Cheaper and Losing: A Layered Model of Cloud Dependence and the Limits of Subsidy

Stéfane Fermigier · Working paper v0.4 · 14 July 2026

Abstract

European cloud providers are several times cheaper than US hyperscalers for raw compute and win independent price-performance benchmarks, yet their share of their own market has roughly halved since 2017 while the three US firms hold about 70%. We explain this with a layered model in which each buyer has an *altitude*, the fraction of its needs that are differentiated (platform and AI services) rather than commodity infrastructure. Europe is cheaper and present in the commodity basement and absent, at competitive prices, in the differentiated layer where value is migrating; buyers sort by altitude, and because firms drift up-stack over their lifecycle the aggregate US share rises over time toward a corner. We show that subsidy is a weak lever at either layer: a commodity subsidy saturates over realistic budgets when firm types are distributed bimodally (only a subsidy scaling with the entire commodity base could restructure the market), while Europe's actual subsidies, which target the differentiated layer, meet a network constraint that capability funding alone does not clear. Interoperability's payoff is *un-bundling* rather than switching; and once cross-provider coupling costs are modelled, the realistic European position is a loosely-coupled *secondary* beachhead whose value is sovereign reversibility and the preservation of a supplier the demand-side screen presupposes. Every result is proved, cross-checked numerically, and stress-tested by Monte-Carlo.

Keywords: cloud computing, network externalities, switching costs, digital sovereignty, industrial policy, coordination.
JEL: L86, L13, D85, H57, L52.

1. Introduction

Europe's dependence on American cloud infrastructure is usually told as a story of inferiority: the hyperscalers are bigger, better, cheaper, further ahead, and European providers cannot keep up. Where the two compete directly, the story is false. For raw compute a European provider such as Hetzner is about six times cheaper than the equivalent Amazon, Microsoft, or Google instance (the European field as a whole about 2.5×; §3); for outbound bandwidth the gap runs from roughly tenfold to eightyfold, with several European providers charging nothing; and in independent price-performance benchmarks the European commodity providers sit, in one tester's words, "on a completely different tier." Yet over the same years in which this price advantage was plainly visible, European providers' share of their *own* regional market fell from about 29% (2017) to about 15% (2022) and has stagnated there, while Amazon, Microsoft, and Google together hold roughly 70% and the five largest cloud providers in Europe are all American (Synergy Research 2025; European Parliament 2025). At market scale the dependence is priced: 80 percent of European business spending on cloud software and services goes to US vendors, some 265 billion euros a year (Asterès 2025), and prices on that base rose 8.7 percent annually over the last three years (Asterès 2026). Cheaper, and losing.

The one-dimensional "who is better" axis cannot hold both facts at once. This paper replaces it with a two-dimensional picture organized by a single primitive. Cloud needs have an *altitude*. At the bottom sits commodity infrastructure (raw virtual machines, block storage, bandwidth), a near-homogeneous good on which European providers compete on price and win. At the top sit differentiated services (managed databases, analytics, serverless and platform services, and most recently managed AI and inference), where the value, the growth, and the margin have migrated, and where European providers are thin to absent. Between the layers runs a **scissors** (Figure 1): as one climbs the stack, Europe's price advantage falls to parity and then reverses (measured), Europe's already-minority presence thins further (a 15% infrastructure share concentrated in IaaS, under 5% in AI), and the value of the layer rises. Europe is cheapest exactly where it competes and absent exactly where the value goes.

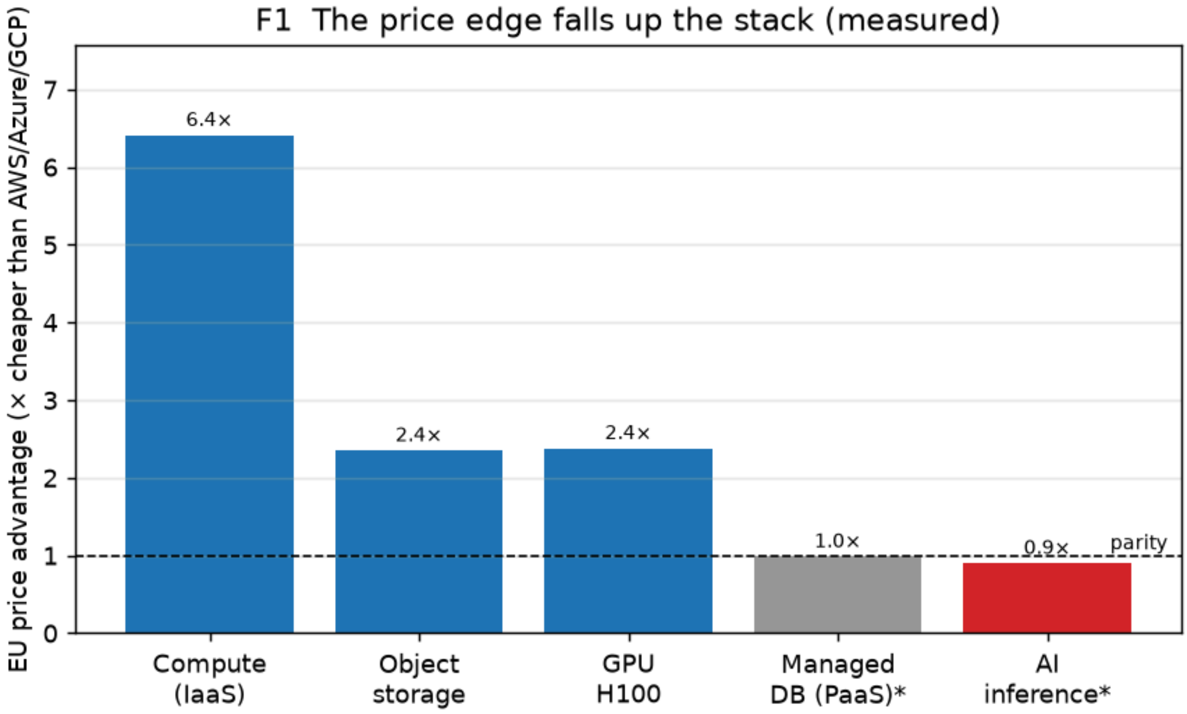

Figure 1. The price edge falls up the stack (measured). EU low-cost leaders (Hetzner, OVH, Scaleway) versus AWS/Azure/GCP: 6× cheaper on compute, 2.4× on object storage and H100 GPUs, at parity on managed databases, and slightly dearer on AI inference. Compute, storage, and GPU are geometric means computed from the getdeploying reference tables (2026, `pricing/`); the PaaS and AI points () are from BenchANT and AWS Bedrock vs Scaleway. Market presence, not published per layer, is in §3.*

We give each firm an altitude $a \in [0, 1]$ (the fraction of its needs that are differentiated) and let it choose one primary provider. A firm's net preference for the American over the European provider is increasing in its altitude, so firms **sort**: commodity-heavy firms take the cheaper European basement; differentiated-heavy firms take the American tower. The observed aggregate share is then a weighted count of firms above a threshold altitude, and four results follow, which are the paper's contribution.

First, dependence is a snapshot on a rising trajectory. Firms are born low-altitude and drift up-stack over their lifecycle: a start-up that begins with a virtual machine later adds a managed database, then analytics, then (recently) an AI service. As a growing population ages, and as the differentiated layer's pull strengthens, mass crosses the threshold and the US share climbs toward a corner. The observed 70% (the model's $m^* \approx 0.75$, a related but distinct object) is a point mid-ascent (Proposition 7, Figure 7).

