Cuts and Options: A Theory of Strategic Autonomy in Production Networks

Stéfane Fermigier · Working paper v0.8 · 14 July 2026

Abstract. Modern production runs through deep vertical chains in which a producer needs all of its input classes and can, within each class, choose among substitutable suppliers. This paper develops the combinatorial theory of coercion and diversification on such structures. Its scope is the channels through which a controller can deny or condition another actor's operation (physical supply, technology and licensing, legal reach, distribution); the slower channels acting on the same structure, influence over rule-setters, capital, and the capture of suppliers, reshape the structure itself over time and are treated as dynamics, in the research programme of Section 6 and in companion work. Two dual quantities organize the statics: the coercion number κ, the size of the smallest coalition of controllers able to cut every feasible realization of a buyer's production, and the option number N, the largest family of realizations sharing no controller. Coercion composes as a weakest link across independent input classes; $N \leq \kappa$ always; and the central result is a possibility theorem for certification failure: system-wide diversification can be strictly below every per-layer option count, with the gap opening exactly where the same controller appears at more than one layer. A per-layer independent-suppliers criterion, the form regulation actually takes, therefore certifies a quantity that can equal two at every layer while a single controller cuts everything. A forkable substrate acts as a standing, self-controlled option and caps exit costs. The graded Leontief operator of applied work coincides with the calculus in support and interpolates it monotonically in magnitude. The primitives are classical, and the mathematics is kin to minimal cut sets in reliability theory and to network interdiction; the control overlay, the duality pair, and the certification theorem aimed at regulatory texts have, to our knowledge, no antecedent in the production-network economics literature. Applied to a measured 49-category graph of the European digital stack, the theory computes: every European demand node has $\kappa = 1$, with five distinct blocs each able to cut alone at full stack depth and the US bloc above the semiconductor floor for every buyer; on the legal channel both the US and China hold unilateral cut capability; and the public-sector buyer passes a per-layer two-suppliers test at every required class while its true option number is one. A supplier-resolution recomputation on pilot data splits both ways: a fully self-controlled productivity stack exists above the silicon floor, while the accelerator class offers seven suppliers under a single controller. The framework gives exact readings to the criteria in European cloud regulation and locates what they can and cannot certify.

JEL: F52, L14, D85, L86, O25. **Keywords:** production networks, strategic autonomy, coercion, diversification, options, supply security, certification, digital sovereignty.

---

1. Introduction

A firm, an administration, or a bloc that consumes a modern digital service stands at the top of a production structure with a specific logical shape: at every stage, all required input classes must be served (compute and software and distribution and silicon beneath them all), while within each class the buyer may choose among whatever substitutable suppliers exist. Strategic questions about such a buyer, who can cut it off, how many independent ways it has of operating, what a certificate of supplier plurality actually certifies, are questions about this AND/OR shape, and they have exact answers. This paper works them out.

The theory is organized by two quantities. The **coercion number** κ of a buyer is the size of the smallest coalition of outside controllers whose joint refusal makes every feasible way of operating infeasible. The **option number** N is the largest number of ways of operating that share no controller at all, the buyer's capacity for simultaneous diversification. These are dual objects: every one of N controller-disjoint realizations must donate a distinct member to any cutting coalition, so $N \leq \kappa$ always (Theorem 2). Coercion, in turn, composes as a weakest link: a coalition cuts the buyer if and only if it fully covers some single input class, so κ is the minimum of the per-class coercion numbers (Theorem 1). The floor result of applied work, that a fully foreign deepest layer defeats every ownership arrangement above it, is the degenerate case in which a necessary node hands $\kappa = 1$ to any of its controllers.

