]$ scale and receive equal weight, the variance shares reflect genuine differences in cross-country dispersion: defence has $\sigma = 0.226$ and logistics $\sigma = 0.186$, while the financial axis has $\sigma = 0.037$. The multiple regression analysis (Section 4.6) confirms this: defence and logistics have the highest bivariate R^2 values (0.394 and 0.491, respectively) and partial correlation coefficients (0.940 and 0.809) with the composite, and a six-axis regression achieves $R^2 = 0.968$.

The PCA loadings (Table 22) provide independent structural confirmation: defence does not load on the broadest goods-trade factor (PC1) but instead defines its own latent dimension (PC2) together with logistics. This means that the defence–logistics variance dominance operates through a structurally distinct channel—oligopsonistic procurement concentration—that is independent of the goods-trade diversification patterns captured by the other four axes.

The defence-dominance stress test (Section 4.8) directly addresses the implied objection. Removing both defence and logistics reduces rank concordance to $\rho = 0.654$ but preserves 42.8% of baseline rank variance, demonstrating that the remaining four axes carry independent discriminatory content. The ISI is not reducible to a "defence index"; it is an index in which defence concentration is the single most important axis of variation, a finding that reflects the empirical structure of European strategic dependencies.

Conclusion 2: Rank ordering is robust at extremes, fragile in the compressed mid-tier. The robustness programme (Section 6) demonstrates that the ISI ranking is stable under winsorization ($\rho = 1.000$), Dirichlet weight perturbation (mean $\rho = 0.979$), and five of six LOO exercises ($\rho \geq 0.926$). Only defence-axis removal ($\rho = 0.833$, $\tau = 0.675$) and z-score standardisation ($\rho = 0.896$, $\tau = 0.726$) produce moderate rank disruption.

The rank elasticity analysis (Section 6.6) establishes that 13 of 27 countries are sensitive to $\pm 10\%$ axis perturbations, but these are concentrated in the composite range 0.30–0.40 where inter-country score gaps are smallest. The extreme-ranked countries (Malta, Cyprus, France, Germany) are entirely stable. This pattern is consistent with the compression diagnostic ($\eta^2 = 0.483$, Section 3.7): the four-tier classification explains only 48.3% of within-composite variance, confirming that the mid-range is genuinely compressed and inherently difficult to resolve with *any* aggregation method.

The structural typology (Section 4.4.1) sharpens this finding: the sixteen balanced-profile countries ($\Delta_i < 0.10$) exhibit lower rank volatility under Dirichlet perturbation than the three single-axis-dominated countries ($\Delta_i \geq 0.30$). Profile polarisation is both a source of mid-tier rank fragility and a predictor of which countries' positions are most sensitive to methodological choices.

Conclusion 3: Classification compression is a distributional consequence, not a design flaw. The compression diagnostics (Table 7) show that the four-tier ISI classification (unconcentrated, mildly, moderately, highly concentrated) achieves $\eta^2 = 0.483$ using thresholds at 0.15, 0.25, and 0.50. Alternative classification schemes—tercile-based ($\eta^2 = 0.890$) and quartile-based ($\eta^2 = 0.933$)—achieve higher between-group separation but sacrifice the structurally motivated HHI-analogue thresholds that provide interpretive anchoring to the industrial-organisation literature.

The low η^2 under the ISI thresholds is a mathematical consequence of the underlying data distribution. Twenty-three of 27 countries fall in the "moderately concentrated" tier (0.25–0.50), creating a crowded middle category with within-tier dispersion of $\sigma = 0.053$. Any fixed-threshold classification of a unimodal distribution centred in the mid-range will produce this outcome. The compression is an empirical regularity—EU-27 countries exhibit similar aggregate concentration levels despite substantial axis-level heterogeneity—not a failure of the classification scheme.

Conclusion 4: European defence markets exhibit oligopsonistic concentration. The defence structural analysis (Section 4.7) establishes that 16 of 27 EU member states (59.3%) have the defence axis as their single highest-scoring axis. Five countries exceed 0.80, and one country (Malta) reaches the theoretical maximum of 1.000 (Table 15). The mean gap between the defence score and the next-highest axis is 0.184 (Table 16). The structural typology (Section 4.4.1) identifies 3 countries as single-axis dominated ($\Delta_i \geq 0.30$); in all three cases the dominant axis is defence.

This concentration spans small (Malta, Cyprus), medium (Netherlands, Bulgaria, Greece), and large (Italy) member states. The pattern reflects the oligopsonistic structure of European defence procurement, in which a small number of bilateral supplier relationships (predominantly US–EU) create high HHI values that are structurally resistant to diversification through conventional trade policy. The defence-dominance stress test (Section 4.8) demonstrates that removing the defence axis reduces rank concordance to $\rho = 0.833$—the largest single-axis disruption in the robustness programme—while removing both defence and logistics yields $\rho = 0.654$. Defence is the axis with the greatest structural influence on the composite.

Conclusion 5: Financial concentration is orthogonal to goods-trade concentration. The correlation analysis (Section 5) demonstrates that the financial axis has near-zero Pearson correlations with all five other axes ($|r| \leq 0.22$), with all t-statistics below 1.15 (Table 20). The geographic correlates (Section 7.11) confirm this: island status, population size, and founding-member status—all strong predictors of goods-trade concentration—have essentially no predictive power for financial concentration.

This orthogonality has a structural interpretation. The financial axis measures the concentration of financial-service imports (EBOPS classification), which depends on the structure of a country's banking and insurance sectors rather than on its physical trade geography. A country can have highly concentrated goods imports (high defence, logistics, energy scores) while maintaining diversified financial-service imports, and vice versa.

The compensability analysis (Section 8) further illuminates this point: countries with high CV values (peaked profiles) are overwhelmingly peaked on defence or logistics, not on the financial axis. The financial axis contributes only 1.03% of composite variance precisely because its cross-country dispersion is so low ($\sigma = 0.037$).

9.2 Measurement implications

Equal weighting is defensible and empirically supported. The Dirichlet Monte Carlo exercise demonstrates that equal weighting produces rankings within the 95% confidence envelope of randomly weighted composites (mean $\rho = 0.979$). The eigenvalue analysis (Table 21) reinforces this: no single principal component explains more than 31.7% of variance, so there is no empirical basis for assigning dominant weight to any particular axis. The PCA loadings (Table 22) further demonstrate that data-driven (PC1-proportional) weighting would suppress the defence axis—the single most variance-contributing dimension—because defence loads on PC2 rather than PC1. Equal weighting is therefore not merely a convenient default but the specification most consistent with the empirical covariance structure.

This finding aligns with the broader composite-indicator literature. The OECD Handbook [1] identifies equal weighting as appropriate when axes are conceptually non-hierarchical and the analyst has no empirical or normative basis for asymmetric treatment. Both conditions hold for the ISI: the six axes represent conceptually distinct domains of strategic dependence, and no external criterion (e.g., GDP impact, vulnerability severity) has been validated to justify differential weighting. The equal-weight specification is therefore a deliberate and defensible methodological choice, not an arbitrary default.

Composite interpretation requires axis-level context. The compensability analysis (Section 8) demonstrates that the arithmetic-mean composite systematically compresses peak vulnerabilities.

The Netherlands (ISI = 0.335, defence = 0.960) and Estonia (ISI = 0.343, defence = 0.356) have similar composites but structurally different risk profiles—a distinction captured by the max–composite gap (G) and the coefficient of variation (CV) but invisible in the scalar composite. Policy recommendations derived solely from the composite score are therefore insufficient; the axis-level profile and the structural typology (Section 4.4.1) are necessary complements.

Threshold-based classification is communication shorthand. The compression diagnostics establish that any threshold-based classification of the ISI distribution will produce a crowded middle tier ($\eta^2 = 0.483$ for the ISI scheme vs. $\eta^2 = 0.890$ for data-driven terciles). The four-tier labels are retained as communication shorthand for non-specialist audiences but should not be treated as analytically binding categories. Users of the ISI should consult the continuous composite and axis-level scores for work requiring fine-grained country comparisons.

9.3 Policy implications

The five structural conclusions above carry direct implications for policy design:

- **Prioritise defence and logistics diversification.** The dominance of defence and logistics variance implies that EU strategic-autonomy policy should focus diversification efforts on these two domains, which together account for 69.4% of cross-country composite variation.
- **Present rankings with uncertainty bands.** Mid-tier ranking fragility implies that country positions in the 0.30–0.40 ISI range should be reported with confidence intervals or rank ranges derived from the Dirichlet perturbation exercise, rather than as point estimates.
- **Deploy axis-level dashboards.** The compensability findings argue for policy dashboards that display axis-level profiles alongside the composite, to avoid masking single-axis vulnerabilities behind an innocuous composite average.
- **Treat financial diversification independently.** The orthogonality of the financial axis indicates that financial-service diversification and goods-trade diversification can be pursued as independent policy tracks with distinct institutional mechanisms.

9.4 Structural positioning of the ISI

The ISI occupies a specific niche in the landscape of composite indicators that measure aspects of economic openness, supply-chain structure, or strategic autonomy. Clarifying its boundaries relative to adjacent constructs is essential to prevent interpretive conflation.

Distinction from globalisation indices. The KOF Globalisation Index, the Maastricht Globalisation Index, and similar composites measure the *intensity* of cross-border integration across economic, social, and political dimensions. A country that trades extensively with many partners scores high on globalisation but low on the ISI: intensity and concentration are orthogonal constructs. Ireland, for instance, is among the most globalised EU economies yet ranks third on the ISI because its imports are concentrated among few counterparties. The ISI does not measure how much a country trades but how narrowly its trade is distributed.

Distinction from supply-chain vulnerability metrics. The European Commission's supply-chain stress tests, OECD critical minerals analyses, and firm-level supply-chain risk indices typically combine concentration with assessments of *partner identity* (geopolitical alignment, governance quality, substitutability). The ISI deliberately excludes partner identity in v1.0: it measures the *geometry* of counterparty concentration—how many partners, how unequal the shares—without scoring whether those partners are allied, adversarial, or neutral. This is a narrower but more precisely defined construct: a high ISI score indicates a concentrated supplier base regardless of who the suppliers are. The advantage is measurement precision and normative neutrality; the limitation is that the ISI cannot distinguish between concentration among allied partners (potentially benign) and concentration among adversarial partners (potentially coercive).

Distinction from strategic autonomy scorecards. Strategic autonomy assessments—such as those produced by the European Council on Foreign Relations or the Istituto Affari Internazionali—incorporate military capability, institutional sovereignty, regulatory capacity, and economic self-sufficiency. The ISI measures none of these. It captures a single structural dimension—counterparty concentration of import flows—that is a necessary *input* to strategic autonomy assessments but not a substitute for them. A country with low ISI (diversified imports) may still lack strategic autonomy if it has no domestic production capacity; a country with high ISI (concentrated imports) may retain autonomy if its concentrated partners are stable allies and switching costs are low.

The ISI as concentration geometry. The unifying concept is that the ISI measures the *geometric structure* of bilateral dependence: the shape of the distribution of import shares across counterparty countries, summarised by the Herfindahl–Hirschman Index. This positions the ISI as a structural diagnostic tool—one input among several required for comprehensive vulnerability assessment—rather than as a standalone policy verdict.

9.5 Methodological limitations

Several limitations constrain the interpretation of these results. They are stated explicitly to delimit the index's evidential scope and to guide the design of future extensions.

1. **Concentration as a proxy for vulnerability.** The ISI measures *counterparty concentration*, which is a necessary but not sufficient condition for strategic vulnerability. A country may have a highly concentrated supplier base composed entirely of stable, allied partners, in which case high concentration does not translate into coercive leverage or supply-disruption risk. The ISI intentionally excludes partner-identity assessment in v1.0; incorporating geopolitical risk weighting is a planned extension.
2. **Static single-vintage design.** The ISI v1.0 is computed for a single reference year (vintage 2024). Temporal dynamics—how concentration changes in response to policy interventions, geopolitical shocks, or trade-diversion effects—cannot be assessed without panel data. A country's current rank may reflect transient conditions rather than persistent structural features.

3. **Small-state structural amplification.** Small economies (Malta, Cyprus, Luxembourg) face inherently fewer bilateral trade partners, mechanically elevating their HHI scores. The ISI captures this size effect faithfully—it *is* a structural feature of small-state dependence—but users should be aware that high ISI scores for small states partly reflect market-size constraints rather than policy choices.
4. **Re-export and transit-trade masking.** Trade statistics record the *declared* partner country, which may be a re-export hub (e.g., the Netherlands, Belgium) rather than the ultimate origin. This can understate the true concentration of ultimate-source dependence for countries whose imports transit through intermediary economies. The magnitude of this effect varies by axis and is not corrected in v1.0.
5. **Equal weighting as a normative choice.** Although the Dirichlet exercise demonstrates rank stability under random weighting, specific policy contexts exist in which defence concentration should receive higher weight than financial concentration. The ISI framework accommodates differential weighting but does not prescribe it. Users with domain-specific priors should apply the axis-level scores directly.
6. **Data-window heterogeneity.** The defence axis uses a six-year rolling window (2019–2024) while all other axes use a three-year window (2022–2024). This reflects the lower frequency of arms-transfer flows but introduces temporal asymmetry: the defence axis may reflect structural patterns that pre-date the observation windows of other axes.
7. **HHI as the sole concentration measure.** The axes use Herfindahl–Hirschman indices or comparable concentration statistics. Alternative measures (entropy, Gini of trade shares, bilateral dependency ratios) might produce different rankings, particularly for countries near the median where score gaps are smallest.
8. **Small N.** With $N = 27$, many pairwise correlations fail to reach statistical significance even when point estimates are moderate (Table 20). The geographic correlates (Section 7.11) are supported by a parsimonious OLS regression ($R^2 = 0.382$, Table 33), but the small sample limits the number of covariates that can be included without overfitting and precludes robust subgroup inference.

