What Moves the System: A Factor Map and Ranking for European 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corpus now carry the published addresses of the papers they name, with the stale version pins corrected and the one unpublished source flagged in place. The canonical edge list was re-synced to the paper after Section 5 reclassified L2 as the capture chain; the four verification scripts now declare their own dependencies. A prose pass brought the document's coordination rate inside the corpus's own range. Revised after two referee rounds, the second by an independent reviewer given only the papers; v0.4 adds the first executed bindingness measurement (Section 7). The edge list is canonical in `verify/edges.py`: the table and the figure are generated from it and asserted equal at render time, conjectured edges included. The evidence tiers grade internal verification, the ranking is presented as a reasoned prioritization with tiered premises. Annex A states the results behind the top-tier edges, so the map is readable as mathematics without the source corpus.

Abstract

The instruments deployed for digital-sovereignty policy score nodes: an organisation's resilience, a supplier's posture, an offer's compliance level. The published indices describe levels. Neither says what makes the sovereignty of the whole ecosystem move. This paper assembles a causal factor map for the European digital ecosystem from one research programme's verified results: twenty-two factors in five families and a canonical edge list of thirty-two directed edges (twenty-three receipted, nine conjectured, two interactions), with a stated evidence tier per edge, six feedback loops, and one directed chain (the capture pull). The tiers grade internal verification; external validation is the map's open flank. Annex A states the results behind the top-tier edges; every edge names what an attack would have to break.

Confronting the map with the deployed instruments yields two findings. First, every substantive criterion of the node instruments corresponds to a system factor, a correspondence whose construction failed twice and forced factor additions, which is its falsifiability credential. Second, conditional on the programme's own selection of what to formalize, the Cloud Sovereignty Framework's published weight vector, confirmed against the primary text, carries no economic objective at all; the official objective sets fold the market mechanisms away.

The paper then argues a prioritization at the current margin: a gate-and-driver structure extracted from the corpus's inertness results, ordinal, with premises tiered T2 and T3 and rank-by-rank comparisons that are explicitly heterogeneous (marginal effect where computed, implementability where documented, scope where proved). A model-scoped elasticity computation seeds the quantitative layer; its headline comparison holds on 59% of the sampled parameter space (85% of qualifying draws), with the excluded and complement regions characterized. The sequencing claim, clear the gates before spending on the drivers, is the paper's output, stated with its falsifiers and the measurements that would upgrade it. a coding of five years of sovereignty-labelled EU procurement notices finds structural control verified in 2 of 64 construct-relevant notices.

Keywords: digital sovereignty, industrial policy, causal factors, procurement, lock-in, open-source software.
JEL: F52, L86, O25, D85.

1. Introduction

Three questions partition the landscape of sovereignty measurement. *How sovereign is this node?* is answered by at least eleven deployed instruments (surveyed and crosswalked in

Scoring the Seller): client-side resilience referentials (the French DRI/IRN), supplier-side frameworks (the EU Cloud Sovereignty Framework and its national counterparts), and verification indices. *How sovereign is the system?* is answered descriptively, by country- and sector-level indices that report research shares, infrastructure counts, or stack control at a date. *What makes the system sovereign, by how much, and does a given policy or organisation help or harm?* goes unanswered. The conceptual literature supplies rationales without operational content (Edler et al. 2023 define technology sovereignty as an ability to act and derive policy frames from it), the evaluation literature supplies a method without a domain model (Mayne 2001). No deployed instrument scores causes.

This paper answers the first half of the open question: the factors, their evidence, and their prioritization. The companion (*The Contribution Criterion*, `contribution-criterion-wp.md`) answers the second half: scoring a given policy or organisation by its effect on the system.

The method is assembly. The programme behind this paper has spent its effort proving and measuring mechanisms one at a time: a layered market model of why price does not decide cloud adoption, a screening theory of what procurement can verify, an exit-pricing instrument scored on real services, a capture study of how control leaves, models of bundling, co-funding, open-source enclosure, and pooled provision, and a measured dependency graph of the European stack with a cut-and-option calculus on top. Each of those results is a causal edge with a receipt. This paper harvests them into one map, types every edge against the programme's channel taxonomy, states the evidence tier of every claim, and only then asks the prioritization question.

A receipt, throughout, means: the claim's proof or measurement lives in the cited source together with its verification suite. The sources are the author's working papers. The map's epistemic status is exactly theirs: internally verified, awaiting external challenge, as Section 2′s tier definitions say explicitly. Two mitigations narrow the gap between internal and inspectable. Annex A states the results behind every top-tier edge, so the edge list is readable as mathematics without access to the corpus. The verification suite that generates the table, the figure and every quoted number is part of the repository, intended as supplementary material. The map is offered as a target as much as a synthesis: every edge names the result an attack on it would have to break.

The contributions follow in order of appearance. First, the map itself: twenty-two factors in five families, a canonical list of thirty-two directed edges with per-edge tiers and stated edge types, six loops and one chain (Sections 4 and 5). Second, the coverage finding: every substantive criterion of the deployed node instruments corresponds to a system factor, with the construction's two documented failures as its falsifiability credential (Section 6). Third, the inversion finding, stated conditionally (Section 6). Fourth, the prioritization: a gate-and-driver structure extracted from the corpus's inertness results, with tiered premises and explicitly heterogeneous comparisons, and a model-scoped cardinal seed (Sections 7 and 8). Fifth, the falsifiers, which Section 9 states.

Scope. "European" is the default and the bloc is a parameter: the map's structure carries over to any bloc; only the receipts and calibrations are Europe-specific.