Second, subsidy is a weak lever at both layers. A commodity subsidy is weak over realistic budgets: when firm types are bimodal, a moderate subsidy pushes the sorting threshold into the sparse middle of the altitude distribution, where there are few marginal firms to win, so the aggregate barely moves (Proposition 4); the subsidy that would flip the market exists, but its cost scales with the entire commodity base, and it only deepens an advantage Europe already holds where firms do not defect. Europe's real subsidies, a smaller strand of the agenda than the regulation, certification and procurement that dominate it, do not target that basement anyway: IPCEI-CIS's €1.2bn and France 2030′s €667M cloud line fund platform, edge, IoT and AI services, the differentiated layer where the value is. Targeting that layer is the right instinct, and the money largely funds capability there; but capability is not the binding constraint. The differentiated layer runs on network effects (the term γ below, and the tipping structure of Proposition 3): value accrues to the incumbent's installed base and ecosystem, so a capable but standalone European offering is still fought by the increasing returns that built the trap. What lowers γ is interoperability, which lets a European service plug into the prevailing ecosystem; a subsidy that leaves γ standing moves neither layer.

Third, interoperability's payoff is un-bundling the stack. Portability will not make a differentiated-heavy firm abandon the American AI layer it needs; that firm uses whoever is best there, and Europe is not. What interoperability does is let the firm stop putting its *whole* stack on one provider, reclaiming for Europe the commodity and data spend of firms whose core is American (Proposition 5). Once we model why stacks are not freely split (cross-provider data movement costs egress and latency), this becomes the realistic European position: a loosely-coupled **secondary** provider holding backups, disaster recovery, and non-mission-critical workloads (Proposition 6). This is where egress fees actually bite, as a tax on running a European secondary.

Fourth, that beachhead is the bridge to the demand side. Its worth exceeds the few points of market share it recovers. It provides sovereign reversibility (a live European copy is the exit option of the companion *Exit or Own*) and the preservation of a European supplier at all, which the companion *Screening for Sovereignty* presupposes must exist before any procurement screen can select it.

What we do not claim. Not that European technology is worse: in the commodity layer it is cheaper. Not that the model wins the AI layer: it cannot, without Europe becoming genuinely competitive there, which is the hard prescription the model yields rather than evades. And not identification: the model *rationalizes* the record and states what would refute it (§10), though the record is consistent with simpler stories and we say which.

Roadmap. Section 2 positions the paper. Section 3 lays out the sourced empirical base. Section 4 states the model. Section 5 proves the equilibrium results; Section 6 the leverage results; Section 7 the coupling variant and the beachhead; Section 8 the dynamics. Section 9 documents the verification. Section 10 discusses limits; Section 11 concludes. Proofs are in the Appendix; all are cross-checked numerically.

2. Related literature

Network externalities and standards. The lock-in of an inferior-for-some-purpose technology through increasing returns is the oldest thread here: Arthur (1989) and David (1985) on path dependence, Katz and Shapiro (1985) and Farrell and Saloner (1985) on installed base and compatibility. Our departure is the *altitude* dimension: lock-in is a property of the buyer's position on the stack, and the technology alone does not fix it; the network externality operates on one layer (the differentiated one) while the other clears competitively on price. This localizes the coordination pathology and produces the sorting and the scissors, which anonymous-share models do not.

Two-sided markets. The chicken-and-egg of platform adoption (Rochet and Tirole 2003; Caillaud and Jullien 2003; Armstrong 2006) is adjacent, but our engine is not cross-side pricing. It is a supply-side scope-and-network gradient over a space of needs; the "two sides," to the extent they exist, are layers of one buyer's own stack.

Multiproduct firms and bundling. The one-stop-shop advantage of a wide provider (Nalebuff 2004, and the bundling-as-entry-barrier tradition) is present in our π_D and in the breadth gap. Our contribution beyond it is the switching-cost distinction of §7 (a core-locking rewrite cost, informal here, and a formalized split-taxing coupling cost) and the un-bundling comparative static it yields.

Switching costs and lock-in. Farrell and Klemperer (2007) survey the field our κ and τ inhabit. We separate the switching cost into two components with *opposite policy targets*: proprietary-API rewrite cost, which locks the differentiated core and is addressed only by API/data portability; and egress/coupling cost, which taxes a split stack and is addressed by egress caps. This "egress rehabilitation" (§7) reconciles two facts that look contradictory: egress pricing is uncorrelated with market share, yet egress matters.

Coordination and equilibrium selection. The tipping structure of §5 uses the standard machinery (Kandori, Mailath and Rob 1993; Young 1993; Carlsson and van Damme 1993 for selection); our contribution is the altitude structure that places the trap on the differentiated layer, together with the up-stack-aging dynamic that makes the trend a demographic fact rather than an assumption; the method itself is standard.

Digital sovereignty and the companion programme. The policy motivation draws on weaponized interdependence (Farrell and Newman 2019) and on the companion trilogy: *Screening for Sovereignty* (Fermigier 2026a) on why structural requirements screen and behavioural ones wash; *Exit or Own* (2026b) on portability against ownership screens; *Two Screenings* (2026c) on the capture of cloud versus 5G screens. That programme takes the sovereign supplier pool as given and asks how to select from it; the present paper asks whether the pool forms and survives.

3. The dependence, in the record

The model is a theory exercise, but its choices are disciplined by public facts. We state them as sourced stylized facts with confidence flags (firm / rough); the full evidence base, its two-pass fact-check, and the refuted-claims log are in the data appendix.

Fact 1 (firm): Europe is cheaper where it competes. Geometric means over the getdeploying reference tables (2026; data and script in `pricing/`): EU low-cost leaders (Hetzner, OVH, Scaleway) undercut the three hyperscalers by 6× on like-for-like compute (the European field as a whole 2.5×), 2.4× on object storage, and 2.4× on H100 GPUs, while egress runs 10–86× cheaper (several EU providers charge nothing). US GPU specialists (Runpod, Oblivus) undercut everyone, so the GPU edge holds only against the hyperscalers. European commodity providers also win independent price-performance benchmarks, yet lose market share to the hyperscalers.

Fact 2 (firm): the advantage narrows and reverses up the stack. On managed PostgreSQL at production scale the gap collapses to parity (AWS RDS ≈1.5% cheaper than OVHcloud in a single, likely OVH-commissioned BenchANT 2025 benchmark whose own throughput figures leave the price-performance edge with OVHcloud); at the AI-inference layer US managed token-metered services are priced comparably or below European offerings (AWS Bedrock vs Scaleway model-as-a-service, vendor pricing pages). GPU rental favours Europe (2.4× on H100), but that reflects a hyperscaler on-demand premium: several non-hyperscaler specialists undercut the European pair, which sits at the expensive end of the GPU-specialist market.

Fact 3 (firm): Europe is a small minority in infrastructure, near-absent in AI. European providers hold 15% of the combined IaaS+PaaS market and no single one exceeds 2% (Synergy); the clean IaaS-only versus PaaS-only split is not published (it sits in paid IDC segment reports), though the European field is concentrated in IaaS/hosting and thin in PaaS. In generative AI, European cloud companies hold under 5% of infrastructure versus 85% for US firms, and captured essentially none of the

€3.1bn 2024 European gen-AI market, which went to OpenAI (30%), Microsoft (25%), AWS (20%), and Google (15%) (European Parliament study PE 778.576, 2025; this vendor split is the study's own inferred estimate, not measured revenue). This per-layer opacity is the paper's load-bearing data gap (§10).

Fact 4 (firm): value is migrating up-stack, and the aggregate is drifting. European providers grew revenue 167% over 2017–2022 yet lost half their share, because the total market grew 6× to €61bn and the new value went to the differentiated layers (Synergy 2025). The PaaS share of European IaaS+PaaS is forecast to rise from 44% to 54% (2020–2028; rough, single BDO/Statista source). Crucially, this decline (2017–2022) *predates* the generative-AI boom (ChatGPT, late 2022), so AI cannot be its cause: the differentiated layer that drove the migration was managed PaaS (databases, analytics, serverless), which the hyperscalers had dominated for years. Generative AI is the newest and most extreme member of that layer, and until about 2025 its footprint sat mostly in SaaS (chat assistants), off this paper's IaaS/PaaS axis; in the model it is a recent accelerant of a long-running drift, not the historical cause.