These limitations are inherent to the v1.0 design and do not invalidate the empirical findings. They define the boundary conditions within which the ISI results should be interpreted and provide a structured agenda for methodological extension.

10 Conclusion

This paper has presented a comprehensive empirical analysis of the International Sovereignty Index (ISI) v1.0 across the 27 member states of the European Union. The ISI aggregates six axis-level concentration scores—financial services, energy, technology, defence, critical inputs, and logistics—into a single composite indicator of strategic import dependence. The analysis has gone substantially beyond descriptive ranking to provide a rigorous statistical characterisation of the index's structure, robustness, and interpretive boundaries.

10.1 Summary of contributions

The paper makes the following analytical contributions:

1. **Complete distributional characterisation.** The ISI composite distribution is unimodal, moderately right-skewed (range 0.236–0.517, mean 0.344, $\sigma = 0.070$), with 23 of 27 countries in the "moderately concentrated" tier (Section 3).
2. **Full variance decomposition with proof.** Defence (35.10%) and logistics (34.26%) jointly account for 69.37% of composite variance; own-variance contributes 64.28% and cross-covariance 35.72% (Section 4.5).
3. **Leave-one-axis-out with Kendall concordance.** LOO analysis with both Spearman ρ and Kendall τ establishes that defence removal ($\rho = 0.833$, $\tau = 0.675$) is the most disruptive perturbation (Section 6.1).
4. **Rank displacement matrix.** Country-by-axis signed rank changes identify Italy, Slovakia, and the Netherlands as the most displacement-sensitive countries (Section 6.2).
5. **Rank elasticity to axis perturbation.** A $\pm 10\%$ axis-level perturbation exercise classifies 13 countries as rank-elastic and 14 as rank-locked, with sensitivity concentrated in the compressed mid-range (Section 6.6).
6. **Axis contribution to rank differentiation.** Bivariate R^2 values, partial correlations, and a six-axis multiple regression ($R^2 = 0.968$) quantify each axis's marginal contribution to the composite (Section 4.6).
7. **Classification compression diagnostics.** The η^2 statistic (0.483) for the four-tier classification is compared with tercile and quartile alternatives (Section 3.7).
8. **Defence axis structural analysis.** Sixteen of 27 member states (59.3%) have defence as their dominant axis; the mean gap between the defence score and the next-highest axis is 0.184 (Section 4.7).
9. **Defence-dominance stress test.** Three counterfactual composites (without defence, without logistics, without both) demonstrate that the ranking retains structural coherence ($\rho \geq 0.654$) even when the two highest-variance axes are entirely excluded (Section 4.8).
10. **Formal structural typology.** Countries are classified by dominance gap Δ_i into single-axis dominated, dual-axis concentrated, and balanced profiles, linking profile polarisation to rank volatility (Section 4.4.1).

11. **Cross-axis significance testing with multiple-testing correction.** Formal t-tests with Bonferroni correction identify only two statistically significant correlations at $\alpha = 0.05$: energy–technology ($r = 0.530$, $t = 3.124$) and critical inputs–logistics ($r = 0.549$, $t = 3.280$) (Section 5.2).
12. **PCA dimensionality diagnostic with component loadings.** Three principal components exceed the Kaiser criterion (76.1% of variance); PC1 captures goods-trade concentration, PC2 captures defence–logistics, and PC3 isolates the financial axis (Table 22).
13. **Compensability analysis.** The max–composite gap, max–min range, and coefficient of variation identify the Netherlands (gap = 0.624, CV = 0.880) as the most internally compensated country, revealing hidden single-axis vulnerabilities masked by the composite (Section 8).
14. **Geographic and structural correlates.** Island status, population size, and founding-member status correlate with the composite ISI but not with the financial axis, reinforcing its orthogonality to goods-trade determinants (Section 7.11).
15. **Monte Carlo and winsorization robustness.** Dirichlet weight perturbation (mean $\rho = 0.979$) and 95th percentile winsorization ($\rho = 1.000$) confirm overall ranking stability (Sections 6.4 and 6.5).
16. **Eigenvalue and dimensionality analysis.** Three principal components exceed the Kaiser criterion, explaining 76.1% of variance, with an effective dimensionality of approximately 4.4 (Table 21).
17. **Geometric-mean aggregation comparison.** Replacing the arithmetic mean with a geometric mean yields $\rho = 0.934$ and $\tau = 0.812$, confirming that the aggregation rule is a substantive but not fragile design choice (Table 29).
18. **OLS geographic correlates.** A four-covariate regression (log population, island dummy, log GDP, EU-6 core) explains 38.2% of composite variance; island status ($p = 0.017$) is the strongest predictor (Table 33).
19. **Formal analytical rebuttal of the "defence-only" critique.** A formal proposition (Proposition 2) demonstrates that high variance contribution under equal weighting is a measurement finding about cross-country dispersion, not an artefact of differential weighting. A Slovakia-exclusion stress test confirms that the zero defence score has negligible effect on the composite's variance structure ($\rho = 0.998$ for the 26-country subsample; Section 4.8.5).

10.2 Policy implications

The five structural conclusions of the discussion (Section 9) carry direct policy implications:

- The dominance of defence and logistics variance implies that EU strategic autonomy policy should prioritise diversification of defence procurement and logistics corridors over financial or technology import diversification, at least from the perspective of reducing cross-country composite variation.

- The mid-tier ranking fragility implies that country rankings in the 0.30–0.40 ISI range should be interpreted with caution and presented with confidence intervals or rank ranges rather than point estimates.
- The compensability findings argue for policy dashboards that display axis-level profiles alongside the composite, to avoid masking single-axis vulnerabilities.
- The orthogonality of the financial axis indicates that financial-service diversification and goods-trade diversification can be pursued as independent policy tracks.

10.3 Future research

Three directions merit further investigation:

1. **Panel extension.** Computing the ISI for multiple years would enable analysis of temporal trends, policy responsiveness, and the impact of specific geopolitical events (e.g., Russia–Ukraine conflict) on concentration dynamics.
2. **Axis-specific deep dives.** Each of the six axes could be decomposed into sub-components (e.g., the defence axis into military equipment categories) to provide finer-grained policy guidance.
3. **Network analysis.** Bilateral trade-flow data underlying the HHI computations could be used to construct supplier-country networks, enabling identification of common single-supplier dependencies across multiple EU member states.

11 Data Availability Statement

The ISI v1.0 dataset used in this paper is available on request from the corresponding author. All axis-level scores, composite scores, and country classifications are reproduced in full in the tables of this paper and its appendices.

The computation script (`isi_results_paper2.py`) that generates all derived tables, robustness statistics, and coordinate files for figures is included in Section D. The script embeds the complete authoritative data snapshot (27 countries $\times$ 6 axes) and has no dependencies beyond NumPy $\geq$ 2.0. All results reported in this paper—including variance decompositions, LOO statistics, Dirichlet Monte Carlo draws, rank elasticities, compensability metrics, compression diagnostics, and geographic correlates—are reproduced deterministically by this single script.

The underlying source data for each axis is drawn from the following publicly available databases:

- **Financial axis:** BIS Locational Banking Statistics ([7]); IMF Coordinated Portfolio Investment Survey ([8]).
- **Energy axis:** IEA World Energy Statistics ([9]).
- **Technology axis:** UN Comtrade ([12]).
- **Defence axis:** SIPRI Arms Transfers Database ([11]).
- **Critical inputs axis:** Eurostat Comext ([10]); European Commission Critical Raw Materials assessment ([5]).
- **Logistics axis:** Eurostat Comext ([10]).

Specific access procedures, vintages, and extraction parameters for each source are documented in the ISI methodology paper ([6]).

References

[1] OECD, European Commission, and Joint Research Centre. *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris, 2008. ISBN 978-92-64-04345-9. doi: 10.1787/9789264043466-en.

[2] Michaela Saisana, Andrea Saltelli, and Stefano Tarantola. Uncertainty and sensitivity analysis techniques as tools for the quality assessment of composite indicators. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 168(2):307–323, 2005. doi: 10.1111/j.1467-985X.2005.00350.x.

[3] Orris C. Herfindahl. *Concentration in the U.S. Steel Industry*. PhD thesis, Columbia University, 1950.

[4] Albert O. Hirschman. The paternity of an index. *The American Economic Review*, 54(5): 761–762, 1964.

[5] European Commission. Study on the critical raw materials for the EU. Technical report, European Commission, Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs, 2023. Final report, 2023 assessment. Accessed 5 February 2025.

[6] International Sovereignty Institute — Research Unit. Methodological foundations of the International Sovereignty Index (ISI), version 1.0, 2026. ISI Paper Series, Paper 1.

[7] Bank for International Settlements. Locational banking statistics. `https://www.bis.org/statistics/bankstats.htm`, 2024. Table A6.2, positions by counterparty country. Accessed 15 January 2025.

[8] International Monetary Fund. Coordinated portfolio investment survey (CPIS). `https://data.imf.org/cpis`, 2024. Derived portfolio positions by counterparty economy, vintage December 2024. Accessed 15 January 2025.

[9] International Energy Agency. World energy statistics and balances. `https://www.iea.org/data-and-statistics/data-product/world-energy-statistics-and-balances`, 2024. 2024 edition. Accessed 12 January 2025.

[10] Eurostat. COMEXT — EU trade since 1988 by HS2–HS6. `https://ec.europa.eu/eurostat/web/international-trade-in-goods/data/focus-on-comext`, 2024. Dataset code DS-045409, HS 2022 revision. Accessed 20 January 2025.

[11] Stockholm International Peace Research Institute. SIPRI arms transfers database. `https://www.sipri.org/databases/armstransfers`, 2024. Trend-indicator values (TIV), coverage 1950–2024. Accessed 10 January 2025.

[12] United Nations. UN Comtrade database. https://comtrade.un.org/, 2024. HS 2017 revision, annual trade data. Accessed 18 January 2025.

[13] U.S. Department of Justice and Federal Trade Commission. Horizontal merger guidelines. Technical report, U.S. Department of Justice and Federal Trade Commission, 2010. Section 5.3 defines HHI thresholds: unconcentrated (below 1,500), moderately concentrated (1,500–2,500), highly concentrated (above 2,500).

A Full Axis-Level Ranking Tables

This appendix presents the full country rankings for each of the six ISI axes, ordered from most concentrated (rank 1) to least concentrated (rank 27).

A.1 Axis 1: Financial

Table 35. Financial axis ranking, EU-27, 2024.

Rk	Country	Score
1	Croatia	0.2729
2	Greece	0.2007
3	Romania	0.2003
4	Italy	0.1875
5	Portugal	0.1830
6	Estonia	0.1817
7	Slovakia	0.1616
8	Czechia	0.1611
9	Bulgaria	0.1605
10	Belgium	0.1534
11	Austria	0.1490
12	Ireland	0.1466
13	Spain	0.1455
14	Malta	0.1399
15	Lithuania	0.1334
16	Poland	0.1331
17	Cyprus	0.1240
18	Denmark	0.1240
19	Latvia	0.1231
20	Hungary	0.1222
21	Netherlands	0.1210
22	Slovenia	0.1191
23	Finland	0.1177
24	Sweden	0.1164
25	Luxembourg	0.1146
26	Germany	0.1018
27	France	0.1003

A.2 Axis 2: Energy

Table 36. Energy axis ranking, EU-27, 2024.

Rk	Country	Score
1	Denmark	0.4702
2	Ireland	0.4647
3	Austria	0.4558
4	Estonia	0.4393
5	Romania	0.4054
6	Slovenia	0.4054
7	Slovakia	0.3999
8	Hungary	0.3858
9	Czechia	0.3808
10	Latvia	0.3652
11	Finland	0.3626
12	Luxembourg	0.3598
13	Bulgaria	0.3571
14	Croatia	0.3512
15	Malta	0.3506
16	Portugal	0.3486
17	Cyprus	0.3476
18	Lithuania	0.3374
19	Greece	0.3354
20	Poland	0.3348
21	Germany	0.3290
22	Sweden	0.3272
23	Netherlands	0.3199
24	France	0.3091
25	Belgium	0.3078
26	Italy	0.3078
27	Spain	0.3018

A.3 Axis 3: Technology

Table 37. Technology axis ranking, EU-27, 2024.

Rk	Country	Score
1	Ireland	0.5869
2	Slovenia	0.3573
3	Romania	0.2969
4	Belgium	0.2945
5	Austria	0.2605
6	Lithuania	0.2417
7	Estonia	0.2271
8	Hungary	0.2253
9	Cyprus	0.2243
10	Slovakia	0.2164
11	Poland	0.2084
12	Malta	0.2054
13	Luxembourg	0.2001
14	Finland	0.1978
15	Portugal	0.1976
16	Spain	0.1926
17	Czechia	0.1824
18	Croatia	0.1734
19	Greece	0.1627
20	Bulgaria	0.1520
21	Denmark	0.1499
22	Italy	0.1410
23	Latvia	0.1330
24	Netherlands	0.1252
25	France	0.1200
26	Sweden	0.1198
27	Germany	0.0976

A.4 Axis 4: Defence

Table 38. Defence axis ranking, EU-27, 2024.

Rk	Country	Score
1	Malta	1.0000
2	Netherlands	0.9595
3	Italy	0.9205
4	Sweden	0.8810
5	Bulgaria	0.8606
6	Cyprus	0.7484
7	Croatia	0.7464
8	Denmark	0.7236
9	Slovenia	0.6666
10	Luxembourg	0.6448
11	Austria	0.5850
12	Germany	0.5826
13	Romania	0.5635
14	Greece	0.5514
15	Hungary	0.5447
16	Czechia	0.4841
17	Ireland	0.4488
18	Belgium	0.4473
19	Poland	0.4432
20	Portugal	0.4337
21	Lithuania	0.4022
22	Finland	0.3896
23	Spain	0.3808
24	France	0.3703
25	Estonia	0.3556
26	Latvia	0.3407
27	Slovakia	0.0000

A.5 Axis 5: Critical Inputs

Table 39. Critical Inputs axis ranking, EU-27, 2024.