2. The outcome variable and the discipline

The outcome. The system's sovereignty is read as a stock of options: the cut-and-option calculus of *Cuts and Options* (coercion and diversification as dual quantities on an AND/OR production structure), evaluated empirically on the measured multi-layer dependency graph of *Path Sovereignty*. This paper defines no new composite index. A new index would assert dimensions and weights, the very assertion problem the corpus's measurement papers diagnose in the deployed frameworks (*Scoring the Seller*; *Counting Dependencies*). Inheriting

the outcome keeps the factor map accountable to an object that already has a formal definition and a measurement.

Evidence tiers. Every edge carries a tier. The tiers grade *internal* verification and say nothing about external validity. T1: proved and numerically verified within the programme (an analytical statement with a written proof, a numeric cross-check, and a Monte-Carlo robustness report). T2: measured or empirically scored by the programme, with sources and stated limits, identification possibly open. T3: documented case or literature-sourced claim. T4: conjectured, no receipt. A slashed tier T_a/T_b records the mechanism's tier over the applicability record's tier: the mechanism is established at T_a and its application to the European case documented at T_b. For inheritance and for any counting, the edge's operative tier is the weaker (right-hand) value; where this paper counts "T1-mechanism edges" it says so explicitly. A composite claim inherits its weakest operative tier. No tier claims peer review or out-of-corpus replication; edges established in the external literature would enter at T3 until re-verified under the programme's gates.

Loops and identification. The factor graph is cyclic; its cycles carry content. Wholesale causal-graph identification is therefore unavailable: the do-calculus guarantees hold on acyclic graphs (Pearl 2009), cyclic structural models lose the guarantees that make graphs sufficient objects (Bongers et al. 2021). The programme's measurement work applies intervention semantics as perturbation-then-recompute on the cyclic structure, reserving identification claims for carved acyclic sub-questions. Three such sub-questions are marked in Section 5.

3. The factors

Supply side. F1 *Integrated supplier breadth*: existence of domestic suppliers at the integrated breadth buyers ask for, against components at price parity (the assembly gap). F2 *Supplier scale and reinvestment*: domestic providers' revenue base and its growth relative to the market. F3 *OSS producer sustainability*: whether open producers can capture value without enclosing. F4 *Commons maintenance funding*: public and pooled funding of maintenance the market undersupplies.

Demand side. F5 *Buyer altitude*: the distribution of procurement evaluation capability across buyers; a buyer's altitude a measures how much of an offer's total cost of ownership and risk it can price beyond the sticker; $a^*(m)$ is the altitude threshold above which a buyer evaluates alternatives rather than defaulting to the incumbent, at incumbent consistency level m. F6 *Demand aggregation*: pooled buying through clubs, anchor buyers, and code dues. F7 *Decision-event generation*: whether procurement produces genuine adoption decisions at all, against suppression by zero-priced bundling. F8 *Screening capacity*: procurement's ability to verify effective control rather than checkable proxies.

Friction and structure. F9 *Exit friction*: per-service exit taxes, architectural coupling, egress and migration costs. F10 *Exit-rights infrastructure*: portability mandates and tested exit paths as installed policy. F11 *Option supply and disjointness*: substitutable options per layer, controller-disjoint at depth, forkable where open. F12 *Measurement integrity*: whether the indices and criteria steering policy identify what they claim to.

Control and governance. F13 *Control retention*: whether control of domestic firms stays domestic through M&A, private equity, and infrastructure deals. F14 *Procurement-channel integrity*: resistance of sovereignty-labelled procurement to washing.

Cross-cutting. F15 *Jurisdictional exposure*: extraterritorial legal reach along dependency paths, one channel among fifteen. F16 *Adopter capability*: the skill and capacity of adopting organisations, which reweights every adoption driver. F17 *Talent*, F18 *Standards authorship*, F19 *Capital-market structure*, F20 *Data gravity and compute capacity*, F21 *Narrative and expectations*, F22 *Rule shaping*: conjectured (T4), retained because the channel taxonomy requires them (Section 6), and unranked for that reason.

4. The edges

The canonical edge list lives in `verify/edges.py`: the table below and Figure 1 are both generated from it, conjectured edges included, with the renderer asserting their equality. Thirty-two directed edges: twenty-three receipted, nine conjectured (one core, eight from the conjectured factors), two of the receipted edges being interactions (marked ×). "OUT" is the outcome.

The edges are typed; the types matter for what can be computed on the map. Most edges are signed causal influences. Two are interactions (a joint effect of two factors on a third). One is a moderator (F12 → OUT: measurement integrity moves the outcome through the quality of instrument choice, never through the market). One is a reweighting (F16 → F5, signed ±: capability shifts the altitude distribution without moving it monotonically). F11 → OUT is partly definitional, since the outcome is defined as option value over F11′s object; its receipt is the calculus content (path-level properties that per-layer certificates cannot certify), never the tautological part. The companion paper's sign-transfer lemma operates only on the well-signed causal subgraph, excluding the moderator and reweighting edges; this is stated there as an explicit restriction.