Fact 5 (firm): the formation and exit channels are American. AWS Activate offers start-ups up to $200,000 in credits, gated on an accepted (largely US-VC/accelerator) provider; hyperscaler marketplaces function primarily as committed-spend drawdown vehicles (enterprises retiring their AWS commitment against third-party software) and less as neutral sales channels; US venture capital supplies 28% of European start-up funding and US corporations 48% of European tech-M&A transaction value (2024, up from 9% in 2015; a value share, sensitive to individual megadeals) (European Parliament 2025; AWS pricing pages; CMA 2024 on committed-spend). The 2026 migration of DeepL (a European sovereignty-branded AI firm) onto AWS as a default sub-processor is the illustrative case: it followed its enterprise customers, who were already on AWS.

What the record cannot show. No identification is claimed. The distinguishing predictions are stated in §10; the record is also consistent with a simple "US got there first" story, which the dynamics of §8 nest rather than refute.

4. The model

Primitives. A unit mass of firms. Each firm has a type, its **altitude** $a \in [0, 1]$, the fraction of its cloud needs that are differentiated (platform/AI) rather than commodity (raw infrastructure); altitudes are distributed with CDF G and density g. Two providers, E (Europe) and U (United States). In the baseline each firm single-homes (picks one primary provider) because splitting a stack across providers carries a coupling cost (relaxed in §7; single-home is that model's $\tau \to \infty$ limit). Let m denote the aggregate share of firms on U.

Payoffs. A firm's net payoff from U over E is

$$\Delta(a, m) = a\,[\pi_D + \gamma(2m - 1)] - (1 - a)\,\pi_C = a\,[\pi_C + \pi_D + \gamma(2m - 1)] - \pi_C,$$

the sum of a commodity term (fraction $1 - a$ of needs, on which E is cheaper by $\pi_C > 0$) and a differentiated term (fraction a, on which U is favoured by $\pi_D > 0$ plus a network/ecosystem benefit $\gamma(2m - 1)$ increasing in U's share). Each parameter maps to a fact of §3: π_C to the 3–7× commodity price gap (Fact 1); π_D to US presence-plus-competitive-price at the top (Facts 2–3); γ to skills, integrations, marketplaces, and committed-spend that reward joining the larger ecosystem (Fact 5).

Assumptions. **(A1)** $\pi_C, \pi_D > 0$, $\gamma \geq 0$. **(A2)** $\gamma < \pi_C + \pi_D$, the "no-degenerate-preferences" condition that keeps the altitude gradient positive. **(A3)** G has connected support on $[0, 1]$ the empirically motivated leading case is *bimodal* (a cluster of commodity-only firms and a cluster of AI-native firms, with a thin middle), the benchmark case a broad unimodal G.

Remarks on the modelling choices. (i) Altitude is the right one-dimensional type because it is exactly the margin on which the two providers' cost advantages diverge: the scissors are its economic content. (ii) Single-home is the $\tau \to \infty$ limit of the general coupling model of §7, recovered when splitting is prohibitively costly, so it imposes no restriction. (iii) The network benefit is placed on the differentiated layer only, because that is where skills, proprietary APIs, and committed-spend actually accumulate; the commodity layer clears on price. This single placement produces the localization of the tipping trap (§5), and Propositions 1 and 3 should be read as consequences of it: sorting follows from the single-crossing payoff, and the trap's confinement to the differentiated layer follows from where γ is placed. (iv) The parameterization is illustrative, not estimated: values are chosen so the model's needs-share $m^* \approx 0.75$ sits near the 70% hyperscaler revenue share of §3, a related but distinct object, and the specific numbers carry no identification.

5. Equilibrium: sorting, existence, and the tipping trap

The coefficient on a in Δ is $B(m) \equiv \pi_C + \pi_D + \gamma(2m - 1)$, positive under (A2), so preferences are single-crossing in altitude.

Proposition 1 (sorting). *Under (A2), for each m there is a unique threshold altitude $a^*(m) = clip(\pi_C / B(m), 0, 1)$, and a firm chooses U iff $a \geq a^*(m)$.* Firms sort: the commodity-heavy take Europe, the differentiated-heavy take the United States (Figure 3). The aggregate consistency condition is $m = \Phi(m) \equiv 1 - G(a^*(m))$.

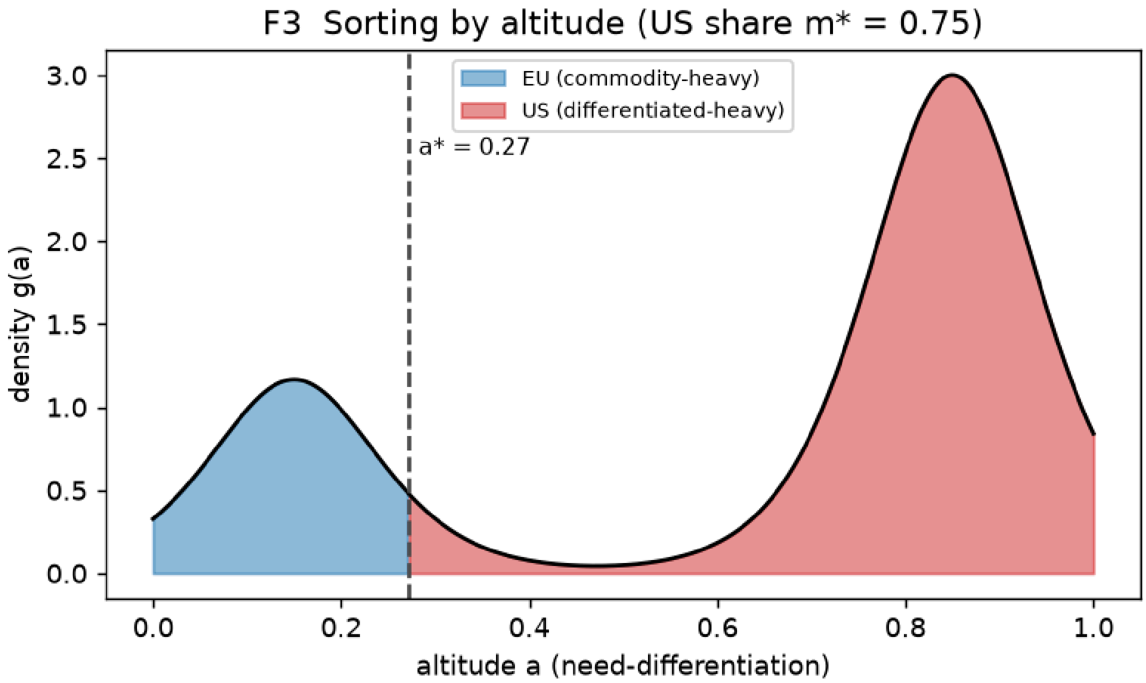

Figure 3. Sorting by altitude at the calibrated equilibrium: Europe wins the commodity cluster below $a^* = 0.27$, the United States the differentiated cluster above it; US share $m^* = 0.75$.

Proposition 2 (existence and lattice). *Φ is continuous and non-decreasing on $[0, 1]$ the set of equilibria is a nonempty complete lattice, with a least and a greatest equilibrium.* (Brouwer for existence; Tarski for the lattice.)

