The Contribution Criterion: Scoring Policies and Organisations by Their Effect on System-Level Digital Sovereignty

Stefane Fermigier (Abilian) · sf@abilian.com

Draft working paper v0.5, 2026-09-18. Prepared for publication: the citations to the programme corpus now carry the published addresses of the papers they name, with ten stale version pins corrected and one broken path repaired; the one source not yet published is flagged as such in place. A prose pass brought the document's coordination and antithesis rates inside the corpus's own range. Revised after two referee rounds, the second by an independent reviewer given only the papers: the proof of Proposition 1 now closes its local-to-global step, with the planner-corner claim corrected to its exact threshold (both verified in `verify/claim1.py`). The sign-transfer lemma is restated with its hypotheses in full. The assembled model is stated completely in this paper. Facts carrying open verification flags are marked in place.

Abstract

Every deployed sovereignty instrument scores a node: an organisation's own resilience, a supplier's own posture, an offer's own compliance level. This paper defines the complementary object: the *contribution* of a policy or an organisation, the change in the system's sovereignty stock under a stated counterfactual replacement of the entity's realized actions. Four rules discipline the definition: only realized mechanisms score, a sign verdict requires agreement across the canonical counterfactuals, conduct is scored over a stated window, and every claim inherits the weakest evidence tier on its path.

The formal core is classical: in a layered adoption market with consistency effects, the equilibrium under-adopts relative to the adopter-welfare optimum, the textbook network-externality mechanism (Katz and Shapiro 1985; Farrell and Saloner 1985; the lineage runs to Rohlfs 1974). The paper adds the closed form of the wedge in this reference parameterization, $2\gamma \int_c^1 a\, dG$ evaluated at the equilibrium cutoff, and exact thresholds for when the result holds under a fuller welfare ledger: a local threshold at the equilibrium and a corner threshold at universal adoption, both verified. The structural reading is that node-scoped instruments cannot register the wedge, because it lives on other nodes, whichever way the scores point. A sign-transfer lemma with explicit hypotheses connects the computable definition to two applied screens. Three worked cases exercise them with stated, sign-checked counterfactuals: a suite-bundling conduct scores negative through a mechanism invisible to every node instrument, with one countervailing lower-tier path recorded; the commitments ending it score positive with a receipted timing discount; and a full-funding maintenance programme scores positive, provisionally, exactly where the standard leverage statistic returns nothing.

Keywords: digital sovereignty, policy evaluation, contribution analysis, network externalities, procurement, bundling.
JEL: F52, L86, D62, O25.

1. Introduction

An organisation can pass every node-level sovereignty screen while making the system it operates in less sovereign, and can fail every screen while supplying the system's most valuable options. Both directions occur in practice: a locally incorporated, locally staffed, certified subsidiary re-badging a foreign-controlled stack improves its scores and deepens the dependency. A one-maintainer open-source project fails any resilience test while keeping an option alive for everyone. These two phenomena are screen-blindness cases, outside the formal model below and inside the worked cases and the queue. The model formalizes a third, subtler failure: even where node scores point the right way, they cannot price what an adoption decision does to *other* nodes. The question that decides policy is the system one: does this entity's behaviour move the *system's* sovereignty, in which direction, and by how much.

This paper defines that object and makes it usable. Section 2 gives the definition, the evidence-tier vocabulary and the four discipline rules. Section 3 states the assembled model in full, places its core on the network-externality literature it instantiates, and proves the under-adoption result with exact ledger thresholds. Section 4 states the sign-transfer lemma with its hypotheses in the open and derives the two applied screens. Section 5 applies the screens to three real conducts with their counterfactuals stated and sign-checked, obtaining both signs, one recorded countervailing path, and one disagreement, by design, with a standard policy statistic. Section 6 locates the criterion among the remedies and states what it is not. The factor map and prioritization the criterion computes on are the companion paper (*What Moves the System*, `factor-map-wp.md`).

The criterion's method sits in the tradition of contribution analysis (Mayne 2001; 2019), which assembles a defensible causal claim that an intervention contributed to an outcome in a complex system, without demanding full attribution. Two additions distinguish the present instrument: the causal steps are edges of a receipted factor map rather than a programme theory drawn for the occasion. The counterfactual is stated and sign-checked on a formal model, under a lemma whose hypotheses are explicit.

2. The criterion

Let s be the system state (the measured multi-layer graph with its factor values) and $V(s)$ the outcome metric: the sovereignty stock as option value under the cut calculus, evaluated at the feedback structure's fixed point. An entity E (a policy, a programme, an organisation, or a bounded conduct of one) has a **footprint** A_E over a stated window: the set of its realized actions, each mapped to the factor-edge perturbations it effects. Write s_{-E} for the counterfactual state: the state recomputed to its fixed point, on the stable branch continued from s (Section 4 makes the branch condition precise), with E's actions replaced by the stated counterfactual. The **contribution** of E is

$$C(E) = V(s) - V(s_{-E}).$$

The semantics are perturbation-then-recompute on the cyclic factor graph; identification-style causal claims are reserved for acyclic sub-questions over bounded horizons.

Evidence tiers (shared with the companion and used by Rule 4): T1, proved and numerically verified within the programme; T2, measured with sources and stated limits; T3, documented case or literature-sourced claim; T4, conjectured. Slashed tiers record mechanism over applicability; the weaker of the two is operative. The tiers grade internal verification and say nothing about external validity.

