Counting Dependencies: Aggregation Rules and What Official Dependency Indices Identify

Stéfane Fermigier

Draft working paper v0.1, 2026-07-18. Three propositions, each with a proof and a numeric cross-check against brute force, plus a Monte-Carlo report on rank stability. The application reconstructs a published index from its own published cells; nothing new is measured.

Abstract

European economic-security policy now rests on official dependency indices. The instrument in current use scores technology areas against a fixed list of dependency dimensions on a two-level ordinal scale, counts the dimensions rated high, and reports the difference between that count and the corresponding count of reverse dependencies. We ask what such an index identifies.

Threshold counting and weakest-link exposure are different functionals and rank the same latent profiles differently, so an index that counts cannot represent a system its own authors describe as failing when one essential component fails. On a two-level scale the weakest-link reading is worse than approximate. Where every unit scores high on some shared dimension, the maximum is constant across units, all ranking content comes from the additive rule, and the choice of rule is doing the work the data are assumed to do. Netting counts of high dependencies against counts of high reverse dependencies is a third distinct object, and because a two-level scale discards intensity, the net count can invert the ordering of bargaining positions it is read as describing.

We reconstruct the index published for the European Commission from its own cells. The reconstruction reproduces twelve of thirteen printed totals; the thirteenth row carries seven high cells and prints six. Applying the index's own stated rule to its own tables contradicts its published prose: the area described as showing net dependency of more than three types scores minus one, a net reverse dependency. Under the strict reading of the accompanying claim, three of the remaining eight areas also fail it. Every one of the thirteen areas scores high on digital infrastructures and software, so the weakest-link reading is degenerate on this data. And under random positive reweighting of the eight dimensions the most-dependent area changes on 49.8% of draws, alternating between two areas. We then state what a path-aware instrument has to supply instead, and connect the requirements to the measured-graph approach developed in the companion papers.

Keywords: composite indicators, economic security, strategic dependency, aggregation, interdependence, measurement.
JEL: C43, F52, L86, O38.

1. Introduction

Policy that aims to reduce strategic dependency needs a measure of it. The European Commission's evidence base for the 2028-2034 programming period contains one: thirteen technology areas scored against eight dependency dimensions on a two-level scale, aggregated by counting the dimensions rated high, and netted against a parallel count of reverse dependencies (European Commission 2026). Investment options are then ranked partly by the result.

This paper asks what that construction identifies. Three results follow from the form of the index alone, before any question about the underlying scores.

Proposition 1 separates two functionals that policy discourse treats as one. Counting dimensions above a threshold and taking the worst dimension are different operations, and they order the same latent exposure profiles differently. A unit with one catastrophic dependency and seven benign ones counts as less dependent than a unit with five moderate ones, whatever the intensity of the catastrophic one.

Proposition 2 sharpens this for coarse scales. On a two-level scale, the weakest-link reading collapses to a single question, whether any dimension is high. Where every unit answers yes, the maximum is constant across units, the ranking is produced entirely by the additive rule, and no part of it remains under the structural reading the index's own authors state in prose.

Proposition 3 concerns netting. A difference of counts treats dependency and reverse dependency as commensurable and substitutable across dimensions. Because a two-level scale discards intensity, the difference can carry the wrong sign relative to the asymmetry of exposure it is read as summarizing.

Section 5 reconstructs the published index from its published cells and reports what the reconstruction shows, including where the published arithmetic does not reproduce. Section 6 states the requirements a path-aware instrument must meet, and connects them to the measured multi-layer graph of the companion papers (`power-network-wp.md`, `theory-wp.md`).

The object of criticism is the aggregation rule. The expert scores may be entirely sound. The claim is that the rule applied to them supports less than it is asked to support, and that the shortfall is repairable.

2. Related literature

The methodology of composite indicators is well developed and its warnings are standard. Nardo et al. (2008) set out the construction sequence and treat weighting and aggregation as modelling choices requiring justification. Saisana, Saltelli and Tarantola (2005) show that uncertainty and sensitivity analysis routinely overturn country rankings built from defensible-looking indicators, and argue that a ranking published without such analysis is not interpretable. The Monte-Carlo of §5.4 is an instance of their prescription applied to a dependency index.

The substantive literature on asymmetric interdependence is older and supplies the object the index is reaching for. Hirschman (1945) established that the political leverage in a trading relationship comes from the asymmetry of adjustment costs rather than from volumes. Keohane and Nye (1977) separate sensitivity from vulnerability, the latter being the cost of adjustment once alternatives are sought, which is the quantity policy cares about. Farrell and Newman (2019) locate contemporary leverage at network chokepoints, where topology rather than aggregate exposure determines who can coerce whom.