The central result is a negative one. Regulation certifies supplier plurality one layer at a time: the French SecNumCloud referential requires the capacity to maintain service from the provider's own competences or from at least two third companies; the European Commission's Cloud Sovereignty Framework defines its highest assurance level by the absence of critical non-EU dependencies, assessed offer by offer. Theorem 3 shows that the quantity such tests inspect, the per-layer option count, does not bound the system quantities: there are structures in which every layer has two independently controlled suppliers and yet no two controller-disjoint ways of operating exist, so that $N = 1$ while every certificate reads two. The gap opens exactly when the same controller appears in the option sets of different layers, and that configuration is ordinary: a bloc that supplies cloud platforms, operating systems, and accelerator silicon at once sits in three option sets simultaneously, which is the measured shape of the actual stack. The shortfall between the certified per-layer counts and the true option number, the **bloc-spanning discount**, is a well-defined, computable quantity, and Section 7 computes it on measured data: the European public-sector buyer passes a per-layer two-suppliers test at every required class while its option number is one and its coercion number is one.

Options also carry prices, and the theory says where they come from. A forkable substrate, an input whose specification and implementation can be self-hosted, is a standing option under the buyer's own control: it removes its layer from every cutting coalition's cheapest targets and caps the buyer's exit cost by the fork-and-operate cost (Theorem 5). This converts a regularity documented in applied work, that domestic alternatives exist almost exactly where substrates are open, into a theorem plus a testable premise. The graded machinery of production-network economics, Leontief propagation with substitutability weights, is recovered as the smooth interpolation of the cut calculus: its support coincides with the reachability calculus exactly, and its magnitudes move monotonically with substitutabilities, which is what licenses its use on cyclic structures where exact enumeration is unavailable (Section 5).

Two further components are developed as a programme with explicit status rather than as results. The interaction of policy levers is measured as partial substitution on the graded functional in every Monte-Carlo draw of the companion measurement work, and the naive theoretical account of that finding is provably wrong: counting fully liberated buyers is not submodular, since completing a buyer's last blocked layer has increasing returns. The refined conjecture, interior substitution with completion complementarity at the floor, is stated in Section 6.1 and left open. The dynamics, options as public goods whose insurance value is non-rival across buyers while their maintenance is paid by whoever exercises them, with the resulting under-provision of diversification by myopic buyers, descends from the second-sourcing literature and is developed jointly with a companion game-theoretic project (Section 6.2).

The contributions, stated against the state of the art, are five. First, the formal object: production as a controlled AND/OR structure, with control entering as channel-typed controller sets on options instead of as node attributes; the neighboring literatures carry either the production network without control identity [Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi 2012; Liu 2019] or control and coercion without the combinatorial production structure [Farrell and Newman 2019;

Clayton, Maggiori and Schreger 2026]. Second, the duality pair: the coercion number and the option number, connected by the series law and weak duality; Elliott, Golub and Leduc [2022] count sourcing multiplicity against random disruption, and we know of no prior formalization of coalition-based denial against controller-disjoint diversification. Third, the certification theorem: per-layer plurality criteria cannot bound the system quantities, with the gap characterized exactly (cross-layer controller sharing) and localized (at least three layers, joint control, or depth); to our knowledge it has no antecedent, and it is the central result because the criterion form it defeats is the form European regulation actually takes. Fourth, the standing-option ceiling, which gives open substrates an exact economic role: a self-controlled option that removes its layer from every cutting coalition and caps exit costs. Fifth, the support correspondence, which licenses the graded Leontief operator of applied work as the smooth version of the calculus. What is claimed for the mathematics is assembly and application: the primitives (AND/OR reachability, Menger-type duality, Neumann series) are classical and every proof is short; the value claimed is that the assembly makes previously rhetorical questions (how diversified is Europe, what does a certificate certify, who can coerce whom) exactly computable, and Section 7 computes them.

The applied counterpart of this paper is the measurement instrument of Fermigier [2026], which builds the 49-category dependency graph of the European digital stack and the three-layer option-set evidence this paper's Section 7 computes on. The relation is tight: that paper measures a structure; this paper is the theory of what such structures imply. The floor, gap, and standing-option results specialize to results reported there; the duality pair and the certification theorem are new here, and Section 7 computes them on that paper's data.