Four rules govern application.

Rule 1 (realized mechanisms only). A_E contains what the entity did. Declared aims, obtained labels, and compliance postures enter only through the edge changes they actually effected. The rule applies symmetrically to entities whose goals the analyst shares, as Section 5′s Case B shows the rule binding against a sympathetic case.

Rule 2 (stated counterfactual, sign-robustness). Two canonical counterfactuals are computed: *removal* and *next-best*. Removal is selected by a stated rule: the minimal counterfactual that deletes the scored actions while preserving the entity's product set (the same goods and services offered; the scored practice absent). This rule keeps a conduct's removal from being confused with the entity's exit. Next-best is the most plausible alternative conduct by the same entity, or instrument by the same authority. $C(E)$ is reported with its counterfactual named. A sign verdict is issued only when the sign agrees across both. Where the two counterfactuals nearly coincide, the case reports it: agreement between near-identical counterfactuals is a weak check.

Rule 3 (windowed conduct, never essence). The criterion scores a footprint over a window. The same entity can carry opposite signs on different conducts or periods, which Section 5 exhibits. Questions outside the scored conduct's footprint are reported as out of scope; entity-level aggregation across conducts is left to the quantitative layer.

Rule 4 (tier inheritance). A contribution claim inherits the weakest operative evidence tier along its path from action to outcome.

Let $N(E)$ be any node score. N and C are orthogonal by construction: N measures the entity's own posture, C the entity's effect on V. Nothing in either definition constrains the other. The next section shows where the freedom binds.

3. The composition result

The classical core. That a positive adoption externality produces equilibrium under-adoption relative to a welfare optimum is a foundational result of the network-effects literature. The critical-mass structure goes back to Rohlfs (1974), the under-adoption mechanism to Katz and Shapiro (1985) and Farrell and Saloner (1985), with the survey in Farrell and Klemperer (2007). Nothing in the mechanics below is new; the proposition is stated as an instantiation. This paper adds (i) the closed form of the wedge in the programme's reference adoption model, with exact thresholds for when the result holds under a fuller welfare ledger, and (ii) the reading that motivates the criterion. The wedge is a sum over *other* adopters' payoffs, so no node-scoped instrument registers it, whichever way the node scores point.

The model, in full. Buyers have altitude $a \sim G$ with density $g > 0$ on $(0, 1)$, where altitude is the share of a buyer's requirement that is differentiated rather than commodity, so that a high-altitude buyer values the incumbent's specialised services most. A fraction $\lambda \in (0, 1]$ of buyers receives a decision event; the alternative carries additive friction $s \geq 0$. An event-receiving buyer adopts iff $a\,B(m) > \pi_C + s$, where m is the adoption mass and $B(m) = \pi_C + \pi_D + \gamma\,(2m - 1)$ is the consistency-adjusted value slope, with $\gamma \geq 0$ the consistency weight and the maintained assumption $\pi_C + \pi_D > \gamma$ so that $B > 0$ on $[0, 1]$. Not adopting pays the buyer 0 (the ledger below revisits this normalization). Cutoff strategies give $m(c) = \lambda(1 - G(c))$ an equilibrium cutoff solves $c\,B(m(c)) = \pi_C + s$ and the stability denominator is $D = B - 2\gamma\lambda c\,g(c) > 0$ at a stable equilibrium. The proposition below specializes to $\lambda = 1$, $s = 0$; the companion's elasticity computations use the full parameterization. One clarification on node scores: deployed node instruments may well reward adoption of the domestic alternative (the DRI/IRN's exposure and exit-path criteria do), so node optimization is not assumed to point against adoption. The point proved here is that whichever way a node score points, the wedge term lives on other nodes and no node-scoped instrument prices it. Adopter-side welfare is $W(c) = \int_c^1 (a\,B(m(c)) - \pi_C) dG$.

Proposition 1 (under-adoption; instantiation). *Let $c_{eq} \in (0, 1)$ be an interior equilibrium cutoff, and let the planner maximize W. Then*

$$W'(c_{eq}) = -\,2\gamma\, g(c_{eq}) \int_{c_{eq}}^{1} a\, dG < 0 \quad \text{whenever } \gamma > 0,$$

and moreover $W(c) < W(c_{eq})$ for every $c > c_{eq}$, so the W-planner's cutoff (interior or corner) lies strictly below c_{eq}, with strictly higher welfare and adoption mass. At $\gamma = 0$ the wedge vanishes and the two cutoffs coincide: for the W-planner, the divergence region is exactly consistency effects present and equilibrium interior. (Proof: Appendix A, including the global step and the multiplicity note.)

Corollary. *Since the W-planner's profile has $m_p > m_{eq}$, any outcome V strictly increasing in m is strictly higher at that profile. The corollary inherits the W-planner's choice: it prices the* *welfare optimum in V units, and is not a statement about a V-maximizing planner (who would trivially prefer $m = 1$ at every γ).*

The wedge $2\gamma \int_{c_{eq}}^{1} a\, dG$, evaluated at the equilibrium cutoff, is the consistency externality the marginal buyer ignores: its adoption would raise m and thereby every inframarginal adopter's payoff. No node score of the buyer registers this term, because the term lives on other nodes. That is the composition failure: the individually rational choice profile is collectively wasteful whenever the wedge is positive, however much node instruments favour adoption.