That literature makes the measurement problem harder for a counting index: it says leverage lives in intensities and in topology, and a count of dimensions above a threshold discards both. The companion papers develop the topological side on a measured graph; this paper isolates what the aggregation rule alone determines.

3. The instrument, in the record

The index appears in a study commissioned by the Commission's Directorate-General for Communications Networks, Content and Technology and published by the Publications Office of the European Union as part of the evidence base for the 2028-2034 multiannual financial framework. Its legal notice records that the study reflects the views of its authors and should not be considered representative of the Commission's official position. What follows therefore treats the index as a published measurement practice on which programming decisions draw. The results below concern the arithmetic of the published study, and nothing turns on its status as a Commission determination.

The index is described in the source as follows: dependencies were "investigated by expert and stakeholder interviews, complemented by desk research", and experts "reviewed the level of dependencies of the EU on other global actors, as well as reverse dependencies" (European Commission 2026). Scores are recorded per technology area against eight dimensions: raw materials; high-tech digital components; technology excellence; testing and experimentation facilities; digital infrastructures and software; advanced skilled workforce; specific actors in the research and development ecosystem; and specific sections of the value chain. The scale has two levels, "Some" and "High". The aggregate reported in each row is the number of dimensions rated High.

The netting rule appears once, in a footnote, in full: "Calculated as Total number of dependencies rated high – total number of reverse dependencies rated high".

The same study describes the structure it is measuring in different terms: "Although the system may appear linear and straightforward, it is in fact a tightly coupled and fragile network, where the absence or failure of one essential component within this pipeline can disrupt the whole system." That sentence states a weakest-link structure. The footnote states an additive one. Propositions 1 and 2 concern the distance between them.

4. What a counting index identifies

4.1 Setup

Let units be indexed by i and dimensions by $j = 1, \ldots, J$. Latent exposure is $e_{ij} \in [0, 1]$, unobserved. The instrument records a thresholded score $x_{ij} = \mathbb{1}[e_{ij} \geq \tau]$ for a fixed cutoff τ, and reports the count

$$D_i = \sum_{j=1}^{J} x_{ij}.$$

The weakest-link functional on the same latent profile is $M_i = \max_j e_{ij}$. A parallel instrument records reverse exposure y_{ij} with count R_i, and the published aggregate is the difference $N_i = D_i - R_i$.

4.2 Results

Proposition 1 (counting is not weakest-link). For every threshold $\tau \in (0, 1)$ and every $J \geq 3$ there exist latent profiles a and b with $M_a > M_b$ and $D_a < D_b$. Threshold counting and the weakest-link functional therefore induce different orderings, and neither refines the other.

Proposition 2 (two-level scales make the weakest-link reading degenerate). Suppose the recorded scale is binary and there exists a dimension j^* with $x_{ij^*} = 1$ for every unit i. Then the recorded weakest-link statistic $\max_j x_{ij}$ equals 1 for every unit, so it induces the trivial ordering, and any non-trivial ranking of the units is produced entirely by the aggregation rule applied on top of the scale.

Proposition 3 (netting counts is not comparing bargaining positions). Let disruption on dimension j impose cost $u_j \geq 0$ on the home party and $v_j \geq 0$ on the counterpart, and let the recorded scores be $x_j = \mathbb{1}[u_j \geq \tau]$ and $y_j = \mathbb{1}[v_j \geq \tau]$. Then for any τ there exist cost profiles with $N = \sum_j x_j - \sum_j y_j > 0$ and $\sum_j v_j > \sum_j u_j$: the net count reports the home party as the more dependent while total adjustment cost falls more heavily on the counterpart. The construction extends to the maximum-cost comparison.

Proofs are in Appendix A.

4.3 Reading

Proposition 1 is elementary and its force is contextual. Nothing prevents a policy instrument from choosing an additive rule, provided the choice is stated and defended. A policy instrument cannot state a weakest-link structure in the analysis and use an additive rule in the index, because the two disagree on the object they are both claimed to measure.

Proposition 2 says the coarseness of the scale is not a matter of precision alone. Where a dimension is universally high, a two-level scale removes the weakest-link reading entirely, leaving the additive rule as the sole source of the published ordering.

Proposition 3 isolates what the netting operation assumes. Subtracting one count from another treats a high dependency on raw materials as offsetting a high reverse dependency in skilled labour, at par, and treats all high scores as equally intense. The interdependence literature of §2 locates leverage in the very intensities this discards.