Proposition 3 (the tipping trap).

$$\Phi'(m) = g(a^*)\,\frac{2\gamma\pi_C}{B(m)^2} \geq 0.$$

If $\sup_m \Phi' < 1$ *the equilibrium is unique and globally stable; multiplicity of stable equilibria requires* $\Phi' > 1$ *somewhere (necessary, not sufficient). The transition is governed by* γ*: the coordination trap is created by the differentiated layer's network strength.* Because $\Phi' \propto \gamma$, raising the differentiated-layer network strength moves the market across the $\sup \Phi' = 1$ boundary into multiplicity, a corner-selecting tipping regime (Figures 2 and 5). The trap lives entirely on the differentiated layer; the commodity layer never tips. At the calibration used below ($\gamma = 0.5$), $\sup_m \Phi' > 1$ yet the equilibrium is unique (the supercritical-but-unique case): the tipping machinery exists in the model but does not bind at the operating point, where the share and its rise come from sorting and aging drift rather than from multiplicity.

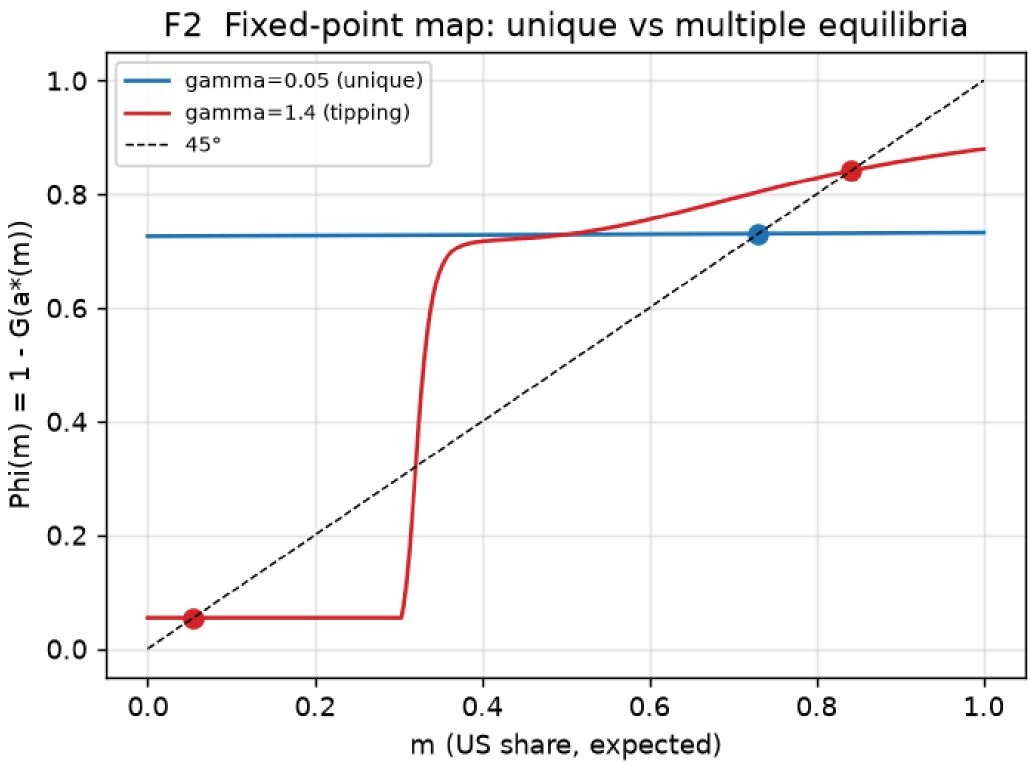

Figure 2. The fixed-point map $\Phi(m)$*: at low* γ *a single stable crossing; at high* γ *the curve steepens through the diagonal and admits multiple stable equilibria.*

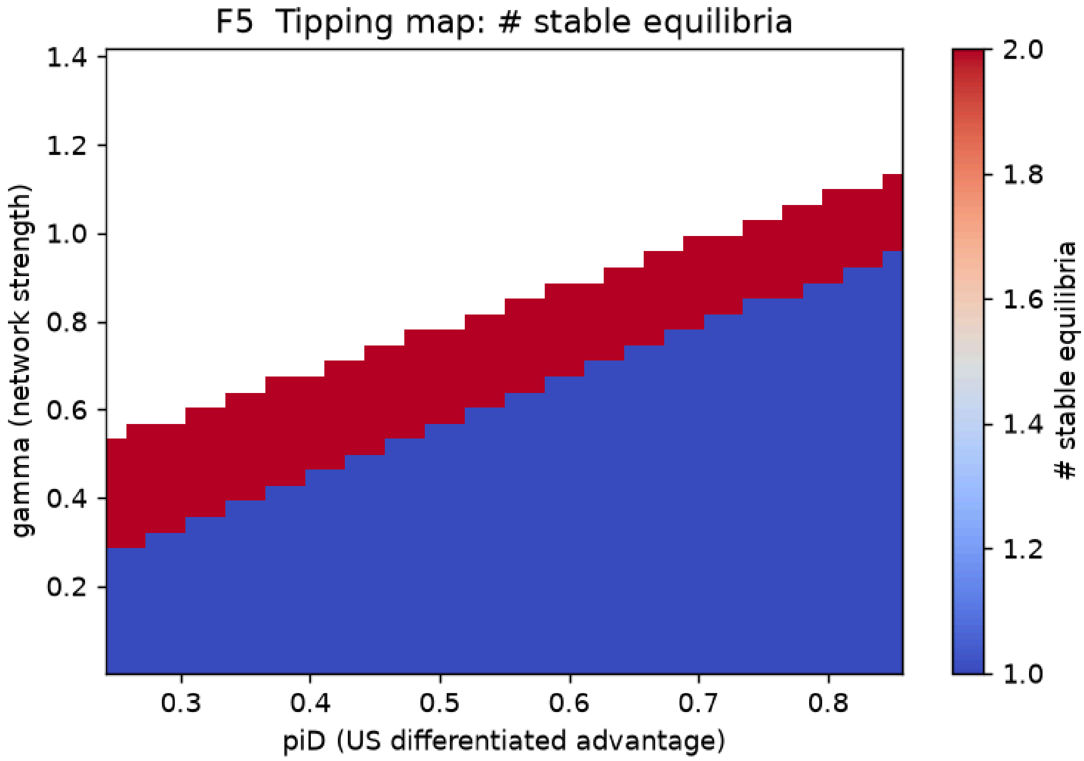

Figure 5. Number of stable equilibria over (π_D, γ): the multiplicity (tipping) region is entered by raising the network strength γ (raising π_D instead exits *it), confirming Proposition 3.*

6. Leverage: why the obvious lever fails

We now price the interventions. Write m^* for a stable equilibrium share and read the levers as comparative statics.

Proposition 4 (the commodity subsidy saturates). *At a stable equilibrium,*

$$\frac{dm^*}{d\pi_C} = \frac{-g(a^*)\,(B - \pi_C)/B^2}{1 - \Phi'} < 0,$$

with magnitude proportional to $g(a^*)$. As a sustained subsidy raises π_C, the threshold a^* rises toward the sparse middle of G, where $g \to 0$, so the net effect of a moderate *subsidy is small: it saturates. This is local: a* large *subsidy drives a^* onto the high cluster, where $g(a^*)$ spikes, so π_C can change the number of equilibria and flip the market, at a budget scaling with the whole commodity base. Winning the differentiated layer ($\pi_D < 0$) or lowering the network/coupling parameters restructures it more cheaply.*

The *local* slope of the commodity subsidy is not small; in our sensitivity analysis it is the largest of the three parameters (§9). Leverage is not local slope. A subsidy is a *sustained, sizable* intervention, and its net effect is governed by the density of marginal firms it can reach; as it pushes the threshold off the commodity cluster into the trough, that density collapses and the aggregate stops moving (verified: the derivative falls from −0.36 to −0.02 along the subsidy path; the intervention is concave on 99% of random configurations). More fundamentally, a *moderate* subsidy slides the equilibrium within the current basin; a large one moves the threshold onto the high cluster and can flip the market, at a budget scaling with the whole commodity base (§9). Interoperability reshapes the market directly and more cheaply. This is close to the distinction Meadows (1999) draws between intervening on a number (low leverage) and on the gain of a feedback loop (high leverage); the differentiated-layer network strength γ (and the coupling cost τ of §7) are the feedback levers, and they are where interoperability acts.