Edge	Sign	Mechanism	Tier	Source
F5 → F2	+	below altitude $a^*(m)$ buyers default to the consistent incumbent; the share is a fixed point with a proved floor	T1	*Cheaper and Losing* P1–P7, Cor. 7b
F2 → F5	+	the share feeds the consistency term $B(m)$: a higher share lowers every buyer's threshold (loop L1′s return edge)	T1	*Cheaper and Losing* P1–P7
F1 → F5	+	the assembly gap sets the consistency premium; component price parity alone does not move the threshold	T1/T2	*Cheaper and Losing* §1, Facts 3–4
F2 → F1	+	scale funds breadth (reinvestment closes the assembly gap); conjectured, no corpus receipt	T4	conjectured
F7 → F2	+	zero-priced bundling suppresses decision events on a characterized set; sub-gap price advantages are inert; itemised pricing and display duties restore events	T1/T3	*The Unpriced Bundle* B1–B3
F7 → F9	−	each suppressed decision deepens the installed base and the next decision's exit tax (loop L3); a composition of two verified halves, tiered as the composition	T3	*The Unpriced Bundle*; *Cloud Services Matrix*
F9 → F5	−	exit taxes and coupling raise the effective price of every alternative, scored per service	T2	*Cloud Services Matrix*; migration guides
F9 → F7	−	accumulated coupling suppresses future decision events	T2	*Cloud Services Matrix*; *The Unpriced Bundle*
F10 × F8 → F14	+	exit mandates destroy lock-in rents in the mimic's prize,	T1	*Exit or Own?* P1–P2

Edge	Sign	Mechanism	Tier	Source
		lowering the accuracy screening must reach: strict complements in a band (interaction)		
F10 → F9	−	installed exit rights (portability mandates, the egress phase-out) remove measured components of the exit tax	T2/T3	*Exit or Own?*; *Cloud Services Matrix* (egress component)
F8 → F14	+	gap-less attributes cannot screen; sparse structural requirement sets on control can	T1	*Screening for Sovereignty* Thm 2–3, Cor. 2
F14 → F2	+	washing prices the pool at the lemon level and diverts sovereignty procurement from genuine suppliers	T1/T3	*Scoring the Seller*
F6 → F3	+	pooled demand and in-kind dues sustain open producers	T1/T3	*Buyers' Clubs* P1–P5
F6 → F4	+	an anchor of sufficient mass eliminates the zero-participation equilibrium	T1	*Buyers' Clubs* P3
F3 → F11	+	an enclosure threshold in the funding-path parameter; fork capability and copyleft deter enclosure	T1	*Value Capture* P1–P4
F4 → F11	+	maintenance funding keeps critical components alive; full-funding designs make leverage statistics inapplicable by design	T1/T2	*Leverage Without Additionality* P1, §6.4–6.5
F2 → F11	+	domestic supplier scale is option supply at its layers, measured on the graph	T2	*Path Sovereignty*
F2 → F13	−	a subsidised champion becomes an acquisition target: scale attracts capture attempts (the capture chain's forward edge)	T3	*Sovereignty as a Flow*; the programme's loop register
F13 → F11	+	control retention keeps the option domestic; capture events (M&A, private equity, infrastructure deals) remove it, and the capture chain makes subsidy a weak lever	T2/T3	*Sovereignty as a Flow*
F8 × F16 → F13	+	screening mode matched to the dependency's structure protects control (interaction)	T3	*Two Screenings*
F16 → F5	±	driver weights differ by capability segment; capability reweights the altitude distribution	T1 (theory)	*Why Driver Weights Differ*
F12 → OUT	+	moderator on instrument choice: aggregation-faithful measurement steers allocation, and official indices can invert rankings	T1	*Counting Dependencies* P1–P3

Edge	Sign	Mechanism	Tier	Source
F11 → OUT	+	the cut-and-option calculus prices coercion and diversification on option sets; per-layer certificates cannot certify path-level properties	T1/T2	*Cuts and Options*; *Path Sovereignty*; *An Open Internet Stack*
F15 → OUT	−	typed jurisdictional-reach edges propagate along paths that per-layer tests do not inspect	T2	*Path Sovereignty*; *Channels of Power*
F17 → F5	+	conjectured channel-coverage edge (Section 3)	T4	conjectured
F17 → F3	+	conjectured channel-coverage edge (Section 3)	T4	conjectured
F21 → F5	±	conjectured channel-coverage edge (Section 3)	T4	conjectured
F18 → F9	−	conjectured channel-coverage edge (Section 3)	T4	conjectured
F18 → F11	+	conjectured channel-coverage edge (Section 3)	T4	conjectured
F20 → F9	+	conjectured channel-coverage edge (Section 3)	T4	conjectured
F19 → F13	+	conjectured channel-coverage edge (Section 3)	T4	conjectured
F22 → F14	−	conjectured channel-coverage edge (Section 3)	T4	conjectured