Rk	Country	Score
1	Finland	0.5469
2	Croatia	0.5239
3	Denmark	0.4362
4	Cyprus	0.4235
5	Malta	0.4087
6	Estonia	0.3862
7	Austria	0.3185
8	Ireland	0.3071
9	Portugal	0.2595
10	Hungary	0.2340
11	Germany	0.2231
12	Slovenia	0.2230
13	Luxembourg	0.2229
14	Greece	0.2220
15	Italy	0.2083
16	Sweden	0.1925
17	Belgium	0.1724
18	Czechia	0.1697
19	Poland	0.1653
20	France	0.1612
21	Romania	0.1581
22	Slovakia	0.1549
23	Latvia	0.1512
24	Bulgaria	0.1484
25	Lithuania	0.1305
26	Netherlands	0.1301
27	Spain	0.1123

A.6 Axis 6: Logistics

Table 40. Logistics axis ranking, EU-27, 2024.

Rk	Country	Score
1	Malta	1.0000
2	Cyprus	0.9410
3	Finland	0.7973
4	Sweden	0.7551
5	Greece	0.7198
6	Denmark	0.6423
7	Luxembourg	0.6185
8	Ireland	0.6137
9	Czechia	0.5294
10	Austria	0.5259
11	Portugal	0.5077
12	Poland	0.4886
13	Slovakia	0.4855
14	Italy	0.4778
15	Estonia	0.4698
16	Hungary	0.4657
17	Slovenia	0.4471
18	Spain	0.4323
19	Latvia	0.3682
20	Bulgaria	0.3629
21	Croatia	0.3579
22	Netherlands	0.3549
23	France	0.3529
24	Belgium	0.3520
25	Lithuania	0.3468
26	Romania	0.2874
27	Germany	0.2729

A.7 Axis dominance detail

Table 41 lists the max-axis and second-max-axis for each country.

Table 41. Axis dominance per country: highest- and second-highest-scoring axes.

Rk	Country	Max	Max Axis	2nd	2nd Axis
1	Malta	1.0000	Defence	1.0000	Logistics
2	Cyprus	0.9410	Logistics	0.7484	Defence
3	Ireland	0.6137	Logistics	0.5869	Technology
4	Denmark	0.7236	Defence	0.6423	Logistics
5	Croatia	0.7464	Defence	0.5239	Crit. Inputs
6	Finland	0.7973	Logistics	0.5469	Crit. Inputs
7	Sweden	0.8810	Defence	0.7551	Logistics
8	Austria	0.5850	Defence	0.5259	Logistics
9	Italy	0.9205	Defence	0.4778	Logistics
10	Slovenia	0.6666	Defence	0.4471	Logistics
11	Greece	0.7198	Logistics	0.5514	Defence
12	Luxembourg	0.6448	Defence	0.6185	Logistics
13	Estonia	0.4698	Logistics	0.4393	Energy
14	Bulgaria	0.8606	Defence	0.3629	Logistics
15	Netherlands	0.9595	Defence	0.3549	Logistics
16	Hungary	0.5447	Defence	0.4657	Logistics
17	Portugal	0.5077	Logistics	0.4337	Defence
18	Romania	0.5635	Defence	0.4054	Energy
19	Czechia	0.5294	Logistics	0.4841	Defence
20	Poland	0.4886	Logistics	0.4432	Defence
21	Belgium	0.4473	Defence	0.3520	Logistics
22	Germany	0.5826	Defence	0.3290	Energy
23	Lithuania	0.4022	Defence	0.3468	Logistics
24	Spain	0.4323	Logistics	0.3808	Defence
25	Latvia	0.3682	Logistics	0.3652	Energy
26	Slovakia	0.4855	Logistics	0.3999	Energy
27	France	0.3703	Defence	0.3529	Logistics

A.8 Covariance matrix

Table 42 reports the full 6×6 covariance matrix of the axis-level scores. Diagonal entries are the own-variances σ_j^2 used in the variance decomposition (Section 4.5).

Table 42. Population covariance matrix of the six ISI axes, EU-27, 2024.

	Fin.	**Ene.**	**Tec.**	**Def.**	**Crit.**	**Log.**
Financial	0.001390	0.000111	0.000267	0.000199	0.000988	−0.001137
Energy	0.000111	0.002295	0.002401	−0.001894	0.002171	0.000997
Technology	0.000267	0.002401	0.008947	−0.004735	0.001113	0.001205
Defence	0.000199	−0.001894	−0.004735	0.050922	0.005251	0.012038
Critical Inputs	0.000988	0.002171	0.001113	0.005251	0.014841	0.012456
Logistics	−0.001137	0.000997	0.001205	0.012038	0.012456	0.034748

B Robustness Detail Tables

This appendix provides the full country-level detail for each robustness exercise summarised in Section 6.

B.1 Leave-one-axis-out: full rank table

Table 43 reports the rank of each country under each LOO variant (baseline, drop axis 1, …, drop axis 6), along with the maximum rank change from baseline.

Table 43. Detailed rank comparison: baseline vs. each leave-one-axis-out variant. Columns show the rank under each LOO composite. $\Delta_{\max}$ is the maximum absolute rank change across all six LOO variants.

Rk	Country	−Fin	−Ene	−Tec	−Def	−Cri	−Log	$\Delta_{\max}$
1	Malta	1	1	1	2	1	1	1
2	Cyprus	2	2	2	3	2	5	3
3	Ireland	4	3	10	1	3	3	7
4	Denmark	3	4	3	5	5	4	1
5	Croatia	7	5	5	8	11	2	6
6	Finland	5	7	6	4	14	13	8
7	Sweden	6	6	4	12	4	11	5
8	Austria	8	10	8	6	8	7	2
9	Italy	10	8	7	19	6	8	10
10	Slovenia	9	11	14	10	7	6	4
11	Greece	12	9	9	9	9	17	6
12	Luxembourg	11	12	11	11	10	15	3
13	Estonia	15	15	15	7	18	14	6
14	Bulgaria	14	14	12	23	12	9	9
15	Netherlands	13	13	13	25	13	10	10
16	Hungary	16	16	16	14	16	16	2
17	Portugal	17	17	17	13	19	18	4
18	Romania	19	19	19	17	15	12	6
19	Czechia	18	18	18	15	17	19	4
20	Poland	20	20	20	18	20	22	2
21	Belgium	21	21	22	20	21	20	1
22	Germany	22	22	21	27	24	21	5
23	Lithuania	23	24	24	21	22	23	2
24	Spain	24	23	23	22	23	24	2
25	Latvia	25	25	25	24	25	25	1
26	Slovakia	27	27	27	16	26	27	10
27	France	26	26	26	26	27	26	1

B.2 Individual LOO axis tables

The following tables report the full score and rank for each country under each individual axis omission.

B.2.1 LOO: Financial axis omitted

Table 44. LOO ranking excluding Financial. Spearman $\rho = 0.9927$, Kendall $\tau = 0.9430$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Malta	0.5929	1	+0
2	Cyprus	0.5370	2	+0
3	Denmark	0.4845	4	+1
4	Ireland	0.4842	3	-1
5	Finland	0.4589	6	+1
6	Sweden	0.4551	7	+1
7	Croatia	0.4306	5	-2
8	Austria	0.4291	8	+0
9	Slovenia	0.4199	10	+1
10	Italy	0.4111	9	-1
11	Luxembourg	0.4092	12	+1
12	Greece	0.3982	11	-1
13	Netherlands	0.3779	15	+2
14	Bulgaria	0.3762	14	+0
15	Estonia	0.3756	13	-2
16	Hungary	0.3711	16	+0
17	Portugal	0.3494	17	+0
18	Czechia	0.3493	19	+1
19	Romania	0.3423	18	-1
20	Poland	0.3281	20	+0
21	Belgium	0.3148	21	+0
22	Germany	0.3010	22	+0
23	Lithuania	0.2917	23	+0
24	Spain	0.2840	24	+0
25	Latvia	0.2717	25	+0
26	France	0.2627	27	+1
27	Slovakia	0.2513	26	-1

B.2.2 LOO: Energy axis omitted

Table 45. LOO ranking excluding Energy. Spearman $\rho = 0.9921$, Kendall $\tau = 0.9430$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Malta	0.5508	1	+0
2	Cyprus	0.4923	2	+0
3	Ireland	0.4206	3	+0
4	Denmark	0.4152	4	+0
5	Croatia	0.4149	5	+0
6	Sweden	0.4130	7	+1
7	Finland	0.4099	6	-1
8	Italy	0.3870	9	+1
9	Greece	0.3713	11	+2
10	Austria	0.3678	8	-2
11	Slovenia	0.3626	10	-1
12	Luxembourg	0.3602	12	+0
13	Netherlands	0.3382	15	+2
14	Bulgaria	0.3369	14	+0
15	Estonia	0.3241	13	-2
16	Hungary	0.3184	16	+0
17	Portugal	0.3163	17	+0
18	Czechia	0.3053	19	+1
19	Romania	0.3012	18	-1
20	Poland	0.2877	20	+0
21	Belgium	0.2839	21	+0
22	Germany	0.2556	22	+0
23	Spain	0.2527	24	+1
24	Lithuania	0.2509	23	-1
25	Latvia	0.2232	25	+0
26	France	0.2209	27	+1
27	Slovakia	0.2037	26	-1

B.2.3 LOO: Technology axis omitted

Table 46. LOO ranking excluding Technology. Spearman $\rho = 0.9683$, Kendall $\tau = 0.8860$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Malta	0.5798	1	+0
2	Cyprus	0.5169	2	+0
3	Denmark	0.4793	4	+1
4	Sweden	0.4544	7	+3
5	Croatia	0.4505	5	+0
6	Finland	0.4428	6	+0
7	Italy	0.4204	9	+2
8	Austria	0.4068	8	+0
9	Greece	0.4059	11	+2
10	Ireland	0.3962	3	-7
11	Luxembourg	0.3921	12	+1
12	Bulgaria	0.3779	14	+2
13	Netherlands	0.3771	15	+2
14	Slovenia	0.3722	10	-4
15	Estonia	0.3665	13	-2
16	Hungary	0.3505	16	+0
17	Portugal	0.3465	17	+0
18	Czechia	0.3450	19	+1
19	Romania	0.3229	18	-1
20	Poland	0.3130	20	+0
21	Germany	0.3019	22	+1
22	Belgium	0.2866	21	-1
23	Spain	0.2746	24	+1
24	Lithuania	0.2701	23	-1
25	Latvia	0.2697	25	+0
26	France	0.2588	27	+1
27	Slovakia	0.2404	26	-1

B.2.4 LOO: Defence axis omitted

Table 47. LOO ranking excluding Defence. Spearman $\rho = 0.8327$, Kendall $\tau = 0.6752$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Ireland	0.4238	3	+2
2	Malta	0.4209	1	-1
3	Cyprus	0.4121	2	-1
4	Finland	0.4045	6	+2
5	Denmark	0.3645	4	-1
6	Austria	0.3419	8	+2
7	Estonia	0.3408	13	+6
8	Croatia	0.3359	5	-3
9	Greece	0.3281	11	+2
10	Slovenia	0.3104	10	+0
11	Luxembourg	0.3032	12	+1
12	Sweden	0.3022	7	-5
13	Portugal	0.2993	17	+4
14	Hungary	0.2866	16	+2
15	Czechia	0.2847	19	+4
16	Slovakia	0.2836	26	+10
17	Romania	0.2696	18	+1
18	Poland	0.2660	20	+2
19	Italy	0.2645	9	-10
20	Belgium	0.2560	21	+1
21	Lithuania	0.2380	23	+2
22	Spain	0.2369	24	+2
23	Bulgaria	0.2362	14	-9
24	Latvia	0.2281	25	+1
25	Netherlands	0.2102	15	-10
26	France	0.2087	27	+1
27	Germany	0.2049	22	-5

B.2.5 LOO: Critical inputs axis omitted

Table 48. LOO ranking excluding Critical Inputs. Spearman $\rho = 0.9414$, Kendall $\tau = 0.8462$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Malta	0.5392	1	+0
2	Cyprus	0.4771	2	+0
3	Ireland	0.4521	3	+0
4	Sweden	0.4399	7	+3
5	Denmark	0.4220	4	-1
6	Italy	0.4069	9	+3
7	Slovenia	0.3991	10	+3
8	Austria	0.3952	8	+0
9	Greece	0.3940	11	+2
10	Luxembourg	0.3876	12	+2
11	Croatia	0.3804	5	-6
12	Bulgaria	0.3786	14	+2
13	Netherlands	0.3761	15	+2
14	Finland	0.3730	6	-8
15	Romania	0.3507	18	+3
16	Hungary	0.3487	16	+0
17	Czechia	0.3476	19	+2
18	Estonia	0.3347	13	-5
19	Portugal	0.3341	17	-2
20	Poland	0.3216	20	+0
21	Belgium	0.3110	21	+0
22	Lithuania	0.2923	23	+1
23	Spain	0.2906	24	+1
24	Germany	0.2768	22	-2
25	Latvia	0.2661	25	+0
26	Slovakia	0.2527	26	+0
27	France	0.2505	27	+0

B.2.6 LOO: Logistics axis omitted

Table 49. LOO ranking excluding Logistics. Spearman $\rho = 0.9261$, Kendall $\tau = 0.7892$.