The ledger, and when the result withstands it. W counts adopters' payoffs with non-adoption normalized to zero: incumbent-side surplus is outside the ledger as stated. Giving every non-adopter a surplus $u_0 > 0$ yields $\widetilde{W}(c) = W(c) + u_0\, G(c)$ and two exact thresholds, both verified numerically on both sides. At the equilibrium margin, $\widetilde{W}'(c_{eq}) = g(c_{eq})\,(u_0 - 2\gamma I_1(c_{eq}))$: the marginal buyer's adoption is socially beneficial exactly when the wedge exceeds the forgone incumbent-side surplus, so the local under-adoption result holds iff $u_0 < 2\gamma I_1(c_{eq})$. At the corner, the planner's optimum remains universal adoption for every $u_0 \leq 2\gamma\, \mathbb{E}[a] - \pi_C$ and leaves the corner above that threshold. At the source model's defaults ($\pi_C = 0.30$, $\pi_D = 0.55$, $\gamma = 0.50$, uniform G) the corner threshold is 0.2 the corner is robust to a non-adopter surplus of up to two-thirds of π_C, and the numeric suite confirms lift-off immediately above the threshold. The defaults' equilibrium numbers are $c_{eq} = 0.2805$, $m_{eq} = 0.7195$, wedge 0.4607, welfare gap $0.375 - 0.277$ against the $u_0 = 0$ corner. One scope statement completes the ledger: producer surplus, the incumbent's lost profit and the entrant's gained profit, remains outside on both sides. This matters for the sign's possibility space, because the network and tying literatures obtain both under- and over-adoption once producer-side terms enter (Katz and Shapiro 1985; Farrell and Saloner's excess-momentum case).

Coordination and the externality, separated. At large γ the fixed-point map can admit multiple equilibria; the proposition compares the planner to a given interior equilibrium (the argument applies at each), and under multiplicity part of the gap is a coordination problem. The distinction matters for remedies: coordination devices (anchor buyers, buyers' clubs, the programme's F6 results) can move the system between equilibria without any scoring instrument, while the instruments of Section 4 price the externality wedge at a given equilibrium. The criterion's necessity is therefore argued comparatively, never proved: it registers conducts that node instruments cannot see (Section 5, Case A), and it evaluates policies beyond coordination's reach.

The supplier-side reading. On the supply side no new proposition is needed. *Screening for Sovereignty* (Thm 2; précis in Annex A) proves that attributes without a type-cost gap cannot screen: residence, staffing, and certifications are verifiable at low cost because both genuine and controlled suppliers can hold them. Node scores built on such attributes are then improvable without any change in effective control, so supplier-side score optimization converges on checkable postures and leaves V unchanged while diverting the procurement spend attached to the scores. This paper states that as a reading of the cited theorem, and claims no new result in it.

4. From the definition to the screens

The criterion is a computable quantity; the screens below are question lists. The bridge is a lemma whose hypotheses are in the open, and whose scope is the map's well-signed causal subgraph: the factor map's moderator edge (F12) and reweighting edge (F16, signed $\pm$) are outside it by construction.

Lemma 1 (sign transfer). *Consider a conduct whose footprint perturbs factors X, and let the perturbed subsystem be the set of factors reachable from X along receipted causal edges, together with those edges. Suppose: (a) scoping: the screen questions left unanswered correspond* *to factors outside the footprint, with no path from X to the outcome leaving the well-signed causal subgraph; (b) path-sign agreement: every simple directed path from a factor in X to the outcome has a sign product (edge signs composed with the outcome gradient's stated signs), and all such products equal σ, with at least one path strict; (c) cycle positivity: every directed cycle within the perturbed subsystem has positive sign product; (d) stable branch: along the homotopy $\theta \in [0,1]$ from s_{-E} to s, the fixed point stays on a continuously differentiable stable branch (spectral radius of the subsystem Jacobian below one throughout; the model's $D > 0$ is the scalar case), which is also what selects s_{-E} under multiplicity. Then* $\operatorname{sign}(C(E)) = \sigma$*, under every counterfactual preserving (a)–(d).* (Proof: Appendix B.)

The increment Lemma 1 adds. Reading a comparative-static sign off the signs of a system's edges is qualitative comparative statics, a genre the lemma belongs to squarely. The sign-solvability tradition asks when the sign pattern of a matrix determines the sign of the solution without its magnitudes (Quirk and Ruppert 1965; the qualitative-matrix line after it, *citation to be verified*). The closest modern relative is the signed-graph rule for causal inference, where the sign of an effect transfers along a directed acyclic graph whose edges all carry monotone signs (VanderWeele and Robins 2010, *citation to be verified*). Against those, the increment here is narrow: the object is a cyclic fixed point, so walks rather than paths carry the perturbation. Hypothesis (c) buys back the acyclic conclusion by making every cycle sign-preserving; hypothesis (d) makes the fixed point differentiable in the first place. A reader who knows the qualitative-comparative-statics literature should read Lemma 1 as that literature's rule, extended to a cyclic system along a stable branch.

One hypothesis needs stating more strongly than the endpoint conditions suggest. The Neumann expansion in Appendix B runs over the whole homotopy, so (b) and (c) must hold at every $\theta \in [0,1]$ and not only at the two endpoints: each edge of the perturbed subsystem keeps its stated sign along the whole branch. This is a global monotonicity of the edges, stronger than sign agreement between s_{-E} and s it licenses treating the signs as constants inside the walk sum. Where an edge changes sign somewhere along the branch, the lemma does not apply and the case goes to the quantitative layer.