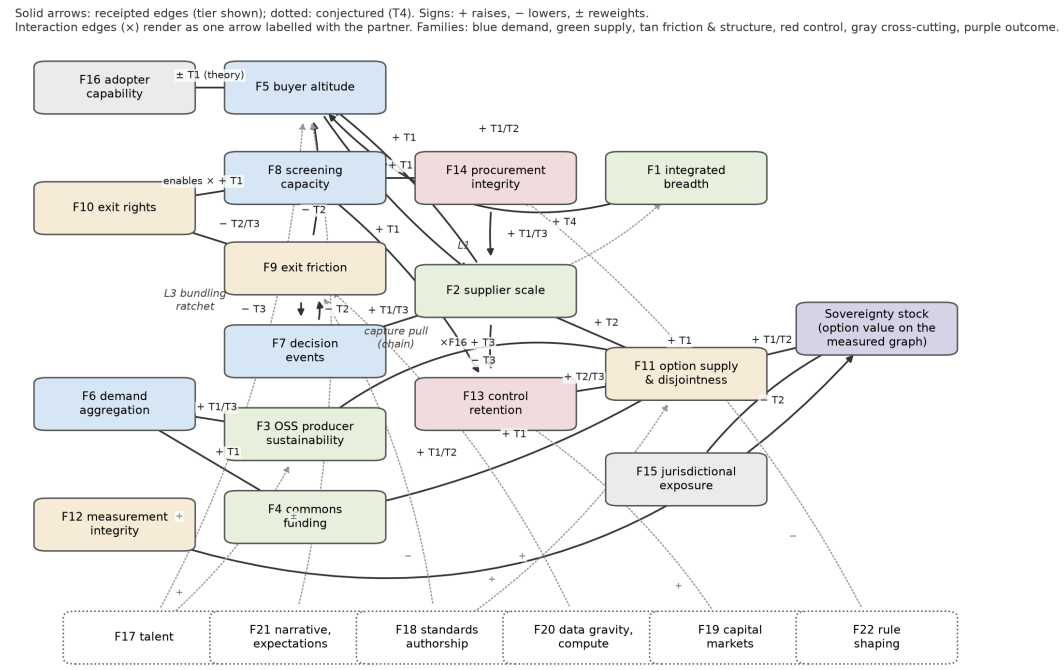

Figure 1: the factor map, generated from the canonical edge list (`verify/render_map.py` asserts figure-table equality, conjectured edges included). Solid arrows are receipted edges, labelled with sign and evidence tier; dotted arrows and boxes are conjectured (T4, signs shown); interaction edges (×) render as one arrow labelled with the partner factor. Families are color-coded (legend in the image); the loop badges mark L1 and L3; the capture pull is a directed chain (Section 5).

5. Loops and the capture chain

The register holds seven feedback structures; Six of the seven are loops on the current map and one is a chain; three of the seven carry receipts.

L1, the consistency loop (F5 ⇄ F2): the share feeds every buyer's consistency term and the threshold feeds the share; the fixed point $m = 1 - G(a^*(m))$ *is* this two-edge loop, proved with its segment floor. A slower route (F2 → F1 → F5: scale funds breadth, breadth lowers the premium) would reinforce it; its first edge is conjectured and the route is marked T4.

The capture pull (F2 → F13 → F11) is a directed chain, named as one here: no receipted edge returns from options or retention to scale, so the "return to dependence" it produces is a reversion of the system's state, never a cycle in the causal graph. The reclassification leaves its policy content intact: scale bought without retention leaks, which makes subsidy a weak lever at the chain's timescale. A closing edge (F11 → F2: richer domestic option supply feeds domestic scale) is plausible and currently unreceipted; if established it would close the chain into a loop. It sits on the conjecture ledger.

L3, the bundling ratchet (F7 ⇄ F9): each suppressed decision deepens the installed base and raises the next decision's exit tax; the loop's tier is the composition's (T3), even though its edges are separately verified.

L4 to L7 (talent, data gravity, standards capture, and the demand chicken-and-egg at bloc scale) are conjectured, with a scale caveat: the club-scale resolution of the chicken-and-egg loop is proved (*Buyers' Clubs* P3); the bloc-scale analogue does not inherit that tier, because the scale jump is itself a claim.

Identification is available where an acyclic sub-question supports a causal claim over a bounded horizon: instruments to decision events to adoption within one procurement cycle (the Teams commitments give a before-and-after event structure); capture events to control state to option count (dated registry events); and the Data Act egress phase-out of January 2027 as a natural experiment for exit rights to friction to migration.

6. Coverage: the map against the deployed instruments

The deployed node instruments were mapped onto the factors. The DRI/IRN mapping is at pillar level with criterion identifiers grouped. The criterion-by-criterion table is owed to the instrument release. For the French DRI/IRN referential (v1.1: eight pillars, thirty criteria, per-asset maturity self-assessment): the Strategic pillar maps to F5 and F16, Economic and Legal to F9, F13, and F15, Data and AI to F11 and F15, Operational to F11, Supply Chain to F11 and F12, Technological to F9 and F3, while Security and Environmental fall outside the sovereignty outcome as a declared boundary.

Finding 1 (node shadows). Every substantive criterion of the deployed node instruments corresponds to a system factor: the instruments and the map describe the same physics at different scopes. The obvious objection is that a 22-factor set maps everything by sheer size. The construction's own record answers it: the correspondence was built as a two-way type-check against the programme's channel taxonomy, and that check *failed twice* during construction, finding two channels (narrative and expectations; rule shaping) with no factor, which forced F21 and F22 into the list as documented coverage bugs. A procedure that can fail, and did, is falsifiable; a criterion requiring a factor outside the current set would be a third failure and would be recorded the same way. The consequence stands: the question "does improving node scores improve the system?" is well posed, because both sides are expressed over the same factor set, which is the question the companion paper analyzes.

Finding 2 (the inversion, conditional). Call the *market-mechanism regrouping* the factors F1, F5, F7, F9, F14: consistency premia, altitude, decision suppression, exit taxes, and washing. The regrouping cuts across four of the five families and is defined here, for this finding. Of the edges whose mechanism tier is T1 (the slashed-tier convention of Section 2), a

clear majority touch this regrouping; the official objective sets represent exactly these factors least. The premise is now confirmed against the primary text (European Commission 2025, read 2026-08-14): the framework's eight Sovereignty Objectives and their published scoring weights are Strategic 15%, Legal and Jurisdictional 10%, Data and AI 10%, Operational 15%, Supply Chain 20%, Technology 15%, Security and Compliance 10%, Environmental 5%, with no economic objective in the vector; vendor lock-in appears only as contributing factors inside the Operational and Technology objectives ("without vendor lock-in"; "reducing dependency on single vendors"), which is the folding this finding names. One condition remains: the count of deeply verified edges is conditioned on the programme's own selection (this corpus formalized economic mechanisms first, so the depth of verification partly measures the research agenda). Even under the condition, the finding still has content: no first-class weight attaches to rents, prices, decision events, or exit costs in the official vector. The mechanisms in question are the ones a procurement authority can act on.