2. The structure

Definition 1 (production structure). A production structure for a buyer u is a finite rooted AND/OR tree: the root is u each node lists its required input classes (conjunctive); each class lists its feasible supplier options (disjunctive); each option either is primitive or carries its own required classes. Cyclic dependency graphs, which occur in measured data, are treated by the greatest-fixed-point semantics of Section 7 or by the graded operator of Section 5; the exact calculus of Sections 3 and 4 is stated on the acyclic structure.

Definition 2 (controllers). A finite set B of controllers (blocs, firms, jurisdictions) and a map c assigning each option node the set $c(v) \subseteq B$ of controllers able to deny or condition its operation on the channel under study. Channels (equity and governance, technology, physical supply, legal reach, standards, distribution, capital) induce different c maps on the same structure; multi-controller nodes are permitted; self-controlled nodes carry $c(v) = \emptyset$.

Scope of the calculus. The cut calculus formalizes the channels with denial-or-conditioning semantics: a controller on such a channel can stop or reprice an option today, which is what Definitions 4 and 5 price. The remaining ways one node influences another in the full measured graph, influence over rule-setters and standards, capital and funding relations, and the acquisition of suppliers, act on the structure itself: they change which options exist, what the defaults are, and who controls what, over time. Those channels are the subject of a planned second theory paper on the dynamics, seeded in Section 6.2 and in the capture companion, and the applied companion carries them as measured overlay layers awaiting that treatment. Two channels are dual citizens by design: equity and governance moves through transactions, which are dynamics, while its resting state, who owns what now, is the controller attribution the statics maps read (the "ownership channel" of Section 7 computes on exactly that attribution); standards likewise bind today's interfaces while their authorship is dynamics. The division is deliberate: the statics answer "who can cut whom, today, given the structure"; the dynamics answer "who is reshaping the structure".

Definition 3 (realization and trace). A realization R of u selects one option per required class, recursively; its trace is $c(R) = \bigcup_{v \in R} c(v)$. $\mathcal{R}(u)$ denotes the set of realizations. A realization is the object a buyer actually runs; the set $\mathcal{R}(u)$ is what its strategic position is made of.

Definition 4 (coercion number). A coalition $S \subseteq B$ cuts u if every $R \in \mathcal{R}(u)$ has $c(R) \cap S \neq \emptyset$ then $\kappa(u) = \min\{| S | : S \text{ cuts } u\}$, with $\kappa(u) = \infty$ if some realization has empty trace.

Definition 5 (option number). $N(u)$ is the maximum size of a family of realizations with pairwise disjoint traces.

Definition 6 (per-class breadth). $N_k(u)$ is the option number of the structure containing only class k at the root. A per-layer certificate inspects N_k (or, in its weakest form, the count of options at class k judged by their own controller sets, ignoring their subtrees).

The single-adversary special case recovers the applied notions: u is path-exposed to b iff $\{b\}$ cuts u, i.e. iff κ restricted to coalitions containing only b is 1, and the sovereignty gap of the applied work is the difference between this and the node-level ownership check.

3. Statics: the duality and the discount

Theorem 1 (series law). Let the root's required classes have node-disjoint option subtrees. A coalition cuts u iff it cuts some class entirely, and $\kappa(u) = \min_k \kappa_k(u)$.

Proof. If S cuts class k, every realization contains some option of k and is traced by S there. If S cuts no class, pick per class an option and a sub-realization avoiding S their combination, which node-disjointness makes feasible, is a realization avoiding S. The formula follows: cutting the cheapest class suffices, and nothing cheaper can work. □

A node contained in every realization (a necessary node) yields $\kappa(u) = 1$ for any of its controllers, whatever happens elsewhere in the structure: the deepest shared layer prices the whole system's coercion resistance. This is the silicon-floor proposition of the applied work in one line.

Theorem 2 (weak duality). $N(u) \leq \kappa(u)$.