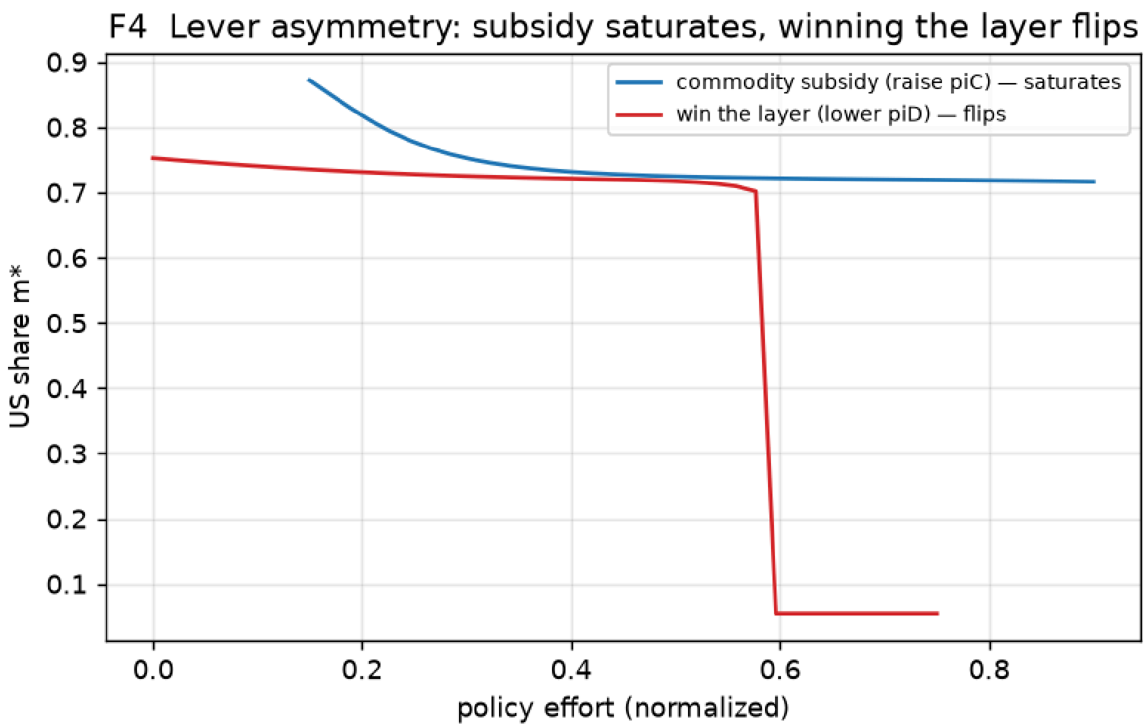

Figure 4. Lever asymmetry: raising the commodity subsidy (π_C) moves $m*$ little and saturates; making Europe genuinely competitive at the differentiated layer (π_D falling below zero) flips the market.

Proposition 5 (un-bundling). *Under single-homing the US needs-share is $m^* = 1 - G(a^*)$ under costless multi-homing (each workload to its best provider) it is $E[a]$. Europe strictly gains from multi-homing exactly when $m^* > E[a]$, by $m^* - E[a]$.* Interoperability's payoff is therefore to let a firm stop bundling, returning to Europe the commodity and data spend of high-altitude firms; a differentiated-heavy firm keeps its American core (in the calibration, the US needs-share falls from 0.75 to 0.65, and Europe rises from 0.25 to 0.35 without winning a single differentiated workload). The lever is real; its target is the bundle. The gain requires $m^* > E[a]$, which holds in the calibration but not universally (it fails on about one market configuration in seven).

The policy reading is a two-line assignment, and we defer the full leverage-points treatment to a companion policy paper. *Subsidy is the weak lever: the basement saturates, and at the differentiated layer, where Europe's money actually goes, network effects defeat a subsidised challenger the same way. Use interoperability to lower γ, un-bundle the stack and preserve a European foothold, and pursue genuine competitiveness where the value is, which a subsidy alone cannot buy.*

7. Coupling and the sovereign beachhead

Proposition 5 assumed frictionless multi-homing. That is unrealistic: moving data across provider boundaries costs egress and latency, so firms do not scatter compute, storage, and AI across clouds. The efficient split follows the **coupling boundary**: keep the tightly-coupled core co-located, and offload only *loosely-coupled* workloads (backups, disaster recovery, dev/test, batch, archival) to a secondary provider. This is the "US-primary, EU-secondary" architecture already practised, and it is where the sovereign foothold actually lives.

Let a cross-provider **coupling cost** τ be incurred per unit of data a workload exchanges with the core, and let a peripheral workload have coupling intensity $c \in [0, 1]$.

Proposition 6 (the beachhead). *A US-primary firm offloads a peripheral workload to a European secondary iff $\pi_C > \tau c$, i.e. $c < \pi_C/\tau$ low-coupling workloads (backups) are captured first. Europe's aggregate needs-share $= (1 - m^*) + \min(1, \pi_C/\tau) \int_{a>a^*} (1 - a)\, dG$ is non-increasing in τ its $\tau \to 0$ limit $E[1 - a] + \int_{a<a^*} a\, dG$* exceeds *the frictionless bound $E[1 - a]$, because the one-directional secondary keeps low-altitude firms' differentiated mass in Europe.* Lowering τ grows the beachhead: as coupling cost falls, Europe reclaims first the archival and backup workloads, then the chattier secondaries (Figure 6).

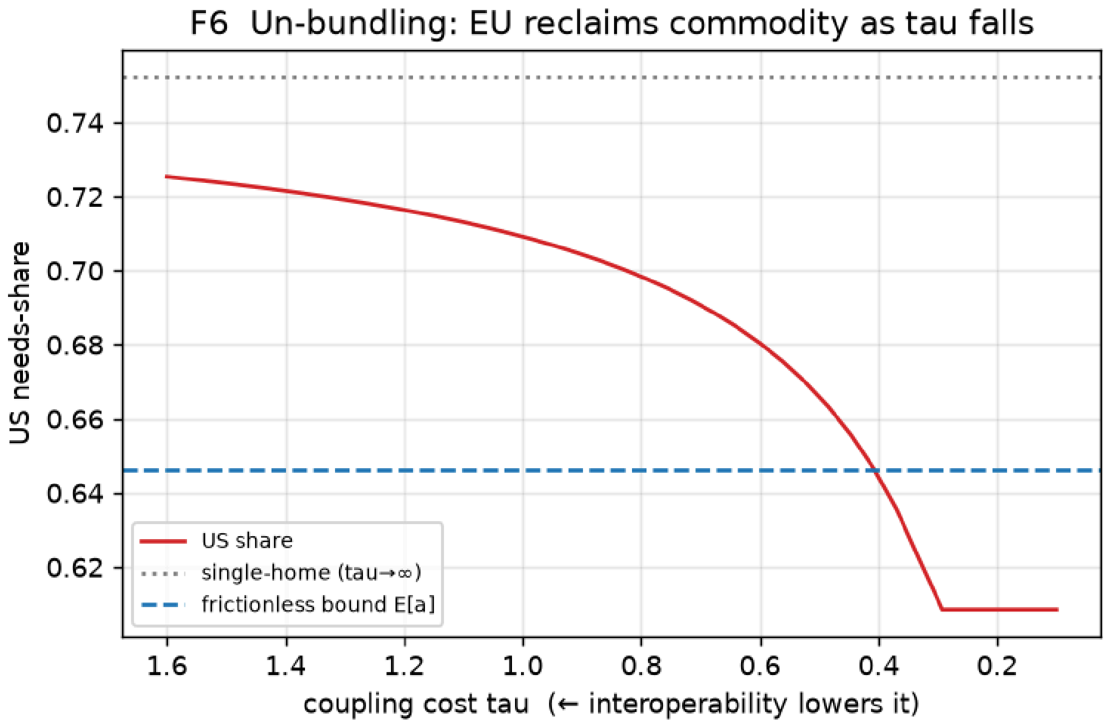

Figure 6. Un-bundling and the beachhead: as the coupling cost τ falls (interoperability), Europe's needs-share rises from the single-home level to a $\tau \to 0$ limit slightly above the frictionless bound $E[1 - a]$.