7. The prioritization: gates before drivers

What is being ranked. This section orders factor groups by where a marginal unit of policy effort moves the outcome most at the current margin. The comparisons behind the order are heterogeneous: marginal effects are used where the model computes them (Section 8), implementability and fiscal cost where the record documents them, and scope where a result proves its own boundaries. The homogeneous version of this ranking, the effect on V per euro of a named instrument over a named horizon, is a different and better object; it has a corpus precedent at instrument-family level (*An Open Internet Stack*) and is queued for the quantitative layer. Until then, this section is a reasoned prioritization with tiered premises, never a theorem.

The structure comes from the corpus's sharpest results, which are inertness results: regions where a factor's marginal effect is provably near zero because another factor gates it. Price advantages below a computed wedge do not trigger adoption events. Component price parity does not move the share when integrated breadth is absent. Public co-funding at a full-funding rate has leverage exactly zero by construction. Sovereignty procurement without control verification selects the lemon. Scale bought without control retention leaks through the capture chain. These results share a structure. **Gates** null or dissipate the contributions of the factors behind them; **drivers** carry continuous marginal effects once the gates pass. The kinship with growth diagnostics is direct (Hausmann, Rodrik and Velasco 2008): identify the binding constraint, sequence accordingly, expect the order to change as constraints clear. The reason marginal driver spending disappoints behind a binding gate is the general lesson of second best (Lipsey and Lancaster 1956).

The premise that the gates bind today is itself a claim, tiered per gate: decision suppression T3, washing T2, capture T3. Decision suppression is documented to commitment level in one large product market (the Teams case, 2025), an existence proof and no more. The washing gate's bindingness is measured. Coding all 157 sovereignty-term-matched TED notices in IT services (CPV 72*: every award notice 2022–2026 plus the hundred most recent calls) finds 42% irrelevant to the construct (buyers named after sovereignty ministries, sovereign bonds, sanctions boilerplate) and 17% procuring work about sovereignty. Among the 64 notices where sovereignty operates on the supplier or the solution, structural control verification appears in 2 (3%). The dominant operative mode is certification, label, or localisation (37 notices, 58%, twelve of them the French label family whose referential itself embeds control criteria since SecNumCloud 3.2); technical-autonomy requirements account for 12, another 11 invoke sovereignty with no operative requirement, and 2 defer all operative content to external tender documents. The record itself contains two accepted re-badgings (hyperscaler-powered products procured as the sovereign offer) and one sovereignty requirement withdrawn mid-procedure when qualified supply ran

thin. An independent second coder on a 20-notice overlap agreed on 18 and assigned both control codes independently. The count is a lower bound on requirements, since half the relevant notices reference external tender documents, while remaining a faithful measure of what buyers surface and verify at notice level. A screen of the software division (CPV 48*, 85 further notices) found one additional structural-control requirement, a CLOUD-Act eligibility exclusion by the German foreign ministry. T2: measured, stated limits, no identification claim. Capture has a documented flow whose identification the source itself holds open. The observed share plateau reads as the model's floor, the source's careful wording, inherited here. The prioritization inherits the weakest of these tiers. The remaining upgrades are specific: the share of procurement volume sitting inside bundled defaults, and the capture-event rate against the champion population.

The order at the current margin, each rank with its receipt and its evidentiary caveat:

1. **F7, decision-event generation.** The demand side acts only through decisions, so a suppressed decision nulls every other demand-side factor for that buyer and cycle. The restoring instruments (itemised pricing, display duties) are regulatory acts of low fiscal cost with a proved mechanism. *Caveat: the bindingness premise is one documented market.*
2. **F8 with F14, screening capacity and procurement integrity.** Whether decisions select genuine control decides whether demand-side spend produces sovereignty or subsidises its imitation. *Caveat: the mechanism is T1; no elasticity for these factors exists in the cardinal seed, and the verification measurement stays notice-level, with the external-documents blind spot stated above.*
3. **F13, control retention.** Gates the durability of supply-side gains. *Caveat: rests on the documented capture flow, identification open; no proved retention instrument exists in the corpus.*
4. **F10 with F9, exit rights and exit friction.** Exit mandates make ownership screens feasible in a characterized band; friction reduction raises every buyer's effective altitude. *Caveat: the friction magnitudes are scored on one provider's services; the mandate side's enforcement is the open margin.*
5. **F5 with F16, buyer altitude and adopter capability.** Moves the share fixed point directly and reweights every driver. *Caveat: slow to build and hard to buy; the capability half is theory plus a survey design without data yet.*
6. **F6, demand aggregation.** Sharp where coordination fails completely: the anchor result eliminates the zero-participation equilibrium in its scope (club-scale pooled provision of open source). *Caveat: the scope is exactly that; bloc-scale extrapolation is a conjectured loop.*
7. **F1 with F2, integrated breadth and supplier scale.** The long-run object itself, capped at the current margin by the gates and by the capture chain. Rises sharply as gates clear; this rank is the sequencing claim's main content. *Caveat: the inertness results that cap it are model results applied through the bindingness premises above, one measured and two documented.*
8. **F3 with F4, OSS producer sustainability and commons funding.** Verified steady effect on option supply, subject to additionality discipline. *Caveat: magnitudes rest on one flagged programme record and club-scale results.*

Outside the market ordering: F12, measurement integrity, multiplies the quality of every instrument choice (official indices can invert rankings) and moves the outcome through allocation. F15, jurisdictional exposure, is managed through disjointness (F11′s composition), and ranking it as a direct instrument target would repeat the per-layer-certificate error. The conjectured tail (F17 to F22) is unranked: ranking it now would be assertion.