The hypotheses are substantive, and each is checkable per case: (a) is Rule 3′s scoping plus the subgraph restriction; (b) and (c) are read off the factor map's signed edges, at the map's tiers, inherited by Rule 4; (d) is the stability condition the elasticity computations instantiate. A case whose paths disagree in sign, or whose perturbations propagate through a negative cycle, falls outside the lemma. Its verdict needs the quantitative layer, or, where the disagreeing path is of strictly lower tier than the verdict's, the case may issue the verdict at its tier with the countervailing path recorded as an explicit caveat. A verdict in Section 5 is an invocation of this lemma with (a)–(d) checked and the path enumeration shown.

Remark 1 (what cycle positivity buys, and what it does not). Hypothesis (c) is sufficient and appears not to be necessary, which matters because a real factor map will contain negative cycles. The suite (`verify/lemma1.py`, seed 20260830) draws random signed systems on a stable branch: on all 590 draws satisfying (b) and (c) the conclusion holds, as the lemma says. The sign is never reversed on the 1,180 draws that satisfy (b) and violate (c). Two facts explain the second number. Under (d) every eigenvalue of $I - J$ has positive real part, so $\det(I - J) > 0$ and the determinant cannot supply a reversal; and a lone negative cycle of weight $c \in (-1, 0)$ multiplies its path's walk sum by $1/(1-c) \in (0,1)$, which attenuates the effect and cannot turn it over. A reversal therefore requires a path cofactor to go negative, which takes a strongly connected remainder near the stability boundary rather than any single negative edge or cycle. The operational reading is unchanged and slightly softer than Section 4 states it: a case violating (c) is not thereby wrong-signed; the lemma simply stops certifying it, and sends it to the quantitative layer.

Remark 2 (why the edge-sign hypothesis is global). The same suite exhibits a three-factor system whose edge from the middle factor to the outcome is positive at both ends of the homotopy and negative in its interior. Both endpoints satisfy (b) and (c), the branch is stable throughout, and the integrated contribution is negative against an endpoint path sign of +1. Checking the hypotheses at s_{-E} and s alone would therefore certify the wrong sign; that is why the hypothesis is stated over the whole branch.

The screens: two question lists, one for policies and one for organisations, gate questions leading; the ordering is presentation, and answers, signs, and tiers are order-invariant. Each question names its factor edge and its receipt where one exists; two questions rest partly on conjectured factors and say so. Answers are signed (+, 0, −) with the evidence tier of the answer. The output is the signed profile plus a sign verdict; there is no scalar, because a weighted composite would re-create the asserted-weights instrument this criterion exists to avoid.

The policy screen.

#	Question	Edge	Receipt
P1	Does it generate or restore genuine adoption decisions, or spend inside suppressed decisions, where sub-threshold advantages are provably inert?	F7	*The Unpriced Bundle* B1–B3
P2	Does it verify effective control, or checkable proxies that cannot screen?	F8	*Screening for Sovereignty* Thm 2, Cor. 2
P3	Is the funded activity additional, or would it have happened anyway?	F4	*Leverage Without Additionality* P1
P4	Does it install enforced exit rights or lower measured exit friction?	F10, F9	*Exit or Own?* P1–P2; *Cloud Services Matrix*
P5	Does it retain control of what it grows, or fatten acquisition targets?	F13	*Sovereignty as a Flow*
P6	Does it aggregate demand toward anchor mass, or fragment it?	F6	*Buyers' Clubs* P3
P7	Does it target a measured factor, or a proxy index whose aggregation can invert rankings?	F12	*Counting Dependencies* P1–P3
P8	Does it create washing surface: a label or premium without a verification channel?	F14	*Scoring the Seller*
P9	Does it make interoperability the multiplier, or re-create lock-in around preferred suppliers?	F9, F18	*Exit or Own?*; F18 conjectured, tier flagged
P10	Are its effects controller-disjoint at depth across the layers it touches, or certified layer by layer?	F11	*Cuts and Options*

The organisation screen. It scores a conduct over a window, never an identity. Questions marked (B) bind mainly for buyer-side entities.

#	Question	Edge	Receipt
O1	Does the offer add a controller-disjoint option at its layers, or re-badge an existing controller behind a local entity?	F11	*Cuts and Options* §6.2; *Path Sovereignty*
O2	Do its products lower or raise the exit tax faced by its customers?	F9	*Cloud Services Matrix*
O3	Will its open components stay open, or is it in the enclosure region?	F3	*Value Capture* P2–P3
O4	Who will control it within the window: ownership, funding path, acquisition exposure?	F13, F19	*Sovereignty as a Flow*; tier capped

#	Question	Edge	Receipt
O5	Does it itemise or bundle: does its pricing generate decision events downstream or suppress them?	F7	*The Unpriced Bundle* B1–B3
O6	Does it feed the commons it depends on, or drain them?	F4, F6	*Buyers' Clubs* P4–P5
O7	(B) Does its procurement screen for control and pool demand, or buy the default and fragment?	F8, F6	*Screening for Sovereignty*; *Buyers' Clubs*
O8	Does its sovereignty marketing exceed its verifiable posture?	F14	*Scoring the Seller*
O9	Does it author open interfaces and standards, or capture them?	F18	conjectured, tier flagged
O10	Does it build capability on open, transferable stacks, or certify dependence?	F16, F17	*Why Driver Weights Differ*; F17 conjectured

Discipline. Facts about named entities must be public and sourced; absent the fact, the question is reported as unanswered. Conjectured-edge answers always carry their flag and never carry a verdict alone. Before release, the profile is re-read under the assumption that the entity is one the analyst favours, and again one the analyst opposes, to confirm no answer rests on aim, identity, or reputation.