5. Application: the published index

5.1 Reconstruction

We transcribe both published tables cell by cell and recompute the row sums (`verify/dependency_index.py`). Twelve of the thirteen printed totals reproduce exactly. The artificial-intelligence row contains seven cells rated High and prints an overall of 6. We carry the reconstructed value forward and flag the discrepancy; nothing below turns on which is correct.

5.2 The published rule against the published tables

Applying the footnote's rule to the two tables yields the net values:

Area	Dependency	Reverse	Net
Data analytics	5	6	−1
Technologies for interoperability	4	3	1
Photonics	4	2	2
Microelectronics	5	2	3
Quantum	4	1	3
Robotics	4	1	3
Blockchain	7	3	4
Cloud	4	0	4
Cybersecurity	5	0	5
High-performance computing	5	0	5
Next-generation internet	6	1	5
Advanced connectivity	6	0	6
Artificial intelligence	7	0	7

The accompanying prose states that data analytics, technologies for interoperability, photonics, quantum and robotics "show a net dependency of more than three types", and that "All other areas have a net dependency rating of more than four".

Neither claim follows. Every member of the named group scores at or below three. Data analytics scores minus one, a net *reverse* dependency, so the sign is inverted relative to the group it is placed in. In the remaining group, microelectronics scores three under any reading. Under the strict reading of "more than four", blockchain and cloud at four also fail, giving three failures of eight; under a charitable reading of the phrase as "four or more", microelectronics alone fails. The data-analytics sign inversion is robust to both readings.

5.3 Degeneracy of the weakest-link reading on this data

All thirteen areas score High on digital infrastructures and software. By Proposition 2 the recorded weakest-link statistic is therefore constant across areas, and the published ranking is entirely a product of the additive rule, which the source does not defend and which its own §3.12 structural description contradicts.

Five of the thirteen areas carry a reverse-dependency count of zero across all eight dimensions: artificial intelligence, advanced connectivity, cloud, cybersecurity and high-performance computing. The reverse-dependency table's entry for digital infrastructures and software, the dimension on which every area scores High in the dependency direction, is empty.

5.4 Rank stability under reweighting

The index weights the eight dimensions equally. Following Saisana, Saltelli and Tarantola (2005), we draw random positive weights and recompute the ranking (`mc_weights`, seed 20260718, 20,000 draws). The most-dependent area changes on 49.8% of draws, alternating between artificial intelligence and blockchain. The identity of the most dependent technology area is thus driven by the equal-weighting convention; it does not follow from the recorded scores.

6. What a path-aware instrument requires

The results imply four requirements, each of which the companion instrument is built to meet.

First, intensity must reach the aggregation step. A scale with more than two levels, or a continuous cost measure, is a precondition for any statement about which dependency binds hardest, and for the netting operation to mean anything.

Second, the aggregation rule must be stated and matched to the claimed structure. Where the argument is that one failing component disrupts the system, the rule is a minimum over the relevant classes, and the ranking it produces is the one to publish.

Third, exposure must be scoped to nodes and composed along paths. Dependency accrues along production paths, and a per-technology score cannot represent the composition; two areas with identical dimension scores can face different exposure because their supply paths traverse different controllers. The companion papers develop the composition rule and the associated cut calculus.

Fourth, coverage must include services. The one dependency indicator the source computes rather than elicits is a goods trade balance, and its stated scope is explicit: "It is to be noted that this only concerns goods, as digital services are not recorded in the same way." Software, cloud and platform layers are outside it, which excludes much of what the exercise is about.

Sensitivity analysis over weights and thresholds should accompany any published ranking, as §5.4 illustrates.

7. Discussion and limits

This paper analyses a rule applied to scores; it does not evaluate the scores. If the expert judgements are sound, a better aggregation of them would be informative, and reconstructing the index is a step toward building one.

The latent-variable framing of §4.1 imputes a continuous exposure behind a two-level record. That is the natural reading of a scale whose levels are labelled "Some" and "High", and Proposition 2 does not depend on it, being a statement about the recorded values alone. Propositions 1 and 3 do depend on it, and they should be read as statements about what the recorded values can and cannot recover from any such latent structure.

The reconstruction depends on a text extraction of the published tables. Cell-level transcription was checked by hand against the source for the rows carrying the results of §5.2, and the full transcription is in `verify/dependency_index.py` for inspection. The single printed-total discrepancy in §5.1 may be an error in the source, an error in extraction, or an intentional adjustment; we do not resolve it and no result rests on it.