8. The cardinal seed

The prioritization is quantified where the assembled model reaches. The model, assembled from *Cheaper and Losing*'s primitives and stated in full in the companion paper's Section 3, has buyers of altitude $a \sim G$ adopting when $a\,B(m) > \pi_C + s$, with consistency term $B(m) = \pi_C + \pi_D + \gamma(2m - 1)$, an event-rate gate λ (the share of buyers whose decision event occurs, so $m = \lambda(1 - G(c))$ at cutoff c), additive friction s, and stability denominator $D = B - 2\gamma\lambda c\, g(c)$. It admits closed-form elasticities of the equilibrium share by the implicit function theorem, verified against numeric derivatives of the fixed point and under Monte-Carlo, at both the full-events and a suppressed baseline.

Three results. First, **the event-gate bound, decomposed.** The elasticity of the share with respect to the event rate is $1 + 2\gamma\lambda c\, g(c)/D$. The unit term is mechanical: the gate enters multiplicatively, so any factor parameterized as a gate would earn it. The content is the strictly positive increment $2\gamma\lambda c\, g(c)/D$, the consistency-loop feedback, which no non-gate parameter of the model earns. Rank 1′s quantitative support is the increment together with the magnitude comparison below, never the bound alone.

Second, **magnitudes, with the coverage stated first.** The Monte-Carlo samples 400 parameter draws; 53 fall outside the model's admissible parameter set ($B > 0$ requires $\pi_C + \pi_D > \gamma$), and 70 unravel to zero-adoption fixed points, at which no interior margin exists and marginal analysis is not the relevant object (the challenger's problem there is survival, the near-collapse regime below). The remaining 277 draws qualify, 69% of the sampled space and 80% of the admissible space. On them: at the source model's defaults with full events ($\lambda = 1$), the elasticities are events +1.36, assembly cost −0.38, quality +0.27, consistency +0.11; the friction effect is reported as the marginal effect $dm/ds = -1.27$ (a different unit from the elasticities, and excluded from the magnitude comparison for that reason). At a suppressed baseline ($\lambda = 0.9$, where "restoring events" has a referent), the ordering is unchanged and the event elasticity rises to +1.40. The event gate has the strictly largest elasticity magnitude on 85% of qualifying draws (236 of 277; 59% of all sampled). The complement is characterized: 40 of its 41 draws sit in the near-collapse regime (equilibrium share below 0.37), where the quality elasticity spikes as the fixed point approaches unravelling; one draw escapes the characterization. The two regimes read naturally: near collapse, survival is a quality problem; established, events dominate. Mapping the regimes onto actual market positions is calibration work and is left open.

Third, **the consistency sign flip.** The share's derivative in the consistency weight is proportional to $2m - 1$, a headwind for a minority challenger and a tailwind past parity, which is the consistency loop's content in one line.

The scope statement travels with the numbers: these are model elasticities at calibrated parameters, covering the factors the assembled model carries. The share is the outcome proxy; any outcome strictly increasing in it inherits the ordering up to its own slope, which the model does not pin by design. Elasticities compare percentage responses, yet one percent means different things across parameters with different policy costs; the per-euro ordering is the queued object (Section 7). The full-map elasticities (enclosure and provision margins, capture hazard, measured-graph option counts) require the measured compute core and the source models' own comparative statics, and are queued with it.

9. What would change the prioritization

The order is state-dependent; its falsifiers are part of its content. Enforced unbundling and display duties at scale would clear the first gate and promote the drivers behind it (F5, F1). Verification infrastructure reaching procurement practice would clear the second and promote breadth and scale. Evidence that capture pressure is lower than the registry suggests, or that retention instruments work, would demote rank 3. Bindingness measurements

(Section 7′s list) could overturn the remaining T3 premises directly; the first executed one, the procurement-verification share, landed on the binding side. The quantitative layer could re-order ranks 4 through 8, where the ordinal argument is tightest; ranks 1 through 3 rest on inertness results and would need those results to fail. A documented case of a suppressed-decision market in which sustained sub-wedge price entry nonetheless moved share would strike directly at rank 1′s receipt.

10. Related work

Node-scoped instruments are surveyed and crosswalked in *Scoring the Seller*; the French DRI/IRN referential (Ascend Partners 2026, v1.1) is the client-side archetype: eight resilience pillars scored per digital asset with maturity levels, keyed to the NIS2, DORA, GDPR, AI-Act, and Data-Act obligations. System-level descriptive indices include the ASPI Critical Technology Tracker (research-output shares across 74 technologies), the Critical-Infrastructure Sovereignty Index (Peter and Udenenwu 2026, sector-level control scores), and the Nextcloud Digital Sovereignty Index (observed self-hosted infrastructure). All describe levels and are silent on causes; none ranks factors or scores entities. The conceptual literature (Edler et al. 2023; the sovereignty-definition survey of Pohle and Thiel 2020) supplies the vocabulary of ability-to-act that the outcome variable formalizes.