Proof. A cutting coalition meets each of N pairwise disjoint traces, which requires N distinct members. □

Theorem 3 (the cross-layer discount). $N(u) \leq \min_k N_k(u)$ always. The inequality can be strict even when every option is singleton-controlled and every class has $N_k \geq 2$, and strictness requires a controller present in the option sets of at least two classes: if no controller appears in two classes, $N(u) = \min_k N_k(u)$.

Proof. The upper bound: a disjoint family at the root restricts to a disjoint family within each class. Strictness by construction: three classes with option controller pairs $\{X, Y\}$, $\{X, Z\}$, $\{Y, Z\}$ give $N_k = 2$ per class and $\kappa = 2$, yet every realization's trace is a 2-subset of $\{X, Y, Z\}$ and any two 2-subsets of a 3-set intersect, so $N = 1$. The equality case: with $m = \min_k N_k$ trace-disjoint options per class and no cross-class controller sharing, combining the i-th option of each class for $i = 1, \ldots, m$ produces m pairwise disjoint realizations. □

Corollary (limits of per-layer certification). A criterion that inspects one layer at a time certifies $\min_k N_k$. For every $m \geq 2$ there are structures with $N_k \geq m$ at every class and $N = 1$. Consequently a stack can satisfy a per-layer two-independent-suppliers requirement everywhere while a coalition of two, or in the presence of a necessary node one, controller cuts everything; and the quantity $\min_k N_k - N \geq 0$, the bloc-spanning discount, measures exactly how much the certificate overstates diversification. The discount is zero when no controller spans layers and is expected positive wherever one bloc supplies several layers at once.

The corollary is aimed at real texts. SecNumCloud §19.6.d requires maintainability “from the provider’s own competences or from at least two third companies”: a per-layer breadth requirement, $N_k \geq 2$ in the weak node-level form. The Cloud Sovereignty Framework’s top level requires “no critical non-EU dependencies” for one offer at a time. Both criteria are meaningful, and neither can bound N or κ: that requires the joint option structure across layers, which is a property no per-offer, per-layer audit observes. Section 7 exhibits the gap on measured data.

Theorem 4 (two-class equality). If the structure has exactly two required classes and every option is primitive and singleton-controlled, then $N = \kappa = \min(|C_1|, |C_2|)$, where C_k is the set of controllers appearing in class k.

Proof. By Theorem 1, $\kappa = \min_k \kappa_k$, and a flat class of singleton-controlled options is cut exactly by the coalitions containing all of its controllers, so $\kappa_k = |C_k|$. For N, let $m = \min(|C_1|, |C_2|)$ with, say, $|C_1| = m$, and let $S = C_1 \cap C_2$. Each shared controller $x \in S$ yields a realization choosing the x-controlled option in both classes, with singleton trace $\{x\}$ these are pairwise disjoint. Pair the $m - |S|$ remaining controllers of C_1 injectively with distinct controllers of $C_2 \setminus S$, which exist because $|C_2 \setminus S| \geq m - |S|$ the resulting two-element traces are pairwise disjoint and avoid S. This exhibits m disjoint realizations, and weak duality caps N at $\kappa = m$. □ The construction is verified exactly, including the closed form, on 500 seeded random instances.

Theorem 4 also locates the discount. Two flat singleton classes give $N_k = |C_k|$ and hence a zero discount, so the certification failure of Theorem 3 requires at least three layers, a jointly controlled option, or depth: exactly the features a real stack has. The general characterization of $N = \kappa$ remains open.

4. Prices: standing options and exit costs

Let each realization carry a cost and let the buyer currently operate R_0. The **exit cost** toward an adversary set $A \subseteq B$ is $E(u, A) = \min\{\mathrm{cost}(R_0 \to R) : R \in \mathcal{R}(u), c(R) \cap A = \emptyset\}$, the cheapest re-routing to an A-free realization, with $E = \infty$ if none exists. The edge-level components of $\mathrm{cost}(R_0 \to R)$ are the switching costs of the industrial-organization literature [Farrell and Klemperer 2007].