The netting critique of Proposition 3 applies to the published rule. It does not follow that the underlying reverse-dependency scores are uninformative, and §5.3′s two observations about them are reported without interpretation for that reason.

8. Conclusion

An index that counts dimensions above a threshold measures neither the worst dependency nor the balance of adjustment costs. On a two-level scale with one universally high dimension, it measures nothing but its own aggregation convention. The index currently in use for European economic-security programming has both properties, its published arithmetic does not reproduce its published prose, and the identity of the technology area it ranks most dependent flips on half of all admissible reweightings. The requirements for an instrument that would support the claims made of it are stated in §6, and are the ones the companion measurement papers implement.

Appendix A. Proofs

Proposition 1. Fix $\tau \in (0,1)$ and $J \geq 3$. Let $a = (1, 0, \dots, 0)$ and $b = (\tau, \dots, \tau, 0)$ with $J-1$ entries equal to τ. Then $M_a = 1 > \tau = M_b$, while $D_a = 1$ and $D_b = J - 1 \geq 2$, so $D_a < D_b$. Since the pair (a, b) is ordered one way by M and the other way by D, neither ordering refines the other. □

Proposition 2. By hypothesis $x_{ij^*} = 1$ for all i, so $\max_j x_{ij} \geq 1$. Since the scale is binary, $\max_j x_{ij} = 1$ for all i. A constant statistic induces the ordering in which all units are tied. Any strict ranking of the units must therefore be generated by a functional other than the recorded maximum, which by construction is the aggregation rule applied to the same recorded scores. □

Proposition 3. Fix $\tau \in (0,1)$ and $J \geq 6$. Let $u_j = \tau$ for $j \leq 5$ and $u_j = 0$ otherwise, and let $v_6 = K$ for some $K > 5\tau$ and $v_j = 0$ otherwise. Then $x_j = 1$ for $j \leq 5$ and $y_j = 1$ for $j = 6$, so $N = 5 - 1 = 4 > 0$, reporting the home party as substantially more dependent. But $\sum_j v_j = K > 5\tau = \sum_j u_j$, so total adjustment cost falls more heavily on the counterpart. Taking $K > \tau$ also gives $\max_j v_j > \max_j u_j$, extending the reversal to the maximum-cost comparison. □

References

- European Commission, Directorate-General for Communications Networks, Content and Technology (2026). *Study on the EU's critical digital capacities deployment beyond 2027: final report.* Prepared by Technopolis Group; K. Izsak, P. Simmonds, M. Ploeg, S. Viscido, A. Lotito, C. Moreno & R. Alasghar (eds.). Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40079-1, doi:10.2759/0673677.
- Farrell, H. & Newman, A. L. (2019). Weaponized interdependence: how global economic networks shape state coercion. *International Security* 44(1):42–79.
- Hirschman, A. O. (1945). *National Power and the Structure of Foreign Trade.* Berkeley: University of California Press.
- Keohane, R. O. & Nye, J. S. (1977). *Power and Interdependence: World Politics in Transition.* Boston: Little, Brown.
- Nardo, M., Saisana, M., Saltelli, A., Tarantola, S., Hoffman, A. & Giovannini, E. (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide.* Paris: OECD Publishing.
- Saisana, M., Saltelli, A. & Tarantola, S. (2005). Uncertainty and sensitivity analysis techniques as tools for the quality assessment of composite indicators. *Journal of the Royal Statistical Society, Series A* 168(2):307–323.

Annex A. Working notes (temporary annex)

This annex lists internal working documents and will be removed from the final version. Paths are relative to the project root.

- `papers/notes/counting-dependencies-wp.md`: internal notes for this paper.

- `papers/power-network-wp.md`, `papers/theory-wp.md`: the companion measurement and formal papers whose requirements §6 states.

Annex B. Reproduction: scripts, data, and figures

One command reproduces every in-text number: `python verify/dependency_index.py` (seconds; python 3, numpy; deterministic, no network).

- `verify/dependency_index.py`: the transcribed tables (`DEPENDENCY`, `REVERSE`, `PUBLISHED_OVERALL`), gate-2 checks `check_D1`…`check_D4` (seed 3 for the search in D4), and gate-3 `mc_weights` (seed 20260718).
- `verify/results/dependency_index.json`: the recomputed counts, the net values of §5.2, the printed-column discrepancy, the failing prose claims, the zero-reverse-dependency areas, and the Monte-Carlo configuration and outcome.

No figures are used in this paper.