5. Worked cases

Each case states its footprint, its two counterfactuals per Rule 2, its path enumeration and screen answers, and its verdict as an invocation of Lemma 1.

Case A: suite bundling of a communications product (2017 to the 2025 commitments). Microsoft distributed Teams inside its productivity suites at no separate price from 2017; a complaint followed in 2020, proceedings in 2023, and commitments made binding in 2025 (European Commission 2025a). *Footprint:* zero-priced inclusion removes the price signal that triggers procurement decision events (F7, T1 mechanism), and each suppressed cycle deepens the installed base and the next cycle's exit tax (L3, T3). *Counterfactuals:* removal, under Rule 2′s selection rule, is the same vendor offering the same products itemised from 2017; next-best conduct is the configuration the vendor itself adopted after 2023, suites with and without the product at distinct prices. The two nearly coincide, and the case reports it: their agreement is a weak check here. The verdict leans on the path analysis, with counterfactual diversity adding little. *Paths:* from F7, the receipted routes to the outcome are F7 → F2 → F11 → OUT and F7 → F2 → F5 → (L1) → F2 → F11 → OUT (both: suppression lowers the share and its option supply; sign product negative for the conduct), F7 → F9 → {F5, F7} (suppression raises accumulated friction; negative), and one countervailing route: F7 → F2 → F13 → F11 → OUT, on which suppression, by keeping the challenger small, also keeps it a lesser capture target (sign product positive for the conduct). The countervailing route runs through two T3 edges; the direct routes carry T1 mechanisms. L1 and L3 are positive-sign cycles, satisfying Lemma 1(c); the countervailing path breaks strict path-sign agreement, so the verdict is issued under the lemma's caveat clause. *Screen:* O5 negative (T1 mechanism, T3 record), O2 negative (T2/T3), O1 zero; O6 and the remaining questions are out of the conduct's footprint under Rule 3 (the same entity's commons contributions would enter a different conduct's footprint). **Verdict: negative, headline tier T3, with one countervailing T3 route recorded and unquantified (suppression also spared the challenger capture exposure by keeping it small).** The mechanism is invisible to every node instrument: no score of the vendor or of any buyer registers a suppressed decision.

Case A′: the commitments as an instrument (2025 onward). *Footprint:* suites without the product at reduced prices, a display duty on the relevant vendor webpages, switching rights, and interoperability and portability commitments running seven years, ten for interoperability and portability (European Commission 2025a). *Counterfactuals:* removal is no remedy (suppression continues on the installed-base path); next-best instrument is a prohibition decision reached earlier in the proceeding. Against removal the restored events are a strictly positive F7 perturbation. Against the earlier-prohibition counterfactual the sign is still positive (later restoration dominates no restoration), while the magnitude is smaller than the alternative would have delivered. That is a statement about size, and Rule 2 requires only the sign to agree. *Paths:* the same enumeration as Case A with the perturbation reversed, including the countervailing capture route, recorded symmetrically. *Screen:* P1 positive (T1 mechanism through a T3 instrument), P4 positive, P2 zero (control verification is outside its object). **Verdict: positive, headline tier T3, expected magnitude modest**, and the discount is receipted: the instrument arrives after years of suppressed cycles whose installed base persists (L3 already ratcheted), so restored events face the accumulated exit tax. The same product thus carries a negative conduct verdict and a positive instrument verdict on adjacent windows, which is Rule 3 operating as designed.

Case B: a full-funding maintenance programme (the Sovereign Tech Agency). *Footprint:* maintenance funding of critical open components at a full-funding rate (F4 to F11, edge tier T1/T2), with programme figures and the selection-criteria description as quoted in *Leverage Without Additionality* §6.5, which carries a "citation to be verified" flag that this case inherits. *Counterfactuals:* removal is the components unfunded, with the thin private counterfactual documented in the source paper; next-best is a co-funding design, against which the case must argue that co-funding would have delivered less maintenance for exactly the deep-dependency components at issue (the market failure the source paper formalizes: no private demand to match), so the sign against next-best is also positive. *Paths:* F4 → F11 → OUT, single receipted route, no countervailing path in the subgraph. *Screen:* P3 positive, with a tension the instrument's own rules expose: the answer rests partly on the programme's *published selection criteria*, a declared rule, and Rule 1 scores realized mechanisms only. The realized mechanism is the funded-component list; the criteria description is the declared half, and the open verification flag has not yet confirmed its realized application. The answer is therefore below the instrument's own standard, and the case reports that. **Verdict: positive, provisionally: sign-robust on the removal and next-best arguments as stated, magnitude unquantified, tier-capped by the open flag, and subject to Rule 1 completion (verifying the realized funding decisions) before the verdict is issued without the qualifier.** The criterion and the standard statistic still disagree for a reason the corpus has formalized: a programme can contribute exactly where leverage returns nothing.