LOO Rk	Country	LOO Comp.	Base Rk	Δ
1	Malta	0.4209	1	+0
2	Croatia	0.4136	5	+3
3	Ireland	0.3908	3	+0
4	Denmark	0.3808	4	+0
5	Cyprus	0.3736	2	-3
6	Slovenia	0.3543	10	+4
7	Austria	0.3538	8	+1
8	Italy	0.3530	9	+1
9	Bulgaria	0.3357	14	+5
10	Netherlands	0.3312	15	+5
11	Sweden	0.3274	7	-4
12	Romania	0.3248	18	+6
13	Finland	0.3229	6	-7
14	Estonia	0.3180	13	-1
15	Luxembourg	0.3084	12	-3
16	Hungary	0.3024	16	+0
17	Greece	0.2944	11	-6
18	Portugal	0.2845	17	-1
19	Czechia	0.2756	19	+0
20	Belgium	0.2751	21	+1
21	Germany	0.2668	22	+1
22	Poland	0.2570	20	-2
23	Lithuania	0.2491	23	+0
24	Spain	0.2266	24	+0
25	Latvia	0.2226	25	+0
26	France	0.2122	27	+1
27	Slovakia	0.1865	26	-1

B.3 Z-score standardisation: full rank comparison

Table 25 in Section 6.3 compares baseline and z-score ranks for all 27 countries. See Table 25 for the full table.

B.4 Dirichlet weight perturbation: full volatility table

Table 26 in Section 6.4 reports mean rank, rank standard deviation, and 5th/95th percentile ranks for all 27 countries. See Table 26 for the full table.

B.5 Winsorization: full rank comparison

Table 27 in Section 6.5 compares baseline and winsorized ranks for all 27 countries. See Table 27 for the full table.

C Reproducibility Protocol

This appendix documents the complete procedure for reproducing all results presented in this paper.

C.1 Software requirements

- **Python** $\geq$ 3.10 (tested with 3.14.0).
- **NumPy** $\geq$ 2.0 (tested with 2.2.6). NumPy is the only external dependency.
- **LaTeX** distribution with pdfLaTeX (tested with TeX Live 2025), including packages: `pgfplots` ($\geq$ 1.18), `tikz`, `siunitx`, `longtable`, `booktabs`, `natbib`, `cleveref`, `threeparttable`, `pdflscape`, `adjustbox`, `subcaption`, `listings`, `xcolor`.

C.2 Directory structure

The project is organised as follows:

```
paper2/
+-- main.tex                    % Master document
+-- references.bib              % Bibliography
+-- sections/
|   +-- 00_executive_summary.tex
|   +-- 01_introduction.tex
|   +-- ...
|   +-- 10_data_availability.tex
+-- tables/                     % Generated .tex table fragments
|   +-- T06_axis_top_bottom.tex
|   +-- T07_axis_dominance.tex
|   +-- ...
|   +-- radar_*.dat
+-- figures/                    % TikZ figure source files
|   +-- F01_histogram.tex
|   +-- ...
|   +-- F15_radar_FR.tex
+-- appendix/
|   +-- appendix_a_tables.tex
|   +-- appendix_b_robustness.tex
|   +-- appendix_c_reproducibility.tex
|   +-- appendix_d_code.tex
+-- scripts/
    +-- isi_results_paper2.py
```

C.3 Step-by-step reproduction

1. **Install Python and NumPy.**

```
python -m venv .venv
source .venv/bin/activate
pip install numpy
```

2. **Run the computation script.**

```
cd paper2/scripts
python isi_results_paper2.py
```

   This generates all `.tex` table fragments and `.dat` coordinate files in the `tables/` directory. The script embeds the full ISI v1.0 data snapshot and performs all computations (LOO, z-score, Dirichlet MC, winsorization, variance decomposition, Lorenz curve, ECDF, histograms, radar coordinates).

3. **Compile the LaTeX document.**

```
cd paper2
pdflatex main.tex
bibtex main
pdflatex main.tex
pdflatex main.tex
```

   Three passes are required to resolve cross-references and the bibliography.

4. **Verify.** The final PDF should compile with 0 errors, 0 warnings, and 0 overfull boxes.

C.4 Random seed

The Dirichlet Monte Carlo exercise uses a fixed random seed (`np.random.seed(42)`) to ensure exact reproducibility of rank volatility statistics. Changing the seed will produce quantitatively similar but not identical results.

C.5 Data integrity

The ISI v1.0 data snapshot is embedded directly in the Python script (`isi_results_paper2.py`, lines 37–77) as a list of tuples. No external data files need to be downloaded or configured. The embedded data constitutes the single authoritative source for all computations in this paper.

D Computation Code

The complete Python script used to generate all derived tables and data files is listed below. The script requires only NumPy and embeds the full ISI v1.0 data snapshot.