Theorem 5 (standing-option ceiling). Suppose class k acquires an option o^* with c-trace $\emptyset$: a self-hostable, forkable substrate with fork-and-operate cost E_f. Then (i) no coalition cuts u through class k and (ii) $E(u, A) \leq E_f + \sum_{k' \neq k} E_{k'}$, where $E_{k'}$ are the re-routing costs at the other classes along the cheapest A-free realization; if k was the unique A-exposed class, $E(u, A) \leq E_f$.

Proof. (i) Realizations through o^* have traces that avoid class k’s controllers entirely, so a cutting coalition must work elsewhere. (ii) The realization that forks at k and re-routes minimally at the other classes is feasible and A-free; E is a minimum, hence at most its cost. □

The theorem separates a conceptual claim from an empirical one, and the separation is the point. That open substrates admit self-hosting is definitional; whether the premise holds at a given layer, that is, whether E_f is finite and bearable, is measurable. The applied work’s cross-section (proprietary layers 18:4 against forkable layers 0:3 on the existence of path-independent alternatives, with the exceptions where a moat is domestically owned or standardization has flattened it) is then evidence about where the premise holds, and the theorem says what follows where it does: a price ceiling on exit that no bilateral negotiation can remove, which is an economic object (it disciplines rents today) and a security object (it bounds coercion tomorrow) at once.

5. The graded correspondence

Measured structures are cyclic and their intensities are uncertain, so applied work runs a graded operator instead of exact enumeration: dependency intensities a_{uv}, substitutabilities s_{uv}, effective weights $a_{uv}(1 - s_{uv})$ collected in a matrix A_{eff}, and foreign dependence $d = (I - M)^{-1} A_{\text{eff}} f_0$ with $M = A_{\text{eff}}(I - \text{diag}(f_0))$, in the Leontief tradition [Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi 2012; Baqaee and Farhi 2024]. The correspondence between the operator and the cut calculus is exact at the level of supports.

Proposition 6 (support correspondence). Assume $\rho(M) < 1$. Then (i) $d(u) > 0$ if and only if a marked node is reachable from u along edges with $a_{uv} > 0$ and $s_{uv} < 1$ (ii) every entry of d is nondecreasing in each intensity and nonincreasing in each substitutability; (iii) if every path from u to the marked set crosses an edge with $s_{uv} = 1$, then $d(u) = 0$.

Proof. Expand d as the Neumann series $\sum_{k \geq 0} M^k A_{\text{eff}} f_0$: each term sums nonnegative weights of walks ending at a marked node with unmarked interior, since marked columns are zeroed in M. If a qualifying path exists, its prefix up to the first marked node contributes a positive term, and all terms are nonnegative; with no such path, every walk carries a zero factor, which gives (i). Every term is a polynomial with nonnegative coefficients in the entries of A_{eff}, which gives (ii). An edge with $s = 1$ has zero effective weight and blocks exactly like a deleted edge, so (iii) follows from (i). □
Verified on 300 random cyclic instances (positivity-reachability equivalence, monotonicity under substitutability perturbations, blocking; all passing).

The reading: the support of graded exposure is the reachability calculus on the graph with standing-option edges deleted, and magnitudes interpolate monotonically between the AND-reading and the OR-escape. Reachability conclusions of the applied companion are therefore support facts of d and inherit the exact calculus; magnitude conclusions rest on the ordinal-band discipline reported there.

6. Two open components

6.1 Bundle interactions

Applied measurement finds policy levers to be partial substitutes: on the graded functional, every overlapping pair of levers reduces exposure by less than the sum of its parts, in all 800 Monte-Carlo draws, while a disjoint pair is additive. The tempting theoretical account, that liberated buyers are a coverage objective and coverage is submodular, is wrong: the count of buyers whose every blocked class has been treated is not submodular, because treating a buyer's last blocked class has increasing returns (a completion complementarity). The refined conjecture is that substitution governs the interior of the graded functional while completion effects, if present anywhere, sit at the floor, where the last blocked class of every buyer is the same one. Identifying conditions under which interior substitution is a theorem is open; the empirical sign structure is MC-grade until then.