Queued. Three cases await their factual records: the Cloud Sovereignty Framework's award criterion itself, whose verdict is conditional (a sovereignty premium is washing surface absent verification and the anchor a verification infrastructure needs); the Data Act egress phase-out, whose natural experiment dates from January 2027; and the class of hyperscaler sovereign-cloud joint ventures, where the sign is open and the factors at issue are re-badging and the control trajectory, and which require an assembled per-entity record before any entity is named.

6. What the criterion is not

Not a composite index: the output is a signed, tiered profile with a verdict. Not a judgment of intent, nationality, or identity: only footprints score, over stated windows. Not a substitute for node instruments: procurement still needs node screens, and the composition result says why the two must coexist. Node screens select postures, and only a system-level reading can say whether the selected postures help. Not the only remedy for the wedge: coordination devices attack the multiplicity part of the gap without any scoring instrument (Section 3). The criterion covers the territory neither node scores nor coordination reaches: registering conducts, and evaluating policies, at the system level. Not an advocacy device: Rule 1 and the symmetry check are the design features that make the criterion inconvenient for its own author; Case B is the demonstration.

7. Limitations

The criterion's precision is bounded by the factor map's tiers, and its magnitudes await the quantitative layer: at present, verdicts are signs with receipts, and magnitude language beyond "modest" or "unquantified" is not licensed. Where footprints overlap, the intended attribution device for the quantitative layer is a Shapley-style share; nothing in this paper computes it, so until something does, overlapping-footprint cases are out of scope. The composition result is an instantiation of classical mechanics in one assembled model with an adopter-side ledger; producer surplus is the open margin on both sides, and richer strategy spaces (dynamic entry, endogenous quality, an explicit supplier game) are open with it. Lemma 1′s hypotheses are checkable but substantive: a conduct whose paths disagree at equal tier, or whose perturbations cross a negative cycle, needs the quantitative layer. The lemma's stable-branch condition means the criterion is currently defined only away from fold points of the fixed-point map. The Monte-Carlo behind the proposition qualifies 266 of 400 draws, the excluded draws being parameter-inadmissible or unravelled zero-adoption fixed points; at an unravelled equilibrium the gap between planner and equilibrium is pure coordination failure, outside the proposition's marginal statement. The worked cases depend on public records. The queued cases show the binding constraint in practice: assembling per-entity factual records to the screen's evidentiary standard is the instrument's main operating cost. On evaluation method, Case A′ belongs to the practice of ex-post evaluation of competition remedies, which the Commission itself applies (its 2005 Merger Remedies Study graded the effectiveness of 85 remedies from 40 decisions). The retrospective literature has sharpened that practice into a caution directly relevant here: conduct remedies are generally found weaker than structural ones at restoring outcomes (Kwoka 2015). The criterion's addition to that practice is the outcome variable: remedy retrospectives measure market outcomes (prices, shares, entry), while the assessment series measures the system-level sovereignty effect along a receipted causal path. The two are complements, and the series adopts the practice's timing convention (effects assessed at stated intervals after the remedy). Case A′'s verdict of positive, modest, and late is consistent with the conduct-remedy caution.

8. Conclusion

Node scores answer a question an organisation can act on alone. Sovereignty is a property of the system, and the system's improvement is not the sum of improved nodes. The wedge $2\gamma \int_{c_{eq}}^{1} a\, dG$ is the classical statement of the gap in this model's units; the two ledger thresholds say exactly when it binds; the worked cases are its practice. The contribution criterion registers what node instruments structurally cannot, with receipts, stated counterfactuals, an explicit path discipline, and a sign lemma whose hypotheses are in the open. It is buildable today: the screens are in this paper, and the assessment format of Section 5 is repeatable on any conduct with a public record.

References

- Ascend Partners (2026). *DRI Referential, version 1.1 (Référentiel IRN).* CC BY-NC-ND 4.0.