Egress, rehabilitated. This dissolves an apparent contradiction. Egress pricing is *uncorrelated* with market share: OVHcloud offers free egress yet holds only a few percent of the market, while AWS charges dearly for egress and holds 30% (egress prices from getdeploying; market shares from Synergy). So egress is not the lock that keeps firms in. Egress *is* the tax on running a European secondary. Two frictions with different targets: a proprietary-API rewrite cost κ (informal here, not part of the formal model) locks the *differentiated core* and is addressed only by API and data portability; the formalized coupling cost τ taxes the *split* and is addressed by egress caps such as the Data Act's. The Data Act's egress phase-out does not free the core; it makes the primary/secondary architecture affordable. This is the precise, egress-independent case for the portability lever, and it sharpens the companion *Exit or Own*: portability's value here is un-bundling, not exit.

The value of a holding action. The beachhead is bounded by the commodity fraction of high-altitude firms (the commodity fraction of AI-native firms is 15% in the calibration), so it recovers single-digit points of share, not the market. Its worth exceeds those points. It provides **sovereign reversibility**: a live European copy is the exercisable exit option of *Exit or Own* and the operational form of

DORA's Article 28 exit strategies. And it provides **capability preservation**: it keeps a European supplier alive and funded, the supplier that *Screening for Sovereignty* presupposes must exist before any procurement screen can select it. This is the bridge between this paper's supply side and the trilogy's demand side: the beachhead is how the sovereign supplier survives long enough to be worth screening for, and, if Europe later climbs the stack (π_D falling), the secondary becomes a primary.

8. Dynamics: why the share rises

The observed 75% is a point on a trajectory rather than an equilibrium to be explained. Firms are born low-altitude and drift up-stack over their lifecycle: a young firm needs a virtual machine; a maturing one adds a managed database and analytics; a scaling one, recently, adds AI.

Proposition 7 (the rising trend). *If every firm's altitude drifts up ($\dot{a}_i > 0$) and none repatriates (high κ), then $m(t)$ is non-decreasing. At subcritical γ the tracking equilibrium is unique; a rise in the differentiated-layer advantage (π_D rising, lowering a^*) raises m further; with finite κ there is a repatriation hysteresis band.* As the population ages the US share climbs from a young-market minority toward the corner, passing through today's level on the way (Figure 7); managed-PaaS adoption (databases, analytics, serverless) drove this climb since well before 2022, and the recent AI surge is only its latest and steepest push; and hysteresis means the ascent does not reverse when conditions ease. The "US got there first" story is nested here as the special case of a market caught early in its aging; the model adds that the drift is structural and, absent intervention, monotone. Figure 7 shows the aging channel at a subcritical γ and a spread of initial altitudes, to isolate the demographic mechanism; at the calibrated $\gamma = 0.5$ the same drift tips the market sharply rather than gently, and the up-stack-drift premise itself is asserted, not measured.

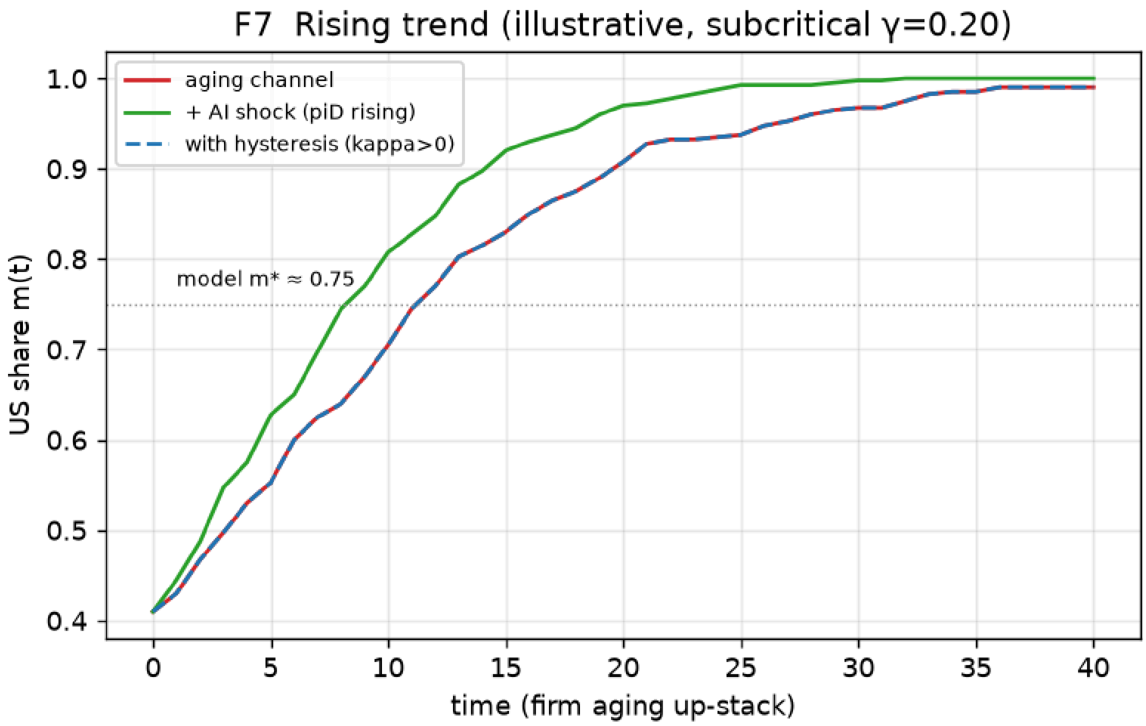

Figure 7. The rising trend (illustrative, at a subcritical $\gamma = 0.20$ to isolate the aging channel): firms aging up-stack (into managed PaaS, and recently AI) drive $m(t)$ from a young-market minority through the model's $m^* \approx 0.75$ toward the corner; a rise in the differentiated-layer advantage

accelerates it; hysteresis prevents reversal. At the calibrated $\gamma = 0.5$ the same drift tips sharply rather than gently.

9. Verification

Because the paper's claims are used to argue against a live policy instrument, we hold them to a three-gate standard and document it. A claim ships only if it clears all three: an analytical **proof** (Appendix A); a **numeric cross-check** of the closed form against brute-force computation; and a **Monte-Carlo** confirmation that the qualitative claim survives random draws of parameters and of the altitude distribution G. One command (`verify/run_all.py`) reproduces every number and figure in about a minute.

All seven propositions pass their cross-checks. The Monte-Carlo results (Figure 8): the finite-population share converges to the continuum equilibrium at the $1/\sqrt{N}$ rate (log-log slope −0.53); sorting is robust on 95% of random distributions and un-bundling helps Europe on 86%; the phase-transition characterization holds exactly in its necessary direction (subcritical $\Longrightarrow$ unique on 100%, multiplicity $\Longrightarrow$ supercritical on 100%); the rising trend is monotone on 100% of dynamic ensembles and the differentiated-layer shock (AI most recently) accelerates it on 100%; and the leverage typology is confirmed: the commodity subsidy is concave (saturating) on 99% of configurations, while interoperability ($\gamma \to 0$) removes the tipping trap on 100% of multi-equilibrium configurations. The one qualitative claim that is *not* universal is reported as such: commodity-subsidy saturation holds on 72% of random configurations, failing where a rising subsidy happens to push the threshold into a denser part of G, the boundary of Proposition 4 acknowledged above, and the reason its rigorous statement is about structure and saturation rather than local slope.