Listing 1. ISI Results Paper 2 — Computation Engine (`isi_results_paper2.py`).

```python
#!/usr/bin/env python3
"""
ISI Results Paper 2 -- Computation Engine (Expanded)
====================================================
Reads the ISI v1.0 EU-27 snapshot (embedded) and produces all derived
tables for Paper 2.

Computations:
  1.  Axis dominance table
  2.  Variance contribution decomposition (full, with own vs cross)
  3.  Leave-one-axis-out (LOO) with Spearman and Kendall tau
  4.  Z-score standardized composite
  5.  Dirichlet weight perturbation (Monte Carlo)
  6.  Winsorization (cap Axis 4 & 6 at P95)
  7.  Top-5 / Bottom-5 per axis
  8.  Lorenz curve coordinates
  9.  Axis rank tables (Appendix A)
  10. Histogram / ECDF coordinates
  11. Radar plot data per country
  12. Boxplot statistics
  13. Rank elasticity to axis perturbation
  14. Axis contribution to rank differentiation (R-squared, partial corr)
  15. Classification compression diagnostics
  16. Defence axis structural analysis
  17. Compensability analysis
  18. Geographic / structural correlates
  19. Correlation t-statistics
  20. Spearman correlation matrix
  21. LOO individual axis tables (6 tables for appendix)
  22. Multiple regression diagnostics

Usage:
    python isi_results_paper2.py

No external dependencies beyond numpy.
"""

import numpy as np
import os

AXIS_NAMES = [
    "Financial", "Energy", "Technology",
    "Defence", "Critical Inputs", "Logistics"
]
AXIS_SLUGS = [
    "financial", "energy", "technology",
    "defence", "critical_inputs", "logistics"
]

# fmt: off
DATA = [
    ("MT", "Malta",        1,  0.51748148, "highly_concentrated",     0.13998786, 0.35068126,
         0.20546991, 1.00000000, 0.40874987, 1.00000000),
    ("CY", "Cyprus",       2,  0.46821264, "moderately_concentrated", 0.12407171, 0.34768252,
         0.22438354, 0.74847998, 0.42359750, 0.94106062),
    ("IE", "Ireland",      3,  0.42802519, "moderately_concentrated", 0.14666620, 0.46472525,
         0.58695939, 0.44887460, 0.30719544, 0.61373023),
    ("DK", "Denmark",      4,  0.42442738, "moderately_concentrated", 0.12401663, 0.47029821,
         0.14995928, 0.72368274, 0.43629935, 0.64230809),
    ("HR", "Croatia",      5,  0.40434768, "moderately_concentrated", 0.27299265, 0.35125471,
         0.17343982, 0.74648317, 0.52392772, 0.35798800),
    ("FI", "Finland",      6,  0.40205136, "moderately_concentrated", 0.11775073, 0.36268783,
         0.19789321, 0.38964518, 0.54699996, 0.79733123),
    ("SE", "Sweden",       7,  0.39871532, "moderately_concentrated", 0.11646504, 0.32721403,
         0.11982797, 0.88107943, 0.19257891, 0.75512656),
    ("AT", "Austria",      8,  0.38249806, "moderately_concentrated", 0.14903467, 0.45581191,
         0.26057910, 0.58506379, 0.31857781, 0.52592111),
    ("IT", "Italy",        9,  0.37386086, "moderately_concentrated", 0.18753526, 0.30781423,
         0.14103384, 0.92057651, 0.20839534, 0.47781001),
    ("SI", "Slovenia",     10, 0.36979186, "moderately_concentrated", 0.11917574, 0.40540014,
         0.35739912, 0.66661871, 0.22301125, 0.44714621),
    ("EL", "Greece",       11, 0.36537350, "moderately_concentrated", 0.20078311, 0.33541623,
         0.16272142, 0.55148593, 0.22200621, 0.71982808),
```

```
    ("LU", "Luxembourg",   12, 0.36016893, "moderately_concentrated", 0.11462818, 0.35980871,
        0.20019721, 0.64483831, 0.22295831, 0.61858285),
    ("EE", "Estonia",      13, 0.34332177, "moderately_concentrated", 0.18172493, 0.43935251,
        0.22713142, 0.35561474, 0.38622776, 0.46987927),
    ("BG", "Bulgaria",     14, 0.34027984, "moderately_concentrated", 0.16052965, 0.35711773,
        0.15204272, 0.86064518, 0.14843767, 0.36290606),
    ("NL", "Netherlands",  15, 0.33515702, "moderately_concentrated", 0.12107826, 0.31992851,
        0.12527375, 0.95955396, 0.13018835, 0.35491930),
    ("HU", "Hungary",      16, 0.32966227, "moderately_concentrated", 0.12222969, 0.38580312,
        0.22538122, 0.54471981, 0.23405214, 0.46578765),
    ("PT", "Portugal",     17, 0.32173430, "moderately_concentrated", 0.18303606, 0.34860414,
        0.19766062, 0.43376254, 0.25955617, 0.50778629),
    ("RO", "Romania",      18, 0.31864192, "moderately_concentrated", 0.20033340, 0.40546401,
        0.29691153, 0.56353195, 0.15814280, 0.28746783),
    ("CZ", "Czechia",      19, 0.31796630, "moderately_concentrated", 0.16116762, 0.38082003,
        0.18242269, 0.48413737, 0.16976911, 0.52948100),
    ("PL", "Poland",       20, 0.29561321, "moderately_concentrated", 0.13314812, 0.33487780,
        0.20845359, 0.44326493, 0.16532767, 0.48860717),
    ("BE", "Belgium",      21, 0.28796781, "moderately_concentrated", 0.15348744, 0.30789144,
        0.29454573, 0.44734854, 0.17248330, 0.35205038),
    ("DE", "Germany",      22, 0.26788694, "moderately_concentrated", 0.10186464, 0.32908790,
        0.09762824, 0.58266751, 0.22312039, 0.27295297),
    ("LT", "Lithuania",    23, 0.26540520, "moderately_concentrated", 0.13349636, 0.33747627,
        0.24172279, 0.40227798, 0.13056067, 0.34689716),
    ("ES", "Spain",        24, 0.26093977, "moderately_concentrated", 0.14554195, 0.30187479,
        0.19263249, 0.38088569, 0.11235234, 0.43235135),
    ("LV", "Latvia",       25, 0.24695137, "mildly_concentrated",     0.12316988, 0.36522282,
        0.13309607, 0.34079636, 0.15120012, 0.36822299),
    ("SK", "Slovakia",     26, 0.23641567, "mildly_concentrated",     0.16163673, 0.39995460,
        0.21649186, 0.00000000, 0.15490434, 0.48550648),
    ("FR", "France",       27, 0.23567126, "mildly_concentrated",     0.10039135, 0.30911838,
        0.12001059, 0.37034617, 0.16122463, 0.35293643),
]
# fmt: on

N = len(DATA)
assert N == 27

ISO   = [d[0] for d in DATA]
NAMES = [d[1] for d in DATA]
RANKS = np.array([d[2] for d in DATA])
COMP  = np.array([d[3] for d in DATA])
CLASS = [d[4] for d in DATA]
AXES  = np.array([[d[5+j] for j in range(6)] for d in DATA])

OUT_DIR = os.path.join(os.path.dirname(os.path.abspath(__file__)), "..", "tables")
os.makedirs(OUT_DIR, exist_ok=True)

def write_table(filename, content):
    path = os.path.join(OUT_DIR, filename)
    with open(path, "w") as f:
        f.write(content)
    print(f"␣␣Written:␣{path}")

def rankdata(x):
    n = len(x)
    order = np.argsort(-x)
    ranks = np.empty(n, dtype=float)
    ranks[order] = np.arange(1, n+1, dtype=float)
    return ranks

def spearman_corr(r1, r2):
    d = r1 - r2
    n = len(r1)
    return 1.0 - (6.0 * np.sum(d**2)) / (n * (n**2 - 1))

def kendall_tau(r1, r2):
    n = len(r1)
    concordant = 0
    discordant = 0
    for i in range(n):
        for j in range(i+1, n):
            s1 = np.sign(r1[i] - r1[j])
            s2 = np.sign(r2[i] - r2[j])
            prod = s1 * s2
            if prod > 0:
                concordant += 1
            elif prod < 0:
                discordant += 1
    pairs = 0.5 * n * (n - 1)
    return (concordant - discordant) / pairs if pairs > 0 else 0.0

def pearson_corr(x, y):
```

```python
    mx, my = np.mean(x), np.mean(y)
    num = np.sum((x - mx) * (y - my))
    den = np.sqrt(np.sum((x - mx)**2) * np.sum((y - my)**2))
    return num / den if den > 0 else 0.0


# ========================================================================
# 1. AXIS DOMINANCE TABLE
# ========================================================================
print("\n" + "="*72)
print("1.␣AXIS␣DOMINANCE")
print("="*72)

max_axis_idx = np.argmax(AXES, axis=1)
second_max_idx = np.zeros(N, dtype=int)
for i in range(N):
    sorted_idx = np.argsort(-AXES[i])
    second_max_idx[i] = sorted_idx[1]

axis_dom_counts = np.zeros(6, dtype=int)
for i in range(N):
    axis_dom_counts[max_axis_idx[i]] += 1

for j in range(6):
    print(f"␣␣{AXIS_NAMES[j]:20s}:␣{axis_dom_counts[j]}")

lines = []
lines.append("%␣Table␣T7:␣Axis␣dominance␣per␣country")
lines.append("\\begin{longtable}{@{}rlS[table-format=1.8]lS[table-format=1.8]l@{}}")
lines.append("\\caption{Axis␣dominance␣per␣country:␣highest-␣and␣second-highest-scoring␣axes.}")
lines.append("\\label{tab:axis-dominance}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Max␣Score}}␣&␣{\\textbf{Max␣Axis}}␣&␣{\\textbf{2nd␣Score}}␣&␣{\\textbf{2nd␣Axis}}␣\\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Axis␣dominance␣per␣country␣(continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Max␣Score}}␣&␣{\\textbf{Max␣Axis}}␣&␣{\\textbf{2nd␣Score}}␣&␣{\\textbf{2nd␣Axis}}␣\\\\")
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
for i in range(N):
    mi = max_axis_idx[i]
    si = second_max_idx[i]
    lines.append(f"␣␣{RANKS[i]:.0f}␣&␣{NAMES[i]}␣&␣{AXES[i,mi]:.8f}␣&␣{AXIS_NAMES[mi]}␣&␣{AXES[i,si]:.8f}␣&␣{AXIS_NAMES[si]}␣\\\\")
lines.append("\\end{longtable}")
write_table("T07_axis_dominance.tex", "\n".join(lines))


# ========================================================================
# 2. VARIANCE CONTRIBUTION DECOMPOSITION (FULL)
# ========================================================================
print("\n" + "="*72)
print("2.␣VARIANCE␣CONTRIBUTION␣(FULL)")
print("="*72)

axis_vars = np.var(AXES, axis=0, ddof=0)
cov_matrix = np.cov(AXES.T, ddof=0)
total_composite_var = np.var(COMP, ddof=0)
var_decomp_check = np.sum(cov_matrix) / 36.0

print(f"␣␣Composite␣variance␣(direct):␣␣␣␣{total_composite_var:.10f}")
print(f"␣␣Composite␣variance␣(from␣cov):␣␣{var_decomp_check:.10f}")

marginal_contrib = np.sum(cov_matrix, axis=1) / 36.0
pct_contrib = marginal_contrib / var_decomp_check * 100.0

own_var_share = np.array([cov_matrix[j, j] / 36.0 for j in range(6)])
cov_share = np.array([sum(cov_matrix[j, k] / 36.0 for k in range(6) if k != j) for j in range(6)])

total_own = sum(cov_matrix[j, j] for j in range(6)) / 36.0
total_cross = sum(cov_matrix[j, k] for j in range(6) for k in range(6) if j != k) / 36.0
print(f"␣␣Total␣own-variance:␣␣␣{total_own:.10f}␣({total_own/var_decomp_check*100:.2f}%)")
print(f"␣␣Total␣cross-covar:␣␣␣␣{total_cross:.10f}␣({total_cross/var_decomp_check*100:.2f}%)")

for j in range(6):
    print(f"␣␣{AXIS_NAMES[j]:20s}:␣var={axis_vars[j]:.8f}␣␣marginal={marginal_contrib[j]:.8f}␣␣pct={pct_contrib[j]:.2f}%")
print(f"␣␣Defence+Logistics:␣␣␣␣{pct_contrib[3]+pct_contrib[5]:.2f}%")
```

```
# T08: summary
lines = []
lines.append("%␣Table␣T8:␣Variance␣contribution␣decomposition")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Variance␣contribution␣decomposition␣of␣the␣ISI␣composite.␣␣Each␣axis's")
lines.append("marginal␣contribution␣equals␣$(1/36)\\sum_{k=1}^{6}\\operatorname{Cov}(A_j,␣A_k)$,")
lines.append("reflecting␣both␣own-variance␣and␣covariance␣with␣other␣axes.}")
lines.append("\\label{tab:variance-decomp}")
lines.append("\\begin{tabular}{@{}lS[table-format=1.8]S[table-format=1.8]S[table-format=2.2]@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Axis}}␣&␣{\\textbf{Var␣(pop.)}}␣&␣{\\textbf{Marginal␣Contrib.}}␣&␣{\\textbf{\\%␣of␣Composite␣Var}}␣\\\\")
lines.append("\\midrule")
for j in range(6):
    lines.append(f"␣␣{AXIS_NAMES[j]}␣&␣{axis_vars[j]:.8f}␣&␣{marginal_contrib[j]:.8f}␣&␣{pct_contrib[j]:.2f}␣\\\\")
lines.append("\\midrule")
lines.append(f"␣␣Total␣&␣{np.sum(axis_vars):.8f}␣&␣{var_decomp_check:.8f}␣&␣{np.sum(pct_contrib):.2f}␣\\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T08_variance_decomp.tex", "\n".join(lines))

# T08b: detailed own vs cross
lines = []
lines.append("%␣Table␣T8b:␣Detailed␣own-variance␣vs␣covariance␣decomposition")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Decomposition␣of␣each␣axis's␣marginal␣contribution␣into␣own-variance␣and")
lines.append("cross-covariance␣components.␣␣Own:␣$(1/36)\\operatorname{Var}(A_j)$.")
lines.append("Cross:␣$(1/36)\\sum_{k␣\\neq␣j}\\operatorname{Cov}(A_j,␣A_k)$.}")
lines.append("\\label{tab:variance-decomp-detail}")
lines.append("\\small")
lines.append("\\begin{tabular}{@{}l␣S[table-format=1.8]␣S[table-format=+1.8]␣S[table-format=1.8]␣S[table-format=2.2]␣S[table-format=+2.2]␣S[table-format=2.2]@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Axis}}␣&␣{\\textbf{Own␣Var}}␣&␣{\\textbf{Cross␣Cov}}␣&␣{\\textbf{Marginal}}␣&␣{\\textbf{\\%␣Own}}␣&␣{\\textbf{\\%␣Cross}}␣&␣{\\textbf{\\%␣Total}}␣\\\\")
lines.append("\\midrule")
for j in range(6):
    pct_own = own_var_share[j] / var_decomp_check * 100
    pct_cross = cov_share[j] / var_decomp_check * 100
    lines.append(f"␣␣{AXIS_NAMES[j]}␣&␣{own_var_share[j]:.8f}␣&␣{cov_share[j]:+.8f}␣&␣{marginal_contrib[j]:.8f}␣&␣{pct_own:.2f}␣&␣{pct_cross:+.2f}␣&␣{pct_contrib[j]:.2f}␣\\\\")
lines.append("\\midrule")
lines.append(f"␣␣Total␣&␣{total_own:.8f}␣&␣{total_cross:+.8f}␣&␣{var_decomp_check:.8f}␣&␣{total_own/var_decomp_check*100:.2f}␣&␣{total_cross/var_decomp_check*100:.2f}␣&␣100.00␣\\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T08b_variance_decomp_detail.tex", "\n".join(lines))

# TE1: Full covariance matrix
lines = []
lines.append("%␣Table␣TE1:␣Full␣6x6␣population␣covariance␣matrix")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Population␣covariance␣matrix␣of␣the␣six␣ISI␣axes,␣EU-27,␣2024.}")
lines.append("\\label{tab:cov-matrix}")
lines.append("\\small")
lines.append("\\sisetup{table-format=+1.6}")
lines.append("\\begin{tabular}{@{}l␣S␣S␣S␣S␣S␣S@{}}")
lines.append("\\toprule")
lines.append("␣&␣{\\textbf{Fin.}}␣&␣{\\textbf{Ene.}}␣&␣{\\textbf{Tec.}}␣&␣{\\textbf{Def.}}␣&␣{\\textbf{Crit.}}␣&␣{\\textbf{Log.}}␣\\\\")
lines.append("\\midrule")
for j in range(6):
    row = f"␣␣{AXIS_NAMES[j]}"
    for k in range(6):
        row += f"␣&␣{cov_matrix[j,k]:+.6f}"
    row += "␣\\\\"
    lines.append(row)
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("TE1_cov_matrix.tex", "\n".join(lines))


# ========================================================================
# 3. LEAVE-ONE-AXIS-OUT (LOO)
```

```python
# ========================================================================
print("\n" + "="*72)
print("3.␣LEAVE-ONE-AXIS-OUT␣(LOO)")
print("="*72)

baseline_ranks = rankdata(COMP)

loo_composites = np.zeros((N, 6))
loo_ranks = np.zeros((N, 6))
loo_spearman = np.zeros(6)
loo_kendall = np.zeros(6)

for j in range(6):
    mask = [k for k in range(6) if k != j]
    loo_comp = np.mean(AXES[:, mask], axis=1)
    loo_composites[:, j] = loo_comp
    lr = rankdata(loo_comp)
    loo_ranks[:, j] = lr
    loo_spearman[j] = spearman_corr(baseline_ranks, lr)
    loo_kendall[j] = kendall_tau(baseline_ranks, lr)
    print(f"␣␣Drop␣{AXIS_NAMES[j]:20s}:␣Spearman={loo_spearman[j]:.6f}␣␣Kendall={loo_kendall[j]:.6f}")

rank_displacement = np.zeros((N, 6), dtype=int)
for i in range(N):
    for j in range(6):
        rank_displacement[i, j] = int(baseline_ranks[i] - loo_ranks[i, j])

max_rank_change = np.max(np.abs(rank_displacement), axis=1)
mean_abs_disp = np.mean(np.abs(rank_displacement), axis=0)
max_abs_disp = np.max(np.abs(rank_displacement), axis=0)

# T09: LOO summary with Kendall
lines = []
lines.append("%␣Table␣T9:␣LOO␣sensitivity␣summary")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Leave-one-axis-out␣(LOO)␣rank␣sensitivity.␣␣Spearman␣$\\rho$␣and␣Kendall")