- European Commission, DG Competition (2005). *Merger Remedies Study.* Public version, October 2005.
- European Commission (2025a). Commission accepts commitments offered by Microsoft to address competition concerns related to Teams. Press release IP/25/2048, 12 September 2025. https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2048.
- European Commission (2025b). *Cloud Sovereignty Framework*, version 1.2.1. DG DIGIT, 20 October 2025. https://commission.europa.eu/document/download/09579818-64a6-4dd5-9577-446ab6219113_en (accessed 2026-08-12).
- Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. In *Handbook of Industrial Organization*, vol. 3. North-Holland.
- Farrell, J., & Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics* 16(1), 70–83. https://doi.org/10.2307/2555589.
- Fermigier, S. (2026a). *Buyers' Clubs for the Open-Source Flow: Collective Provision in the FLOSS Value Network.* Draft working paper v0.8, 2026-09-07. Abilian Econ Lab. https://econ.lab.abilian.com/papers/collective-provision/
- Fermigier, S. (2026b). *Cheaper and Losing: A Layered Model of Cloud Dependence and the Limits of Subsidy.* Draft working paper v0.6, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/cheaper-and-losing/
- Fermigier, S. (2026c). *Counting Dependencies: Aggregation Rules and What Official Dependency Indices Identify.* Draft working paper v0.1, 2026-07-18. Abilian Econ Lab. https://econ.lab.abilian.com/papers/counting-dependencies/
- Fermigier, S. (2026d). *Cuts and Options: A Theory of Strategic Autonomy in Production Networks.* Draft working paper v0.10, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/theory/
- Fermigier, S. (2026e). *Exit or Own? Portability Mandates and Ownership Screens in Sovereign Procurement.* Draft working paper v0.7, 2026-09-16. Abilian Econ Lab. https://econ.lab.abilian.com/papers/exit-or-own/
- Fermigier, S. (2026f). *Leverage Without Additionality: What Co-Funding Ratios Identify in Public Technology Programmes.* Draft working paper v0.2, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/leverage-without-additionality/
- Fermigier, S. (2026g). *Path Sovereignty: Dependence, Control, and Options in the European Digital Stack.* Draft working paper v0.17, 2026-08-30. Abilian Econ Lab. https://econ.lab.abilian.com/papers/power-network/
- Fermigier, S. (2026h). *Scoring the Seller: An Economic Foundation for a Supplier-Side Sovereignty Index.* Draft working paper v0.7, 2026-09-16. Abilian Econ Lab. `rsi/papers/scoring-the-seller-wp.md` (internal)
- Fermigier, S. (2026i). *Screening for Sovereignty: Optimal Procurement when Compliance is Cheap to Fake and Costly to Verify.* Draft working paper v0.16, 2026-09-16. Abilian Econ Lab. https://econ.lab.abilian.com/papers/screening-for-sovereignty/
- Fermigier, S. (2026j). *Sovereignty as a Flow: Capture and the Durability of European Technology.* Draft working paper v0.5, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/capture/
- Fermigier, S. (2026k). *The Unpriced Bundle: Decision Suppression in Software Procurement.* Draft working paper v0.1, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/unpriced-bundle/
- Fermigier, S. (2026l). *A Two-Dimensional Taxonomy of Cloud Service Integrations: Modeling Architectural Coupling and Portability Friction in Managed Infrastructures.* Draft working paper v1.1, 2026-08-30. Abilian Econ Lab. https://econ.lab.abilian.com/papers/cloud-services-matrix/

- Fermigier, S. (2026m). *Value Capture in Commercial Open Source: Services, Funding Paths, and the FLOSS Value Network (Europe, 2026).* Draft working paper v0.6, 2026-09-03. Abilian Econ Lab. https://econ.lab.abilian.com/papers/value-capture/
- Fermigier, S. (2026n). *What Moves the System: A Factor Map and Ranking for European Digital Sovereignty.* Draft working paper v0.5, 2026-09-18. Abilian Econ Lab. https://econ.lab.abilian.com/papers/factor-map/
- Fermigier, S. (2026o). *Why Driver Weights Differ: A Cost-Based Microfoundation for Segmented Technology Acceptance.* Draft working paper v0.3, 2026-09-02. Abilian Econ Lab. https://econ.lab.abilian.com/papers/acceptance-structure/
- Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–440.
- Kwoka, J. E. (2015). *Mergers, Merger Control, and Remedies: A Retrospective Analysis of U.S. Policy.* MIT Press.
- Mayne, J. (2001). Addressing attribution through contribution analysis: using performance measures sensibly. *Canadian Journal of Program Evaluation* 16(1), 1–24. https://doi.org/10.3138/cjpe.016.001.
- Mayne, J. (2019). Revisiting contribution analysis. *Canadian Journal of Program Evaluation* 34(2), 171–191. https://doi.org/10.3138/cjpe.68004.
- Nalebuff, B. (2004). Bundling as an entry barrier. *Quarterly Journal of Economics* 119(1), 159–187. https://doi.org/10.1162/003355304772839551.
- Quirk, J. & Ruppert, R. (1965). Qualitative economics and the stability of equilibrium. *Review of Economic Studies* 32(4), 311–326. (citation to be verified)
- Rohlfs, J. (1974). A theory of interdependent demand for a communications service. *Bell Journal of Economics and Management Science* 5(1), 16–37. https://doi.org/10.2307/3003090.
- Sovereign Tech Agency (2026). Programme pages: funded technologies and selection criteria. https://www.sovereign.tech/ (citation to be verified; figures quoted via *Leverage Without Additionality* §6.5).
- VanderWeele, T. J. & Robins, J. M. (2010). Signed directed acyclic graphs for causal inference. *Journal of the Royal Statistical Society: Series B* 72(1), 111–127. (citation to be verified)
- Whinston, M. D. (1990). Tying, foreclosure, and exclusion. *American Economic Review* 80(4), 837–859.

Appendix A: Proof of Proposition 1

Write $I_1(c) = \int_c^1 a\,dG$ and $I_0(c) = 1 - G(c)$, so that $W(c) = B(m(c))\,I_1(c) - \pi_C\,I_0(c)$ with $m(c) = 1 - G(c)$ and $B(m) = \pi_C + \pi_D + \gamma(2m - 1)$. Differentiating, $m'(c) = -g(c)$, so $\frac{d}{dc}B(m(c)) = -2\gamma\,g(c)$, while $I_{(1)'}(c) = -c\,g(c)$ and $I_{(0)'}(c) = -g(c)$. Hence

$$W'(c) = -2\gamma\,g(c)\,I_1(c) - c\,g(c)\,B(m(c)) + \pi_C\,g(c) = -\,g(c)[\,c\,B(m(c)) - \pi_C + 2\gamma\,I_1(c)\,].$$

At an interior equilibrium cutoff, $c_{eq}B(m(c_{eq})) = \pi_C$ by definition, so the bracket's first two terms cancel and $W'(c_{eq}) = -2\gamma\,g(c_{eq})\,I_1(c_{eq})$, which is strictly negative whenever $\gamma > 0$ (an interior cutoff with $g > 0$ gives $I_1 > 0$ automatically), and zero when $\gamma = 0$.