Four further literatures border the exercise. Growth diagnostics (Hausmann, Rodrik and Velasco 2008) is the closest methodological relative of Section 7: binding constraints, sequencing, state-dependent re-ordering. The addition here is the receipted edge list under the constraint structure. The general-equilibrium caution behind "driver spending dissipates behind gates" is the theory of second best (Lipsey and Lancaster 1956). The new industrial-policy economics (Juhász, Lane and Rodrik 2024) supplies the evaluation discipline and the cautionary history for targeted intervention; the prioritization here is exactly the kind of claim that literature asks to see stated with falsifiers. The network-effects canon (Katz and Shapiro 1985; Farrell and Saloner 1985) underlies the assembled model's consistency externality; the companion paper carries the formal treatment and the lineage (including Rohlfs 1974). On the jurisdictional factor, Farrell and Newman (2019) is the standard reference for coercion through network positions, and the corpus's measured typed-reach edges are its operationalization at stack level. On method, Meadows's leverage-points hierarchy (1999) is the folk prior for the prioritization, with the notable convergence that rules and information rank above parameter spending; the systems-mapping tradition behind the loop register runs through Sterman (2000). Mayne's contribution analysis (2001; 2019) anchors the companion paper's scoring method.

11. Limitations

The map's evidence is one programme's corpus; the tier system grades internal verification only: no edge has yet faced external referees, replication, or out-of-corpus data. Every claim in this paper should be read under that ceiling. Annex A narrows the inspection gap (the results behind the top-tier edges are stated, so the mathematics is readable without the corpus); publication of the source papers is the remedy the annex cannot substitute for. Within the corpus, the tier system shows the dependence on sources: two edges rest on a paper whose own banner holds causal identification open, and are capped accordingly. The map's completeness claim is relative to the channel taxonomy, whose own completeness argument lives in its paper. The elasticities are model-scoped, and the assembly-cost elasticity mixes a cost with a level effect inherited from the source model's payoff form. The prioritization is local to the current margin by design, its premises are tiered (the washing gate measured at T2 with notice-level scope, the others T3), and its comparisons are heterogeneous by declaration. It

states where the system is on the current record, and re-orders as the system moves or the record improves.

12. Conclusion

The question "what should Europe do about digital sovereignty" decomposes into a factor map and a prioritization, both buildable from verified parts, with the verification's internal scope stated. The map shows the deployed instruments to be node-level shadows of system factors, by a construction that documented its own failures. The inversion finding says that, under two stated conditions, official objective sets weight least exactly where the programme's verification runs deepest. The prioritization says the current margin belongs to the gates on the current record: whether decisions occur, whether screens verify control, whether control stays. Spending on the long-run drivers while the gates bind is the corpus's best-documented way to dissipate effort. The sequencing claim, clear the gates first, is the paper's output, and Section 9 states its falsifiers.

References

- Ascend Partners (2026). *DRI Referential, version 1.1 (Référentiel IRN).* CC BY-NC-ND 4.0.
- ASPI (2026). *Critical Technology Tracker.* Australian Strategic Policy Institute. https://www.aspi.org.au/programs/critical-technology-tracker/ (accessed 2026-08-12).
- Bongers, S., Forré, P., Peters, J., & Mooij, J. M. (2021). Foundations of structural causal models with cycles and latent variables. *Annals of Statistics* 49(5), 2885–2915. https://doi.org/10.1214/21-AOS2064.
- Edler, J., Blind, K., Kroll, H., & Schubert, T. (2023). Technology sovereignty as an emerging frame for innovation policy. Defining rationales, ends and means. *Research Policy* 52(6), 104765. https://doi.org/10.1016/j.respol.2023.104765.
- European Commission (2025). *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, H., & Newman, A. L. (2019). Weaponized interdependence: how global economic networks shape state coercion. *International Security* 44(1), 42–79.
- 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). *An Open Internet Stack: What to Build, in What Order, and Why.* Draft working paper v0.7, 2026-09-07. Abilian Econ Lab. https://econ.lab.abilian.com/papers/tbb-criterion/
- Fermigier, S. (2026b). *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. (2026c). *Channels of Power in Production Networks: A Taxonomy for Measuring Technological Dependence.* Draft working paper v0.7, 2026-08-04. Abilian Econ Lab. https://econ.lab.abilian.com/papers/channels/
- Fermigier, S. (2026d). *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. (2026e). *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. (2026f). *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. (2026g). *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. (2026h). *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. (2026i). *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. (2026j). *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. (2026k). *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. (2026l). *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. (2026m). *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. (2026n). *Two Screenings: Why Europe Screened Huawei but not the Hyperscalers.* Draft working paper v0.7, 2026-07-16. Abilian Econ Lab. https://econ.lab.abilian.com/papers/two-screenings/
- Fermigier, S. (2026o). *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. (2026p). *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. (2026q). *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/
- Hausmann, R., Rodrik, D., & Velasco, A. (2008). Growth diagnostics. In N. Serra & J. E. Stiglitz (eds.), *The Washington Consensus Reconsidered: Towards a New Global Governance*, 324–355. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199534081.003.0015.
- Juhász, R., Lane, N., & Rodrik, D. (2024). The new economics of industrial policy. *Annual Review of Economics* 16, 213–242. https://doi.org/10.1146/annurev-economics-081023-024638.
- Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–440.
- Lipsey, R. G., & Lancaster, K. (1956). The general theory of second best. *Review of Economic Studies* 24(1), 11–32. https://doi.org/10.2307/2296233.
- 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.
- Meadows, D. H. (1999). *Leverage Points: Places to Intervene in a System.* Hartland, VT: The Sustainability Institute.

- Nextcloud (2026). *Digital Sovereignty Index.* https://github.com/nextcloud/Digital-Sovereignty-Index (accessed 2026-08-12).
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.
- Peter, A., & Udenenwu, O. (2026). Digital soil sovereignty: a comparative analysis of state cyber postures using the CISI framework. *Frontiers in Political Science* 8. https://doi.org/10.3389/fpos.2026.1771498.
- Pohle, J., & Thiel, T. (2020). Digital sovereignty. *Internet Policy Review* 9(4). https://doi.org/10.14763/2020.4.1532.
- 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.
- Sterman, J. D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World.* McGraw-Hill.