lines.append("$\\tau$␣measure␣rank-order␣agreement␣between␣the␣baseline␣and␣each␣LOO␣ranking.")
lines.append("Mean␣$|\\Delta|$␣and␣Max␣$|\\Delta|$␣report␣the␣average␣and␣maximum␣absolute␣rank")
lines.append("displacement␣across␣the␣27␣countries.}")
lines.append("\\label{tab:loo-summary}")
lines.append("\\begin{tabular}{@{}l␣S[table-format=1.6]␣S[table-format=1.6]␣S[table-format=1.2]␣r@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Excluded␣Axis}}␣&␣{$\\boldsymbol{\\rho}$␣\\textbf{(Spearman)}}␣&␣{$\\boldsymbol{\\tau}$␣\\textbf{(Kendall)}}␣&␣{\\textbf{Mean}␣$|\\Delta|$}␣&␣{\\textbf{Max}␣$|\\Delta|$}␣\\\\")
lines.append("\\midrule")
for j in range(6):
    lines.append(f"␣␣{AXIS_NAMES[j]}␣&␣{loo_spearman[j]:.6f}␣&␣{loo_kendall[j]:.6f}␣&␣{mean_abs_disp[j]:.2f}␣&␣{int(max_abs_disp[j])}␣\\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T09_loo_summary.tex", "\n".join(lines))

# TE2: Full rank displacement matrix
lines = []
lines.append("%␣Table␣TE2:␣Full␣rank␣displacement␣matrix␣(signed)")
lines.append("\\begin{landscape}")
lines.append("\\begin{longtable}{@{}rl␣rrrrrr␣r@{}}")
lines.append("\\caption{Signed␣rank␣displacement␣matrix:␣baseline␣rank␣minus␣LOO␣rank␣for␣each␣axis.")
lines.append("Positive␣values␣indicate␣the␣country␣ranks␣higher␣(more␣concentrated)␣when␣the␣axis␣is␣removed;")
lines.append("negative␣values␣indicate␣it␣ranks␣lower.}")
lines.append("\\label{tab:rank-displacement}\\\\")
lines.append("\\toprule")
hdr = "{\\textbf{Rk}}␣&␣{\\textbf{Country}}"
for j in range(6):
    short = AXIS_NAMES[j][:4]
    hdr += f"␣&␣{{$\\Delta(-${short})}}"
hdr += "␣&␣{$|\\Delta|_{\\max}$}␣\\\\"
lines.append(hdr)
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Rank␣displacement␣matrix␣(continued).}\\\\")
lines.append("\\toprule")
lines.append(hdr)
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    row = f"␣␣{int(baseline_ranks[i])}␣&␣{NAMES[i]}"
```

```
    for j in range(6):
        row += f"␣&␣{rank_displacement[i,j]:+d}"
    row += f"␣&␣{int(max_rank_change[i])}␣\\\\"
    lines.append(row)
lines.append("\\end{longtable}")
lines.append("\\end{landscape}")
write_table("TE2_rank_displacement.tex", "\n".join(lines))

# TB1: LOO detailed ranks
lines = []
lines.append("%␣Appendix␣Table:␣LOO␣detailed␣ranks")
lines.append("\\begin{longtable}{@{}rl" + "r"*6 + "r@{}}")
lines.append("\\caption{Detailed␣rank␣comparison:␣baseline␣vs.\\␣each␣leave-one-axis-out␣variant.")
lines.append("Columns␣show␣the␣rank␣under␣each␣LOO␣composite.␣␣$\\Delta_{\\max}$␣is␣the")
lines.append("maximum␣absolute␣rank␣change␣across␣all␣six␣LOO␣variants.}")
lines.append("\\label{tab:loo-detail}\\\\")
lines.append("\\toprule")
hdr = "{\\textbf{Rk}}␣&␣{\\textbf{Country}}"
for j in range(6):
    short = AXIS_SLUGS[j][:4].capitalize()
    hdr += f"␣&␣{{\\textbf{{$-${short}}}}}"
hdr += "␣&␣{$\\boldsymbol{\\Delta_{\\max}}$}␣\\\\"
lines.append(hdr)
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{LOO␣detailed␣ranks␣(continued).}\\\\")
lines.append("\\toprule")
lines.append(hdr)
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
for i in range(N):
    row = f"␣␣{int(baseline_ranks[i])}␣&␣{NAMES[i]}"
    for j in range(6):
        row += f"␣&␣{int(loo_ranks[i,j])}"
    row += f"␣&␣{int(max_rank_change[i])}␣\\\\"
    lines.append(row)
lines.append("\\end{longtable}")
write_table("TB1_loo_detail.tex", "\n".join(lines))

print("␣␣Most␣axis-sensitive␣countries:")
sens_order = np.argsort(-max_rank_change)
for idx in sens_order[:10]:
    print(f"␣␣␣␣{NAMES[idx]:15s}:␣max_displacement={int(max_rank_change[idx])}")


# ========================================================================
# 4. Z-SCORE STANDARDIZED COMPOSITE
# ========================================================================
print("\n" + "="*72)
print("4.␣Z-SCORE␣STANDARDIZED␣COMPOSITE")
print("="*72)

axis_means = np.mean(AXES, axis=0)
axis_stds_pop = np.std(AXES, axis=0, ddof=0)
Z = (AXES - axis_means) / axis_stds_pop
z_composite = np.mean(Z, axis=1)
z_ranks = rankdata(z_composite)
z_spearman = spearman_corr(baseline_ranks, z_ranks)
z_kendall = kendall_tau(baseline_ranks, z_ranks)
print(f"␣␣Spearman␣rho:␣{z_spearman:.6f}")
print(f"␣␣Kendall␣tau:␣␣{z_kendall:.6f}")

lines = []
lines.append("%␣Table␣T10:␣Z-score␣standardized␣composite␣rank␣comparison")
lines.append("\\begin{longtable}{@{}rlS[table-format=1.8]rS[table-format=2.8]rr@{}}")
lines.append("\\caption{Rank␣comparison:␣baseline␣ISI␣vs.\\␣z-score␣standardized␣composite.")
lines.append("The␣z-score␣composite␣re-centres␣each␣axis␣to␣mean␣zero␣and␣unit␣population")
lines.append("standard␣deviation␣before␣averaging.␣␣Spearman␣$\\rho␣=␣" + f"{z_spearman:.6f}" + "$,")
lines.append("Kendall␣$\\tau␣=␣" + f"{z_kendall:.6f}" + "$.}")
lines.append("\\label{tab:zscore-ranks}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Baseline}}␣&␣{\\textbf{Rk␣(base)}}␣&␣{\\textbf{Z-Composite}}␣&␣{\\textbf{Rk␣(z)}}␣&␣{$\\boldsymbol{\\Delta}$}␣\\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Z-score␣rank␣comparison␣(continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Baseline}}␣&␣{\\textbf{Rk␣(base)}}␣&␣{\\textbf{Z-Composite}}␣&␣{\\textbf{Rk␣(z)}}␣&␣{$\\boldsymbol{\\Delta}$}␣\\\\")
lines.append("\\midrule")
lines.append("\\endhead")
```

```python
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    delta = int(abs(baseline_ranks[i] - z_ranks[i]))
    lines.append(f"␣␣{int(baseline_ranks[i])}␣&␣{NAMES[i]}␣&␣{COMP[i]:.8f}␣&␣{int(baseline_ranks[i])}␣&␣{z_composite[i]:.8f}␣&␣{int(z_ranks[i])}␣&␣{delta}␣\\\\")
lines.append("\\end{longtable}")
write_table("T10_zscore_ranks.tex", "\n".join(lines))


# ========================================================================
# 5. DIRICHLET WEIGHT PERTURBATION (MONTE CARLO)
# ========================================================================
print("\n" + "="*72)
print("5.␣DIRICHLET␣WEIGHT␣PERTURBATION")
print("="*72)

np.random.seed(20240222)
N_MC = 10000
alpha = np.ones(6) * 10.0

mc_ranks = np.zeros((N, N_MC))
for t in range(N_MC):
    w = np.random.dirichlet(alpha)
    mc_comp = AXES @ w
    mc_ranks[:, t] = rankdata(mc_comp)

mean_rank = np.mean(mc_ranks, axis=1)
std_rank = np.std(mc_ranks, axis=1, ddof=0)
p5_rank = np.percentile(mc_ranks, 5, axis=1)
p95_rank = np.percentile(mc_ranks, 95, axis=1)
prob_top5 = np.mean(mc_ranks <= 5, axis=1) * 100.0
prob_top10 = np.mean(mc_ranks <= 10, axis=1) * 100.0

lines = []
lines.append("%␣Table␣T11:␣Dirichlet␣weight␣perturbation␣rank␣volatility")
lines.append("\\begin{longtable}{@{}rl␣S[table-format=2.2]␣S[table-format=1.2]␣S[table-format=2.1]␣S[table-format=2.1]␣S[table-format=3.1]␣S[table-format=3.1]@{}}")
lines.append("\\caption{Rank␣volatility␣under␣Dirichlet␣weight␣perturbation␣($\\alpha_j=10$,␣$N_{\\text{MC}}=" + str(N_MC) + "$).")
lines.append("Mean␣rank,␣standard␣deviation,␣5th--95th␣percentile␣band,␣and␣probability␣of")
lines.append("falling␣in␣the␣top~5␣or␣top~10␣are␣reported␣for␣each␣country.}")
lines.append("\\label{tab:dirichlet-volatility}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Mean␣Rk}}␣&␣{$\\boldsymbol{\\sigma}$}␣&␣{\\textbf{P5}}␣&␣{\\textbf{P95}}␣&␣{\\textbf{\\%Top5}}␣&␣{\\textbf{\\%Top10}}␣\\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Dirichlet␣rank␣volatility␣(continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Mean␣Rk}}␣&␣{$\\boldsymbol{\\sigma}$}␣&␣{\\textbf{P5}}␣&␣{\\textbf{P95}}␣&␣{\\textbf{\\%Top5}}␣&␣{\\textbf{\\%Top10}}␣\\\\")
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    lines.append(f"␣␣{int(baseline_ranks[i])}␣&␣{NAMES[i]}␣&␣{mean_rank[i]:.2f}␣&␣{std_rank[i]:.2f}␣&␣{p5_rank[i]:.1f}␣&␣{p95_rank[i]:.1f}␣&␣{prob_top5[i]:.1f}␣&␣{prob_top10[i]:.1f}␣\\\\")
lines.append("\\end{longtable}")
write_table("T11_dirichlet_volatility.tex", "\n".join(lines))

mc_spearmans = np.array([spearman_corr(baseline_ranks, mc_ranks[:, t]) for t in range(N_MC)])
print(f"␣␣Mean␣Spearman:␣{np.mean(mc_spearmans):.6f}␣(std={np.std(mc_spearmans):.6f})")
print(f"␣␣Min:␣{np.min(mc_spearmans):.6f}␣␣Max:␣{np.max(mc_spearmans):.6f}")

vol_order = np.argsort(-std_rank)
for idx in vol_order[:5]:
    print(f"␣␣␣␣{NAMES[idx]:15s}:␣std={std_rank[idx]:.2f}␣band=[{p5_rank[idx]:.0f},{p95_rank[idx]:.0f}]")


# ========================================================================
# 6. WINSORIZATION
# ========================================================================
print("\n" + "="*72)
print("6.␣WINSORIZATION")
print("="*72)

AXES_W = AXES.copy()
for j_cap in [3, 5]:
```

```python
    p95 = np.percentile(AXES[:, j_cap], 95)
    print(f"  {AXIS_NAMES[j_cap]} P95 cap: {p95:.8f}")
    AXES_W[:, j_cap] = np.minimum(AXES[:, j_cap], p95)

w_composite = np.mean(AXES_W, axis=1)
w_ranks = rankdata(w_composite)
w_spearman = spearman_corr(baseline_ranks, w_ranks)
print(f"  Spearman rho: {w_spearman:.6f}")

lines = []
lines.append("% Table T12: Winsorization sensitivity")
lines.append("\\begin{longtable}{@{}rlS[table-format=1.8]rS[table-format=1.8]rr@{}}")
lines.append("\\caption{Rank comparison: baseline vs.\\ winsorized composite (Axis~4 and Axis~6")
lines.append("capped at their respective 95th percentiles).  Spearman $\\rho = " + f"{w_spearman:.6f}" + "$.}")
lines.append("\\label{tab:winsorization}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Baseline}} & {\\textbf{Rk (base)}} & {\\textbf{Winsorized}} & {\\textbf{Rk (win)}} & {$\\boldsymbol{\\Delta}$} \\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Winsorization (continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Baseline}} & {\\textbf{Rk (base)}} & {\\textbf{Winsorized}} & {\\textbf{Rk (win)}} & {$\\boldsymbol{\\Delta}$} \\\\")
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    delta = int(abs(baseline_ranks[i] - w_ranks[i]))
    lines.append(f"  {int(baseline_ranks[i])} & {NAMES[i]} & {COMP[i]:.8f} & {int(baseline_ranks[i])} & {w_composite[i]:.8f} & {int(w_ranks[i])} & {delta} \\\\")
lines.append("\\end{longtable}")
write_table("T12_winsorization.tex", "\n".join(lines))


# =========================================================================
# 7. TOP-5 / BOTTOM-5 PER AXIS
# =========================================================================
print("\n" + "="*72)
print("7. TOP-5 / BOTTOM-5 PER AXIS")
print("="*72)

lines = []
lines.append("% Table T6: Top-5 and Bottom-5 per axis")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Top-5 (most concentrated) and Bottom-5 (least concentrated) countries per axis.}")
lines.append("\\label{tab:axis-top-bottom}")
lines.append("\\footnotesize")
lines.append("\\begin{tabular}{@{}lllll@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Axis}} & {\\textbf{Top-5 (most conc.)}} & {\\textbf{Score}} & {\\textbf{Bottom-5 (least conc.)}} & {\\textbf{Score}} \\\\")
lines.append("\\midrule")
for j in range(6):
    sorted_idx = np.argsort(-AXES[:, j])
    top5 = sorted_idx[:5]
    bot5 = sorted_idx[-5:][::-1]
    for k in range(5):
        prefix = AXIS_NAMES[j] if k == 0 else ""
        ti = top5[k]
        bi = bot5[k]
        lines.append(f"  {prefix} & {NAMES[ti]} & {AXES[ti, j]:.8f} & {NAMES[bi]} & {AXES[bi, j]:.8f} \\\\")
    if j < 5:
        lines.append("\\addlinespace")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T06_axis_top_bottom.tex", "\n".join(lines))


# =========================================================================
# 8. LORENZ CURVE COORDINATES
# =========================================================================
print("\n" + "="*72)
print("8. LORENZ CURVE")
print("="*72)
```

```python
sorted_comp = np.sort(COMP)
cum_share = np.cumsum(sorted_comp) / np.sum(sorted_comp)
pop_share = np.arange(1, N+1) / N
gini = 1.0 - 2.0 * np.trapezoid(cum_share, pop_share)
print(f"␣␣Gini:␣{gini:.6f}")

lines = []
lines.append("%␣Lorenz␣curve␣coordinates")
lines.append("%␣population_share␣cumulative_score_share")
lines.append("0.0000␣0.0000")
for k in range(N):
    lines.append(f"{pop_share[k]:.6f}␣{cum_share[k]:.6f}")
write_table("lorenz_coords.dat", "\n".join(lines))


# ========================================================================
# 9. AXIS RANK TABLES
# ========================================================================
print("\n" + "="*72)
print("9.␣AXIS␣RANK␣TABLES")
print("="*72)

for j in range(6):
    ax_ranks = rankdata(AXES[:, j])
    sorted_idx = np.argsort(ax_ranks)
    lines = []
    lines.append(f"%␣Appendix:␣{AXIS_NAMES[j]}␣axis␣ranking")
    lines.append("\\begin{longtable}{@{}rlS[table-format=1.8]@{}}")
    lines.append(f"\\caption{{{AXIS_NAMES[j]}␣axis␣ranking,␣EU-27,␣2024.}}")
    lines.append(f"\\label{{tab:axis-rank-{AXIS_SLUGS[j]}}}\\\\")
    lines.append("\\toprule")
    lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Score}}␣\\\\")
    lines.append("\\midrule")
    lines.append("\\endfirsthead")
    lines.append(f"\\caption[]{{{AXIS_NAMES[j]}␣axis␣ranking␣(continued).}}\\\\")
    lines.append("\\toprule")
    lines.append("{\\textbf{Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{Score}}␣\\\\")
    lines.append("\\midrule")
    lines.append("\\endhead")
    lines.append("\\bottomrule")
    lines.append("\\endlastfoot")
    for i in sorted_idx:
        lines.append(f"␣␣{int(ax_ranks[i])}␣&␣{NAMES[i]}␣&␣{AXES[i,j]:.8f}␣\\\\")
    lines.append("\\end{longtable}")
    write_table(f"TA_{AXIS_SLUGS[j]}_ranking.tex", "\n".join(lines))
    print(f"␣␣Written:␣{AXIS_NAMES[j]}")


# ========================================================================
# 10. HISTOGRAM / ECDF COORDINATES
# ========================================================================
print("\n" + "="*72)
print("10.␣HISTOGRAM/ECDF")
print("="*72)

sorted_comp_ecdf = np.sort(COMP)
lines = []
lines.append("%␣ECDF␣coordinates")
lines.append("%␣composite_score␣ecdf_value")
for k in range(N):
    lines.append(f"{sorted_comp_ecdf[k]:.8f}␣{(k+1)/N:.8f}")
write_table("ecdf_coords.dat", "\n".join(lines))

bin_edges = np.linspace(0.20, 0.55, 8)
hist_counts, _ = np.histogram(COMP, bins=bin_edges)
lines = []
lines.append("%␣Histogram␣coordinates␣(bin_center,␣count)")
for k in range(len(hist_counts)):
    center = (bin_edges[k] + bin_edges[k+1]) / 2.0
    lines.append(f"{center:.4f}␣{hist_counts[k]}")
write_table("histogram_coords.dat", "\n".join(lines))


# ========================================================================
# 11. RADAR PLOT DATA
# ========================================================================
print("\n" + "="*72)
print("11.␣RADAR␣PLOT␣DATA")
print("="*72)

selected = ["MT", "CY", "IE", "NL", "DE", "SK", "FI", "FR"]
for iso in selected:
    idx = ISO.index(iso)
```

```
    lines = []
    lines.append(f"% Radar data for {NAMES[idx]} ({iso})")
    lines.append("% axis_index axis_name score")
    for j in range(6):
        lines.append(f"{j} {AXIS_SLUGS[j]} {AXES[idx,j]:.8f}")
    write_table(f"radar_{iso}.dat", "\n".join(lines))
    print(f"  {NAMES[idx]}")


# ==========================================================================
# 12. BOXPLOT STATS
# ==========================================================================
print("\n" + "="*72)
print("12. BOXPLOT STATISTICS")
print("="*72)

for j in range(6):