The global step. Local decrease at c_{eq} does not by itself locate the global maximizer, and $W' > 0$ regions above c_{eq} can exist in this model, so the exclusion of maximizers above c_{eq} needs an argument. For any $c > c_{eq}$: $m(c) < m(c_{eq})$, and B is increasing in m, so

$$W(c) = \int_c^1 \big(a\,B(m(c)) - \pi_C\big)\,dG \le \int_c^1 \big(a\,B(m_{eq}) - \pi_C\big)\,dG < \int_{c_{eq}}^1 \big(a\,B(m_{eq}) - \pi_C\big)\,dG = W(c_{eq}),$$

the last strict inequality because the integrand at m_{eq} is strictly positive on $(c_{eq}, 1)$ (for $a > c_{eq}$, $a\,B(m_{eq}) > c_{eq}B(m_{eq}) = \pi_C$) and $g > 0$. So every $c > c_{eq}$ is strictly dominated by c_{eq}, which is itself strictly dominated by points just below it (local decrease). The maximizer c_p

over $[0,1]$ therefore satisfies $c_p < c_{eq}$ strictly, interior or at the corner $c_p = 0$; under equilibrium multiplicity the argument applies verbatim at each interior equilibrium cutoff. Then $m(c_p) > m(c_{eq})$ because G is strictly increasing, and $W(c_p) > W(c_{eq})$.

For $\gamma = 0$, B is constant, $W'(c) = -g(c)(cB - \pi_C)$ vanishes exactly at $c = \pi_C/B = c_{eq}$, is positive below and negative above, so $c_p = c_{eq}$.

The ledger extension. $\widetilde{W}(c) = W(c) + u_0 G(c)$ gives $\widetilde{W}'(c_{eq}) = g(c_{eq})(u_0 - 2\gamma I_1(c_{eq}))$, negative iff $u_0 < 2\gamma I_1(c_{eq})$: a local statement at the equilibrium margin. At the corner, $\widetilde{W}'(0) = -g(0)\,(2\gamma\,\mathbb{E}[a] - \pi_C - u_0)$, so the corner remains a local optimum, and (by the same monotone-comparison argument) the global one, for $u_0 < 2\gamma\,\mathbb{E}[a] - \pi_C$, and the planner leaves it above that threshold. Both thresholds are verified numerically on both sides. ∎

Appendix B: Proof of Lemma 1

Parameterize the conduct by $\theta \in [0,1]$, scaling the footprint's factor perturbations from the counterfactual state ($\theta = 0$) to the actual state ($\theta = 1$), and let $x(\theta)$ be the factor vector at the fixed point on the stable branch of hypothesis (d). The branch's continuity and differentiability follow from the implicit function theorem at each θ: the spectral radius condition $\rho(J(\theta)) < 1$ keeps $I - J(\theta)$ invertible, and the branch condition excludes folds. Then $\frac{dx}{d\theta} = (I-J)^{-1}\delta$, where δ is the direct perturbation vector supported on the footprint's factors (hypothesis (a)). Expand $(I-J)^{-1} = \sum_{n\geq 0} J^n$, convergent by $\rho(J) < 1$: each entry of $\frac{dx}{d\theta}$ is a sum over directed walks from perturbed factors, each walk contributing the product of its edge weights. Every walk decomposes into a simple path plus a collection of traversed cycles. By hypothesis (c) every cycle's sign product is positive, so each walk carries the sign of its underlying simple path, and by hypothesis (b) every simple path into the outcome, composed with the outcome gradient's stated signs, carries the sign σ, with at least one strict. Hence $\frac{dV}{d\theta} = \nabla V \cdot \frac{dx}{d\theta}$ has sign σ for every $\theta \in [0,1]$, and $C(E) = V(x(1)) - V(x(0)) = \int_0^1 \frac{dV}{d\theta}\,d\theta$ has sign σ. The argument holds for any counterfactual state that preserves (a)–(d), which is Rule 2′s sign-robustness clause. In the assembled model, the elasticity computations instantiate exactly this structure, with $D > 0$ as the scalar case of $\rho(J) < 1$. ∎

Annex A: Statements of the results the screens lean on hardest (précis)

The full précis annex for all top-tier edges is the companion paper's Annex A; the three results this paper's screens and cases lean on hardest are restated here for self-containment. **Screening (*Screening for Sovereignty*, Thm 2):** when a tender verifies attributes both genuine and controlled suppliers can hold at similar cost (residence, staffing, certifications) and leaves effective control unexamined, no verification accuracy separates the types. **Bundling (*The Unpriced Bundle*, Propositions 1–3, checks B1–B3):** zero-priced bundling suppresses adoption on a characterized set of buyers (initiation requires clearing a wedge $w = rk/\theta$); subsidies below a suppressed buyer's gap are inert; itemisation restores the adoption margin for every buyer simultaneously. **Co-funding (*Leverage Without Additionality*, P1):** measured co-funding per unit of public contribution equals $\sum_j \omega_j(1-s_j)/s_j$, a deterministic function of the rate schedule; at a full-funding rate leverage is exactly zero by construction, so the statistic identifies schedule composition, never crowd-in.