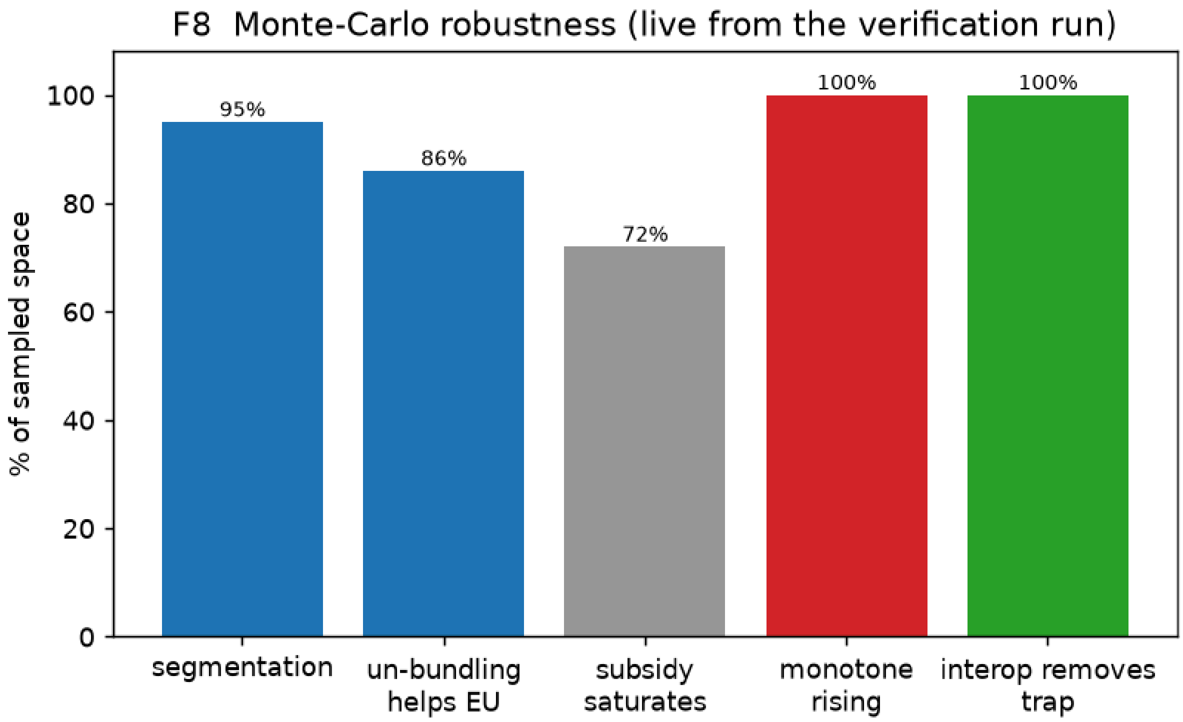

Figure 8. Monte-Carlo robustness of the qualitative claims across the sampled parameter and distribution space (numbers read live from the verification run; the 72% subsidy-saturation is the weakest, as flagged in §9).

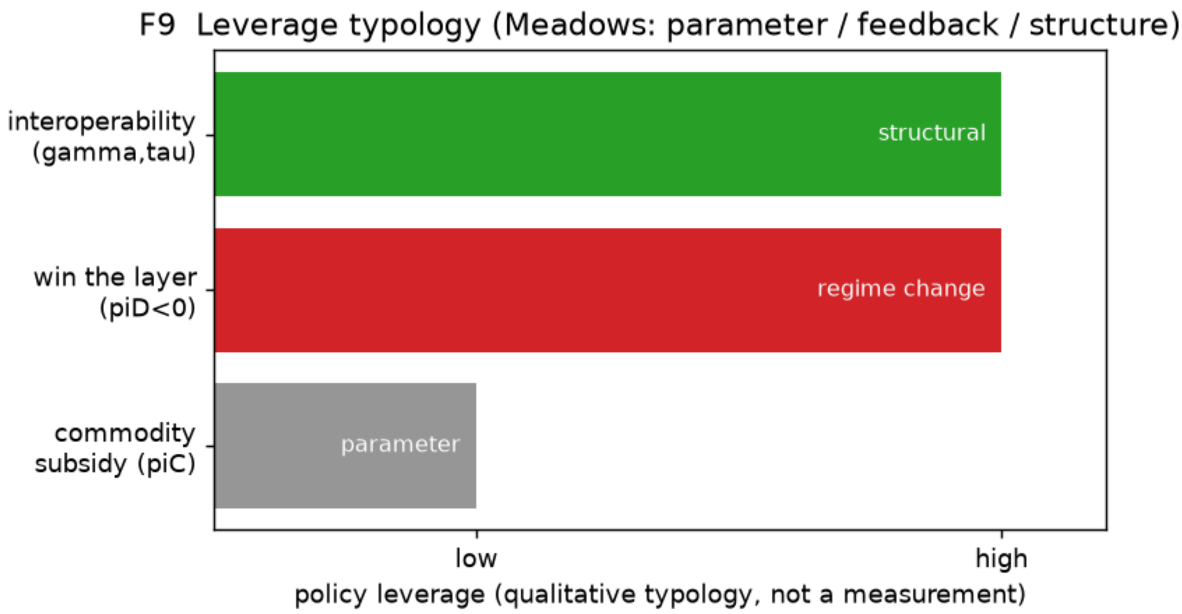

Figure 9. The leverage typology grounding the policy reading (qualitative, not a measurement): the commodity subsidy is a parameter (saturates over realistic budgets); winning the layer is a regime change; interoperability is structural (removes the trap).

10. Discussion and limits

The distribution is load-bearing. The strength of Proposition 4 depends on the bimodality of G: a thin middle starves the saturating subsidy of marginal firms. The clean numeric split of European share within each layer, and the shape of G itself, are the data this paper most wants and does not have; a modest survey of firms' service-mix would supply the second, and is the highest-value next datum. We prove the results for general G and flag where bimodality does work.

Modelling choices. Single- versus multi-homing is handled by making the coupling model of §7 primitive and single-home its limit. Competition among US providers, endogenous π_D (Europe could climb), and the full auction are deferred; the one-provider-per-side case is the conservative one for our claims. Pricing figures are time-sensitive (reserved discounts, mid-2026 Hetzner increases); we use signs and orders of magnitude, and date every figure.

What would kill the paper. Three observations would refute the mechanism. Evidence that provider choice is *not* sorted by altitude (commodity-heavy firms on the American providers at the same rate as AI-heavy firms) would break Proposition 1. A durable *interior* plateau of the US share, instead of a continued climb toward the corner, would break Proposition 7. And an egress-fee ban *alone* reversing the trend would break the κ/τ distinction of §7, which predicts API/data-portability standardisation as the binding lever. Each is a clean, near-term test.

11. Conclusion

Europe's cloud problem is not that its technology is worse; in the commodity basement it is cheaper. The problem is that the value has moved to a layer where the American providers are price-competitive and European ones are barely present, and that firms climb into that layer as they grow. Subsidy cannot fix this: the basement saturates, and at the differentiated layer, where Europe's subsidies actually go, the barrier is the network and no budget dissolves it. The levers that move are interoperability, which un-bundles the stack and makes a European secondary affordable, and the hard one: becoming genuinely competitive where the value is. The second is a diagnosis rather than a lever (the model takes Europe's up-stack weakness as a primitive), so the actionable content is the interoperability and beachhead argument. Until then, the achievable European position is a beachhead (backups, reversibility, a supplier kept alive), and that beachhead is exactly what the demand-side screen of the companion programme needs to have something sovereign to select. Cheaper is not enough; altitude is the battle.