    vals = AXES[:, j]
    q1 = np.percentile(vals, 25)
    med = np.median(vals)
    q3 = np.percentile(vals, 75)
    iqr = q3 - q1
    wlo = np.min(vals[vals >= q1 - 1.5*iqr])
    whi = np.max(vals[vals <= q3 + 1.5*iqr])
    print(f"  {AXIS_NAMES[j]:20s}: Q1={q1:.4f} Med={med:.4f} Q3={q3:.4f} Lo={wlo:.4f} Hi={whi:.4f}")


# ==========================================================================
# 13. RANK ELASTICITY TO AXIS PERTURBATION
# ==========================================================================
print("\n" + "="*72)
print("13. RANK ELASTICITY")
print("="*72)

epsilons = [0.05, 0.10, 0.20]  # 5%, 10%, 20% shocks
rank_sens_by_eps = {}

for epsilon in epsilons:
    rank_sens_pos = np.zeros((N, 6))
    for j in range(6):
        Ap = AXES.copy()
        Ap[:, j] = np.clip(AXES[:, j] * (1 + epsilon), 0.0, 1.0)
        rp = rankdata(np.mean(Ap, axis=1))
        rank_sens_pos[:, j] = baseline_ranks - rp
    rank_sens_by_eps[epsilon] = rank_sens_pos

# Use 10% shock for the main table
rank_sens_pos = rank_sens_by_eps[0.10]
total_sens = np.sum(np.abs(rank_sens_pos), axis=1)
fragile_order = np.argsort(-total_sens)

print("  Top 10 rank-fragile (sum |delta| from +1%):")
for idx in fragile_order[:10]:
    print(f"    {NAMES[idx]:15s}: total={total_sens[idx]:.0f}  per_axis={rank_sens_pos[idx,:]}")

locked = np.sum(total_sens == 0)
print(f"  Structurally locked-in countries: {locked}")
for idx in range(N):
    if total_sens[idx] == 0:
        print(f"    {NAMES[idx]}")

# T13: Rank elasticity table
lines = []
lines.append("% Table T13: Rank sensitivity to +10% axis perturbation")
lines.append("\\begin{landscape}")
lines.append("\\begin{longtable}{@{}rl rrrrrr r@{}}")
lines.append("\\caption{Rank change per country under a $+10\\%$ multiplicative perturbation to each axis.")
lines.append("Positive values indicate rank improvement (toward rank 1);")
lines.append("negative values indicate deterioration.  $\\Sigma|\\Delta|$ is the total absolute displacement.}")
lines.append("\\label{tab:rank-elasticity}\\\\")
lines.append("\\toprule")
hdr = "{\\textbf{Rk}} & {\\textbf{Country}}"
for j in range(6):
    hdr += f" & {{\\textbf{{{AXIS_NAMES[j][:4]}}}}}"
hdr += " & {$\\Sigma|\\Delta|$} \\\\"
lines.append(hdr)
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Rank elasticity (continued).}\\\\")
lines.append("\\toprule")
lines.append(hdr)
```

```python
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    row = f"  {int(baseline_ranks[i])} & {NAMES[i]}"
    for j in range(6):
        row += f" & {int(rank_sens_pos[i,j]):+d}"
    row += f" & {int(total_sens[i])} \\\\"
    lines.append(row)
lines.append("\\end{longtable}")
lines.append("\\end{landscape}")
write_table("T13_rank_elasticity.tex", "\n".join(lines))


# ==========================================================================
# 14. AXIS CONTRIBUTION TO RANK DIFFERENTIATION
# ==========================================================================
print("\n" + "="*72)
print("14. AXIS-RANK CONTRIBUTION")
print("="*72)

axis_rank_corr = np.zeros(6)
axis_rank_r2 = np.zeros(6)
for j in range(6):
    r = pearson_corr(AXES[:, j], -baseline_ranks)
    axis_rank_corr[j] = r
    axis_rank_r2[j] = r**2
    print(f"  {AXIS_NAMES[j]:20s}: r={r:.6f}  R2={r**2:.6f}")

partial_corr = np.zeros(6)
for j in range(6):
    other = [k for k in range(6) if k != j]
    X_o = np.column_stack([np.ones(N), AXES[:, other]])
    y = -baseline_ranks.astype(float)
    x_t = AXES[:, j]
    beta_y = np.linalg.lstsq(X_o, y, rcond=None)[0]
    beta_x = np.linalg.lstsq(X_o, x_t, rcond=None)[0]
    resid_y = y - X_o @ beta_y
    resid_x = x_t - X_o @ beta_x
    partial_corr[j] = pearson_corr(resid_x, resid_y)
    print(f"  {AXIS_NAMES[j]:20s}: partial_r={partial_corr[j]:.6f}")

lines = []
lines.append("% Table T14: Axis contribution to rank differentiation")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Axis contribution to composite rank differentiation.  Bivariate $R^2$:")
lines.append("share of rank variance explained by each axis alone.  Partial $r$: association with")
lines.append("rank after controlling for all other axes.}")
lines.append("\\label{tab:axis-rank-contrib}")
lines.append("\\begin{tabular}{@{}l S[table-format=+1.4] S[table-format=1.4] S[table-format=+1.4]@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Axis}} & {\\textbf{Pearson $r$}} & {$\\boldsymbol{R^2}$} & {\\textbf{Partial $r$}} \\\\")
lines.append("\\midrule")
for j in range(6):
    lines.append(f"  {AXIS_NAMES[j]} & {axis_rank_corr[j]:+.4f} & {axis_rank_r2[j]:.4f} & {partial_corr[j]:+.4f} \\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T14_axis_rank_contrib.tex", "\n".join(lines))


# ==========================================================================
# 15. CLASSIFICATION COMPRESSION DIAGNOSTICS
# ==========================================================================
print("\n" + "="*72)
print("15. COMPRESSION DIAGNOSTICS")
print("="*72)

tiers = {
    "Highly concentrated": (0.50, 1.01),
    "Moderately concentrated": (0.25, 0.50),
    "Mildly concentrated": (0.15, 0.25),
    "Unconcentrated": (0.00, 0.15),
}

grand_mean = np.mean(COMP)
ssw = 0.0
ssb = 0.0
```

```python
tier_info = {}

for tname, (lo, hi) in tiers.items():
    mask = (COMP >= lo) & (COMP < hi)
    if np.sum(mask) == 0:
        continue
    vals = COMP[mask]
    tmean = np.mean(vals)
    tvar = np.var(vals, ddof=0) if len(vals) > 1 else 0.0
    tier_info[tname] = (len(vals), tmean, tvar)
    ssw += np.sum((vals - tmean)**2)
    ssb += len(vals) * (tmean - grand_mean)**2

total_ss = np.sum((COMP - grand_mean)**2)
eta_sq = ssb / total_ss if total_ss > 0 else 0.0

print(f"␣␣Grand␣mean:␣{grand_mean:.8f}")
print(f"␣␣SS_within={ssw:.10f}␣␣SS_between={ssb:.10f}␣␣SS_total={total_ss:.10f}")
print(f"␣␣Eta-squared:␣{eta_sq:.6f}")
print(f"␣␣Gini:␣{gini:.6f}")

for tname, (n, m, v) in tier_info.items():
    print(f"␣␣{tname:30s}:␣n={n}␣␣mean={m:.8f}␣␣var={v:.10f}")

# Tercile and quartile alternatives
tercile_edges = np.percentile(COMP, [33.33, 66.67])
quartile_edges = np.percentile(COMP, [25, 50, 75])
print(f"␣␣Terciles:␣{tercile_edges}")
print(f"␣␣Quartiles:␣{quartile_edges}")

# T15: Compression diagnostics
lines = []
lines.append("%␣Table␣T15:␣Classification␣compression␣diagnostics")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Classification␣compression␣diagnostics.␣␣The␣official␣four-tier␣scheme")
lines.append("is␣compared␣with␣data-driven␣tercile␣and␣quartile␣alternatives.␣␣$\\eta^2␣=␣SS_B␣/␣SS_T$")
lines.append("measures␣between-tier␣variance␣share.}")
lines.append("\\label{tab:compression-diagnostics}")
lines.append("\\small")
lines.append("\\begin{tabular}{@{}l␣r␣S[table-format=1.8]␣S[table-format=1.10]␣l@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Tier}}␣&␣{$n$}␣&␣{\\textbf{Mean}}␣&␣{\\textbf{Var␣(within)}}␣&␣{\\textbf{Scheme}}␣\\\\")
lines.append("\\midrule")
for tname in ["Highly␣concentrated", "Moderately␣concentrated", "Mildly␣concentrated"]:
    if tname in tier_info:
        n, m, v = tier_info[tname]
        lines.append(f"␣␣{tname}␣&␣{n}␣&␣{m:.8f}␣&␣{v:.10f}␣&␣Official␣\\\\")
lines.append("\\addlinespace")

# Terciles
trc_bounds = [(0, tercile_edges[0]), (tercile_edges[0], tercile_edges[1]), (tercile_edges[1], 1.0)]
trc_labels = ["T1␣(low)", "T2␣(mid)", "T3␣(high)"]
for idx, (lo, hi) in enumerate(trc_bounds):
    mask = (COMP >= lo) & (COMP < hi) if idx < 2 else (COMP >= lo)
    vals = COMP[mask]
    lines.append(f"␣␣{trc_labels[idx]}␣&␣{len(vals)}␣&␣{np.mean(vals):.8f}␣&␣{np.var(vals,␣ddof=0):.10f}␣&␣Tercile␣\\\\")
lines.append("\\addlinespace")

# Quartiles
qrt_bounds = [(0, quartile_edges[0]), (quartile_edges[0], quartile_edges[1]),
              (quartile_edges[1], quartile_edges[2]), (quartile_edges[2], 1.0)]
qrt_labels = ["Q1␣(low)", "Q2", "Q3", "Q4␣(high)"]
for idx, (lo, hi) in enumerate(qrt_bounds):
    mask = (COMP >= lo) & (COMP < hi) if idx < 3 else (COMP >= lo)
    vals = COMP[mask]
    lines.append(f"␣␣{qrt_labels[idx]}␣&␣{len(vals)}␣&␣{np.mean(vals):.8f}␣&␣{np.var(vals,␣ddof=0):.10f}␣&␣Quartile␣\\\\")
lines.append("\\midrule")
lines.append(f"␣␣\\multicolumn{{5}}{{@{{}}l@{{}}}}{{$\\eta^2$␣(official)␣$=␣{eta_sq:.6f}$;␣Gini␣$=␣{gini:.6f}$}}␣\\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T15_compression_diagnostics.tex", "\n".join(lines))


# ============================================================================
# 16. DEFENCE AXIS STRUCTURAL ANALYSIS
# ============================================================================
```

```
print("\n" + "="*72)
print("16. DEFENCE STRUCTURAL ANALYSIS")
print("="*72)

defence = AXES[:, 3]
defence_is_max = int(np.sum(max_axis_idx == 3))
near_monopoly = int(np.sum(defence > 0.80))
sk_idx = ISO.index("SK")

# Gap: defence minus second-highest (for defence-max countries)
# or defence minus max (for logistics-max countries)
defence_gap = np.zeros(N)
for i in range(N):
    sorted_vals = np.sort(AXES[i])[::-1]
    if max_axis_idx[i] == 3:
        defence_gap[i] = defence[i] - sorted_vals[1]
    else:
        defence_gap[i] = defence[i] - AXES[i, max_axis_idx[i]]

print(f"  Defence is max axis: {defence_is_max}/27 ({defence_is_max/N*100:.1f}%)")
print(f"  Score > 0.80: {near_monopoly}")
print(f"  Score = 0.00: {int(np.sum(defence == 0))}")
print(f"  Mean gap (def-max countries): {np.mean(defence_gap[max_axis_idx==3]):.6f}")

# Skewness and kurtosis
def_mean = np.mean(defence)
def_std = np.std(defence, ddof=0)
def_skew = float(np.mean(((defence - def_mean)/def_std)**3))
def_kurt = float(np.mean(((defence - def_mean)/def_std)**4) - 3)

# T16: Defence structural
lines = []
lines.append("% Table T16: Defence axis structural analysis")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Structural properties of the defence axis distribution, EU-27, 2024.}")
lines.append("\\label{tab:defence-structural}")
lines.append("\\begin{tabular}{@{}lr@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Statistic}} & {\\textbf{Value}} \\\\")
lines.append("\\midrule")
lines.append(f"  Countries where defence is max axis & {defence_is_max} of 27 ({defence_is_max/N*100:.1f}\\%) \\\\")
lines.append(f"  Countries with score $> 0.80$ & {near_monopoly} ({near_monopoly/N*100:.1f}\\%) \\\\")
lines.append(f"  Countries with score $= 0.00$ & {int(np.sum(defence==0))} \\\\")
lines.append(f"  Mean (pop.) & {def_mean:.8f} \\\\")
lines.append(f"  Median & {np.median(defence):.8f} \\\\")
lines.append(f"  Std.\\ dev.\\ (pop.) & {def_std:.8f} \\\\")
lines.append(f"  Range & {np.max(defence) - np.min(defence):.8f} \\\\")
lines.append(f"  IQR & {np.percentile(defence, 75) - np.percentile(defence, 25):.8f} \\\\")
lines.append(f"  Skewness & {def_skew:.4f} \\\\")
lines.append(f"  Excess kurtosis & {def_kurt:.4f} \\\\")
lines.append(f"  Mean def--2nd gap (def-max) & {np.mean(defence_gap[max_axis_idx==3]):.8f} \\\\")
lines.append(f"  Median def--2nd gap (def-max) & {np.median(defence_gap[max_axis_idx==3]):.8f} \\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T16_defence_structural.tex", "\n".join(lines))

# T16b: Defence gap per country
lines = []
lines.append("% Table T16b: Defence gap per country")
lines.append("\\begin{longtable}{@{}rl S[table-format=1.8] l S[table-format=+1.8]@{}}")
lines.append("\\caption{Defence axis score, max axis identity, and defence-to-max gap per country.}")
lines.append("\\label{tab:defence-gap}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Def.\\ Score}} & {\\textbf{Max Axis}} & {\\textbf{Gap}} \\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Defence gap (continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Def.\\ Score}} & {\\textbf{Max Axis}} & {\\textbf{Gap}} \\\\")
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    lines.append(f"  {int(baseline_ranks[i])} & {NAMES[i]} & {defence[i]:.8f} & {AXIS_NAMES[max_axis_idx[i]]} & {defence_gap[i]:+.8f} \\\\")
lines.append("\\end{longtable}")
```

```python
write_table("T16b_defence_gap.tex", "\n".join(lines))


# ========================================================================
# 17. COMPENSABILITY ANALYSIS
# ========================================================================
print("\n" + "="*72)
print("17. COMPENSABILITY ANALYSIS")
print("="*72)

max_scores = np.max(AXES, axis=1)
min_scores = np.min(AXES, axis=1)
axis_range = max_scores - min_scores
compensation_gap = max_scores - COMP

country_cv = np.zeros(N)
for i in range(N):
    m = np.mean(AXES[i])
    s = np.std(AXES[i], ddof=0)
    country_cv[i] = s / m if m > 0 else 0

comp_order = np.argsort(-compensation_gap)
print("  Most compensated (largest max-composite gap):")
for idx in comp_order[:5]:
    print(f"    {NAMES[idx]:15s}: max={max_scores[idx]:.4f} comp={COMP[idx]:.4f} gap={compensation_gap[idx]:.4f} CV={country_cv[idx]:.4f}")

# T17: Compensability
lines = []
lines.append("% Table T17: Compensability analysis")
lines.append("\\begin{longtable}{@{}rl S[table-format=1.8] S[table-format=1.8] S[table-format=1.8] S[table-format=1.8] S[table-format=1.4]@{}}")
lines.append("\\caption{Compensability analysis.  Gap $=$ Max $-$ Composite.  CV: coefficient of")
lines.append("variation of the six axis scores per country.  High Gap and CV indicate strong")
lines.append("internal compensation where low axes offset high axes.}")
lines.append("\\label{tab:compensability}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Max}} & {\\textbf{Min}} & {\\textbf{Range}} & {\\textbf{Gap}} & {\\textbf{CV}} \\\\")
lines.append("\\midrule")
lines.append("\\endfirsthead")
lines.append("\\caption[]{Compensability (continued).}\\\\")
lines.append("\\toprule")
lines.append("{\\textbf{Rk}} & {\\textbf{Country}} & {\\textbf{Max}} & {\\textbf{Min}} & {\\textbf{Range}} & {\\textbf{Gap}} & {\\textbf{CV}} \\\\")
lines.append("\\midrule")
lines.append("\\endhead")
lines.append("\\bottomrule")
lines.append("\\endlastfoot")
order = np.argsort(baseline_ranks)
for i in order:
    lines.append(f"  {int(baseline_ranks[i])} & {NAMES[i]} & {max_scores[i]:.8f} & {min_scores[i]:.8f} & {axis_range[i]:.8f} & {compensation_gap[i]:.8f} & {country_cv[i]:.4f} \\\\")
lines.append("\\end{longtable}")
write_table("T17_compensability.tex", "\n".join(lines))


# ========================================================================
# 18. GEOGRAPHIC / STRUCTURAL CORRELATES
# ========================================================================
print("\n" + "="*72)
print("18. GEOGRAPHIC CORRELATES")
print("="*72)

ISLAND = {"MT", "CY", "IE"}
SMALL_POP = {"MT", "CY", "LU", "EE", "LV", "LT", "SI", "HR"}
LARGE_POP = {"DE", "FR", "ES", "IT", "PL", "RO", "NL"}
CORE = {"DE", "FR", "BE", "NL", "LU", "AT"}
PERIPHERY = {"MT", "CY", "EL", "PT", "FI", "EE", "LV", "LT", "BG", "RO", "HR", "IE"}

non_island = set(ISO) - ISLAND
mid_pop = set(ISO) - SMALL_POP - LARGE_POP

groups = [
    ("Island/near-island", ISLAND),
    ("Continental", non_island),
    ("Small ($<$5M pop.)", SMALL_POP),
    ("Medium (5--20M)", mid_pop),
    ("Large ($>$20M pop.)", LARGE_POP),
    ("Core EU-6", CORE),
    ("Periphery", PERIPHERY),
]