Annex A: Statements of the results the argument rests on (précis)

Compact statements of the results behind the edges whose mechanism tier is T1, so the edge list is readable without the source corpus. Each entry is a précis in this paper's words; the formal statements, proofs, and verification suites are in the cited sources.

Layered market (*Cheaper and Losing*, edges F5 → F2, F2 → F5, F1 → F5). Buyers of altitude a adopt the alternative when $a\,B(m) > \pi_C$, with $B(m) = \pi_C + \pi_D + \gamma(2m-1)$ increasing in the alternative's adoption mass m. P1–P7 (sorting; existence and lattice structure of equilibria; the tipping trap; commodity-subsidy saturation; un-bundling; the beachhead; the rising trend) establish the cutoff structure $a^*(m)$, the fixed point $m = 1 - G(a^*(m))$, and the subsidy-inertness results: price subsidies below the consistency premium do not move the fixed point. Corollary 7b (the segment floor): when terminal altitudes are bounded with distribution H, the incumbent's consistency measure rises to the interior floor $m_\infty = 1 - H(a^*(m_\infty)) < 1$, bounding the challenger's slide.

Bundling (*The Unpriced Bundle*, edge F7 → F2; the paper's Propositions 1–3, verified by checks B1–B3). P1 (suppression): in the bundle regime adoption requires a champion-initiated evaluation, $\Delta q - p_e \geq w$ with wedge $w = rk/\theta$, so a suppression region of buyers who would adopt under itemised pricing exists as a difference set. P2 (price inertness): for a suppressed buyer with gap g, adoption is invariant to every subsidy below g and flips at g. P3 (restored events): itemisation restores the separate-regime margin for every buyer simultaneously, recovering the suppressed surplus.

Screening (*Screening for Sovereignty*, edge F8 → F14). Theorem 2 (gap-less attributes cannot screen): when the tender verifies attributes both genuine and controlled suppliers can hold at similar cost (residence, staffing, certifications) and leaves effective control unexamined, no accuracy level separates the types. Theorem 3 (screening on control): with a verifiable control attribute, a type-cost gap, and price at the participation floor, separation is feasible. Corollary 2 (sparse structural sets): an attribute contributes to screening only through the product of its type-cost gap and its verifiability, so optimal requirement sets are small and structural.

Exit and screening complementarity (*Exit or Own?*, interaction F10 × F8 → F14). Lemma 1: the exit mandate carries no information about the supplier's type (both types comply identically), so it does not screen. P1 (the complementarity band): the mandate destroys lock-in rents in a mimic's prize, lowering the verification accuracy an ownership screen must reach; a non-empty accuracy band exists in which the screen is infeasible without the mandate and feasible with it. P2 (policy regions): with both fixed costs high and accuracy in the band, mandate-plus-screen is optimal and neither instrument alone is.

Washing (*Scoring the Seller*, edge F14 → F2). Sovereignty procurement without verification is a market for lemons: effective control is hidden, cheap to fake, and costly to verify, so the pool prices at the washed level and label premia divert spend from genuine suppliers; verification infrastructure restores the separating margin. (The paper's aggregation results are stated there with their own suite.)

Enclosure (*Value Capture*, edge F3 → F11). P1: the open firm's optimal contribution intensity has an interior closed form shifted by the funding path β. P2 (the enclosure threshold): if the service margin covers the appropriability gap, OPEN is optimal for every β; otherwise a threshold β^* exists below which the firm encloses. P3 (fork deterrence): β^* falls with fork capability and rises with switching costs and appropriability. P4: the enclosure region expands with appropriable service margins and maturity.

Pooled provision (*Buyers' Clubs*, edges F6 → F3, F6 → F4). P1 (participation partition): given a fee, firms join, pay in code or cash, or ride, on a distribution-free partition of the type space. P3 (the anchor): an anchor buyer of mass above a computable threshold eliminates the zero-participation equilibrium; in-kind dues extend the result. P4–P5: the rider band shrinks as the club grows, and compliance costs recruit.

Co-funding identification (*Leverage Without Additionality*, edge F4 → F11′s flagged half). P1 (schedule identity): under compliance, measured co-funding per unit of public contribution is $\sum_j \omega_j(1 - s_j)/s_j$, a deterministic function of the rate schedule; at a full-funding rate ($s = 1$) leverage is exactly zero by construction, so leverage statistics are inapplicable to full-funding designs and identify schedule composition, never crowd-in.

Aggregation (*Counting Dependencies*, edge F12 → OUT). P1: counting is not weakest-link (profiles exist where the count and the maximum disagree in ranking). P2: binary recorded scales make the weakest-link reading degenerate whenever one dimension is universally present. P3: netting counts does not compare bargaining positions when disruption costs are asymmetric. Jointly: official dependency indices can invert rankings relative to the objects they claim to measure.

Certification (*Cuts and Options*, edge F11 → OUT's calculus content). The coercion number (smallest coalition of outside controllers whose joint refusal halts operation) and the option number are dual path-level quantities; per-layer certificates of supplier plurality do not bound the coercion number, because controller overlap across layers defeats layer-by-layer plurality. Options are priced: forkable substrates are standing options under the buyer's control.

Capability reweighting (*Why Driver Weights Differ*, edge F16 → F5). Under a cost-based microfoundation, acceptance-driver weights differ across adopter capability segments, so capability shifts the effective altitude distribution, with no monotone adoption effect implied; the empirical half is a registered design awaiting data.

This annex is a reading aid: the précis carry no proofs, so the sources remain the objects of record.