Appendix A. Proofs

Throughout, $B(m) = \pi_C + \pi_D + \gamma(2m-1)$, $a^*(m) = clip(\pi_C/B(m), 0, 1)$, $\Phi(m) = 1 - G(a^*(m))$. Each proof is verified against independent brute-force computation by the named check in `verify/`; the full statements and the numerical logs are in `notes/proofs.md`.

Proposition 1. $\partial\Delta/\partial a = B(m) \geq \pi_C + \pi_D - \gamma > 0$ by (A2), so $\Delta(\cdot, m)$ is strictly increasing in a. $\Delta(0, m) = -\pi_C < 0$ where $B(m) > \pi_C$, $\Delta(1, m) = B(m) - \pi_C > 0$, giving the unique interior zero $a^* = \pi_C/B(m)$. If $B(m) \leq \pi_C$ every firm prefers E and $a^* = 1$. (Verified: `check_P1`.) ∎

Proposition 2. B is affine and positive (A2), so $a^* = \pi_C/B$ is continuous and non-increasing in m G continuous non-decreasing (A3) makes $\Phi = 1 - G(a^*)$ continuous non-decreasing on $[0, 1] \to [0, 1]$. Brouwer gives a fixed point; Tarski, applied to the monotone Φ on the complete lattice $[0, 1]$, gives that the fixed-point set is a nonempty complete lattice with least and greatest elements. (Verified: `check_P2`.) ∎

Proposition 3. $a^{*\prime} = -\pi_C \cdot 2\gamma/B^2$, so $\Phi' = -g(a^*)a^{*\prime} = g(a^*)2\gamma\pi_C/B^2 \geq 0$. If $\sup \Phi' < 1$, Φ is a sup-norm contraction: unique fixed point, globally stable. Two stable fixed points (each with $\Phi' < 1$) force, by the intermediate-value theorem on $\Phi(m) - m$, an interior point with $\Phi' \geq 1$ hence multiplicity $\Rightarrow \exists\Phi' > 1$. The converse fails where the steep region is off the diagonal (verified: 24% of supercritical grid points remain unique). Since $\Phi' \propto \gamma$, raising γ crosses the boundary. (Verified: `check_P3`, `mc3`.) ∎

Proposition 4. Implicit differentiation of $\Phi(m; \pi_C) - m = 0$ gives $dm^*/d\pi_C = -(\partial\Phi/\partial\pi_C)/(1 - \Phi')$ with $1 - \Phi' > 0$ at a stable point. $\partial a^*/\partial\pi_C = (B - \pi_C)/B^2 > 0$, so $\partial\Phi/\partial\pi_C = -g(a^*)(B - \pi_C)/B^2 < 0$, hence the sign and the $g(a^*)$ proportionality. Raising π_C raises a^* into the trough where $g \to 0$, collapsing the derivative (verified: −0.36⟶−0.02; 99% concave). n_stable is governed by $\sup \Phi' \propto \gamma$ (P3), which π_C cannot drive below 1 in general, whereas $\gamma \to 0 \Rightarrow \Phi' \equiv 0 \Rightarrow$ unique (verified: removes multiplicity on 100% of multi-equilibrium configs). (Verified: `check_P4`, `mc5`.) ∎

Proposition 5. Single-home: a firm above a^* puts its whole unit of needs on U, so the U mass is $1 - G(a^*) = m^*$. Multi-home: each firm sends its differentiated fraction a to U and its commodity fraction $1 - a$ to E, so the U mass is $\int a\, dG = E[a]$. Europe's share rises from $1 - m^*$ to $1 - E[a]$, a gain of $m^* - E[a] > 0$ iff $m^* > E[a]$. (Verified: `check_P5`.) ∎

Proposition 6. Offloading a peripheral workload trades the U price for the E price (gain π_C) against coupling penalty τc profitable iff $\pi_C > \tau c \Leftrightarrow c < \pi_C/\tau$, a set increasing as $\tau \downarrow$ and ordered by c. A US-primary firm of core altitude $\bar{a}$ yields Europe $(1 - \bar{a}) \min(1, \pi_C/\tau)$ aggregating over the m^* US-primary and $1 - m^*$ EU-primary firms gives the stated share, non-increasing in τ, bounded by $E[1 - a]$ as $\tau \to 0$. (Verified: `check_P6`.) ∎

Proposition 7. At subcritical γ the equilibrium map is a contraction (P3), so m is a continuous function of the altitude profile's a^*-crossing mass. Each a_i non-decreasing in t makes $\{i : a_i(t) > x\}$ non-decreasing for every x evaluated at the (weakly falling) threshold, $\{a_i > a^*\}$ is non-decreasing, so $m(t)$ is non-decreasing. A rise in the differentiated-layer advantage (an AI surge, most recently) lowers a^* (via B rising in π_D), enlarging the set. Finite κ creates a $[a^* - \kappa/B, a^*]$ non-repatriating band (hysteresis) and preserves monotonicity a fortiori. (Verified: `check_P7`, `mc4`.) ∎

References

Citation details for the industrial-organization literature are at working-paper level and will pass the full verification protocol before submission; empirical sources are verified against primary records with confidence flags in the data appendix.

- Asterès (2025). *Technological dependence on American software and cloud services: an assessment of the economic consequences in Europe.* Study commissioned by Cigref, May 2025.
- Asterès (2026). *From technological dependency to economic capture: the cost of cloud and software services inflation to Europe.* Study commissioned by Cigref, May 2026.
- Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics* 37(3):668–691.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal* 99(394):116–131.
- Caillaud, B. & Jullien, B. (2003). Chicken & egg: competition among intermediation service providers. *RAND Journal of Economics* 34(2):309–328.
- Carlsson, H. & van Damme, E. (1993). Global games and equilibrium selection. *Econometrica* 61(5):989–1018.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review* 75(2):332–337.
- Farrell, H. & Newman, A. L. (2019). Weaponized interdependence. *International Security* 44(1):42–79.
- Farrell, J. & Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. In *Handbook of Industrial Organization*, vol. 3, Elsevier.
- Farrell, J. & Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics* 16(1):70–83.
- Fermigier, S. (2026a). Screening for sovereignty: optimal procurement when compliance is cheap to fake and costly to verify. Working paper (companion).
- Fermigier, S. (2026b). Exit or own? Portability mandates and ownership screens in sovereign procurement. Working paper (companion).
- Fermigier, S. (2026c). Two screenings: why Europe screened Huawei but not the hyperscalers. Working paper (companion).
- Kandori, M., Mailath, G. J. & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica* 61(1):29–56.
- Katz, M. L. & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3):424–440.
- Meadows, D. H. (1999). *Leverage Points: Places to Intervene in a System.* The Sustainability Institute.
- Nalebuff, B. (2004). Bundling as an entry barrier. *Quarterly Journal of Economics* 119(1):159–187.
- Rochet, J.-C. & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association* 1(4):990–1029.
- Young, H. P. (1993). The evolution of conventions. *Econometrica* 61(1):57–84.

Empirical and market sources (dated, with confidence flags, in the data appendix): Synergy Research Group (2022, 2025) on European cloud market shares; European Parliament study PE 778.576 (2025) on cloud and generative-AI market structure; IDC MarketScape European public cloud IaaS (2024); UK CMA cloud services market investigation (2024) on egress and committed-spend; provider price lists and independent benchmarks (Hetzner, OVHcloud, Scaleway, AWS, Azure, GCP;

vpsbenchmarks, dev.to, BenchANT, getdeploying); AWS Activate and Bedrock pricing pages; DeepL infrastructure announcement (2026).

Data and code. The evidence base (`notes/figures.md`), the model and proofs (`notes/layered-model.md`, `notes/proofs.md`), and the verification suite (`verify/`, one-command reproduction) are available from the author.