```

```python
lines = []
lines.append("%␣Table␣T18:␣Geographic␣and␣structural␣correlates")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{ISI␣composite␣by␣geographic␣and␣structural␣category.}")
lines.append("\\label{tab:geo-correlates}")
lines.append("\\begin{tabular}{@{}l␣r␣S[table-format=1.6]␣S[table-format=1.6]␣S[table-format=1.6]@{}}")
lines.append("\\toprule")
lines.append("{\\textbf{Category}}␣&␣{$n$}␣&␣{\\textbf{Mean}}␣&␣{\\textbf{Std␣(pop.)}}␣&␣{\\textbf{Median}}␣\\\\")
lines.append("\\midrule")
for glabel, gset in groups:
    mask = np.array([ISO[i] in gset for i in range(N)])
    vals = COMP[mask]
    n = len(vals)
    m = np.mean(vals)
    s = np.std(vals, ddof=0)
    med = np.median(vals)
    lines.append(f"␣␣{glabel}␣&␣{n}␣&␣{m:.6f}␣&␣{s:.6f}␣&␣{med:.6f}␣\\\\")
    print(f"␣␣{glabel:25s}:␣n={n}␣mean={m:.6f}␣std={s:.6f}")
lines.append("\\midrule")
lines.append(f"␣␣EU-27␣(all)␣&␣27␣&␣{np.mean(COMP):.6f}␣&␣{np.std(COMP,␣ddof=0):.6f}␣&␣{np.median(COMP):.6f}␣\\\\")
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("T18_geo_correlates.tex", "\n".join(lines))


# ==========================================================================
# 19. CORRELATION T-STATISTICS
# ==========================================================================
print("\n" + "="*72)
print("19.␣CORRELATION␣T-STATISTICS")
print("="*72)

pearson_matrix = np.corrcoef(AXES.T)
t_matrix = np.zeros((6, 6))
for j in range(6):
    for k in range(6):
        if j == k:
            t_matrix[j, k] = np.inf
        else:
            r = pearson_matrix[j, k]
            denom = 1 - r**2
            t_matrix[j, k] = r * np.sqrt((N - 2) / denom) if denom > 0 else np.inf

for j in range(6):
    for k in range(j+1, 6):
        r = pearson_matrix[j, k]
        t = t_matrix[j, k]
        sig = "*" if abs(t) > 2.060 else ""
        print(f"␣␣␣␣{AXIS_NAMES[j]:15s}␣vs␣{AXIS_NAMES[k]:15s}:␣r={r:+.4f}␣t={t:+.4f}␣{sig}")

# TE3: Correlation t-statistics
lines = []
lines.append("%␣Table␣TE3:␣Correlation␣t-statistics")
lines.append("\\begin{table}[htbp]")
lines.append("\\centering")
lines.append("\\caption{Pearson␣correlations␣(below␣diagonal)␣and␣$t$-statistics␣(above␣diagonal).")
lines.append("$t␣=␣r\\sqrt{(N-2)/(1-r^2)}$,␣$N=27$,␣df$=25$.␣Critical␣value:␣$t_{0.025,25}␣=␣2.060$;")
lines.append("*␣indicates␣significance␣at␣5\\%.}")
lines.append("\\label{tab:corr-tstat}")
lines.append("\\small")
lines.append("\\sisetup{table-format=+1.3}")
lines.append("\\begin{tabular}{@{}l␣S␣S␣S␣S␣S␣S@{}}")
lines.append("\\toprule")
lines.append("␣&␣{\\textbf{Fin.}}␣&␣{\\textbf{Ene.}}␣&␣{\\textbf{Tec.}}␣&␣{\\textbf{Def.}}␣&␣{\\textbf{Crit.}}␣&␣{\\textbf{Log.}}␣\\\\")
lines.append("\\midrule")
for j in range(6):
    row = f"␣␣{AXIS_NAMES[j]}"
    for k in range(6):
        if j == k:
            row += "␣&␣{--}"
        elif j > k:
            r = pearson_matrix[j, k]
            sig = "*" if abs(t_matrix[j, k]) > 2.060 else ""
            row += f"␣&␣{{{r:+.3f}{sig}}}"
        else:
            t = t_matrix[j, k]
            row += f"␣&␣{{{t:+.3f}}}"
    row += "␣\\\\"
```

```
    lines.append(row)
lines.append("\\bottomrule")
lines.append("\\end{tabular}")
lines.append("\\end{table}")
write_table("TE3_corr_tstat.tex", "\n".join(lines))


# ========================================================================
# 20. LOO INDIVIDUAL AXIS TABLES
# ========================================================================
print("\n" + "="*72)
print("20.␣LOO␣INDIVIDUAL␣TABLES")
print("="*72)

for j in range(6):
    order = np.argsort(loo_ranks[:, j])
    lines = []
    lines.append(f"%␣LOO␣excluding␣{AXIS_NAMES[j]}")
    lines.append("\\begin{longtable}{@{}rl␣S[table-format=1.8]␣r␣r@{}}")
    lines.append(f"\\caption{{LOO␣ranking␣excluding␣{AXIS_NAMES[j]}.␣␣Spearman␣$\\rho␣=␣{loo_spearman[j]:.6f}$,␣Kendall␣$\\tau␣=␣{loo_kendall[j]:.6f}$.}}")
    lines.append(f"\\label{{tab:loo-{AXIS_SLUGS[j]}}}\\\\")
    lines.append("\\toprule")
    lines.append("{\\textbf{LOO␣Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{LOO␣Comp.}}␣&␣{\\textbf{Base␣Rk}}␣&␣{$\\boldsymbol{\\Delta}$}␣\\\\")
    lines.append("\\midrule")
    lines.append("\\endfirsthead")
    lines.append(f"\\caption[]{{LOO␣excl.\\␣{AXIS_NAMES[j]}␣(cont.).}}\\\\")
    lines.append("\\toprule")
    lines.append("{\\textbf{LOO␣Rk}}␣&␣{\\textbf{Country}}␣&␣{\\textbf{LOO␣Comp.}}␣&␣{\\textbf{Base␣Rk}}␣&␣{$\\boldsymbol{\\Delta}$}␣\\\\")
    lines.append("\\midrule")
    lines.append("\\endhead")
    lines.append("\\bottomrule")
    lines.append("\\endlastfoot")
    for i in order:
        delta = int(baseline_ranks[i] - loo_ranks[i, j])
        lines.append(f"␣␣{int(loo_ranks[i,j])}␣&␣{NAMES[i]}␣&␣{loo_composites[i,j]:.8f}␣&␣{int(baseline_ranks[i])}␣&␣{delta:+d}␣\\\\")
    lines.append("\\end{longtable}")
    write_table(f"TC_loo_{AXIS_SLUGS[j]}.tex", "\n".join(lines))
    print(f"␣␣Written:␣LOO␣excl␣{AXIS_NAMES[j]}")


# ========================================================================
# 21. MULTIPLE REGRESSION
# ========================================================================
print("\n" + "="*72)
print("21.␣MULTIPLE␣REGRESSION")
print("="*72)

y = -baseline_ranks.astype(float)
X = np.column_stack([np.ones(N), AXES])
beta = np.linalg.lstsq(X, y, rcond=None)[0]
y_hat = X @ beta
ss_res = np.sum((y - y_hat)**2)
ss_tot = np.sum((y - np.mean(y))**2)
r2_full = 1 - ss_res / ss_tot
adj_r2 = 1 - (1 - r2_full) * (N - 1) / (N - 7)
print(f"␣␣R2={r2_full:.6f}␣␣adj_R2={adj_r2:.6f}")
print(f"␣␣Coefficients:␣intercept={beta[0]:.4f},␣" + ",␣".join(f"{AXIS_NAMES[j]}={beta[j+1]:.4f}" for j in range(6)))


# ========================================================================
print("\n" + "="*72)
print("ALL␣COMPUTATIONS␣COMPLETE")
print(f"Output␣written␣to:␣{OUT_DIR}")
print("="*72)
```