6.2 Options as public goods

An option's value has a private part, the price discipline it exerts on the buyer's incumbent supplier today, and a public part: its existence insures every buyer at the layer, in every future period, against coercion and repricing. Its maintenance, by contrast, is paid by whoever exercises it, in switching costs and forgone incumbent quality. Myopic buyers therefore under-maintain diversification, and the wedge grows with lock-in: this is the second-sourcing problem [Farrell and Gallini 1988; Shepard 1987] lifted from one buyer-seller pair to a production network, with the insurance component priced as a real option [Dixit and Pindyck 1994]. The equilibrium model (option survival hazards decreasing in exercised demand, buyers choosing suppliers on price, a planner internalizing the cross-buyer and cross-period insurance value) is being developed jointly with a companion game-theoretic project and is stated here as the programme's welfare core: procurement preference and second-sourcing mandates are, in this reading, the planner's payment for option maintenance, and their layer-by-layer justification is exactly where the public-good wedge is largest. The full treatment, together with the other structure-reshaping channels (capital, influence over rule-setters, capture), is the subject of the planned dynamics paper announced in Section 2.

7. Application: the European digital stack

The applied companion measures a 49-category dependency graph of the European digital stack (96 dependency and distribution edges, 19 jurisdictional-reach edges, five European demand nodes; revised 2026-07-13 to expose the client, enterprise-demand, and IoT/edge layers), with per-edge sources and adversarial verification. Encoding it as a production structure (children grouped into functional classes: same class substitutable, distinct classes jointly required; greatest-fixed-point feasibility, which handles the graph's cycles; mixed and contested categories coded under a strict and a lenient controller assignment, with reasons recorded) makes every quantity of Section 3 computable exactly, since the controller universe is small. All results below are identical under both codings, with one exception noted where it arises.

Coercion. Every European demand node has $\kappa = 1$ on the ownership channel at full stack depth, and the list of solo cutters is five blocs long: the United States, and through the semiconductor floor Taiwan (fabrication), Japan (materials and packaging), Korea (memory), and China (raw-material processing). The unique minimal trace is $\{$US, China, Japan, Korea, Taiwan$\}$: at category resolution, every feasible realization of every European buyer touches all five. Above the floor (the semiconductor floor layers waived, isolating what instruments other than capacity investment can move), $\kappa = 1$ still, with the United States a solo cutter for every buyer; the cut binds through the classes where the graph records a single option, network hardware and identity for every buyer, and for some buyers the SaaS or distribution class itself. The enterprise buyer, new in the revision, is additionally China-cuttable above the floor through the contested IoT platform class, under the strict coding only: the one coding-sensitive result. On the legal channel, where controllers are jurisdictions with recorded extraterritorial instruments, both US law and Chinese law are unilateral cutters at full depth (the former through fourteen reached categories including EU-owned lithography, the latter through materials and raw-material processing), and US law alone above the floor.

The certificate against the truth. The public-sector demand node passes a per-layer independent-suppliers test at every required class: its cloud class holds three suppliers under distinct control (a European provider category, the EU-incorporated sovereign ventures, the US hyperscalers) and its productivity class two. Its certificate, $\min_k N_k$ in the weak node-level form, is 2. Its option number is 1, and its coercion number is 1 on both channels. The bloc-spanning discount is positive on measured data, and the mechanism is the one Theorem 3 isolates: the per-layer count registers the EU-incorporated venture as an independent supplier, while every realization through it runs on a licensed US stack, and every realization not through it meets US control at the operating-system, identity, or network-hardware layer. The other four demand nodes fail even the certificate, each at its own single-supplier classes: consumer distribution for the consumer buyer, the productivity class for the sovereign-critical and critical-infrastructure buyers, and the integrator channel and the IoT platform for the enterprise buyer, the majority of private demand. A per-layer test, where Europe would pass it, certifies a diversification the joint structure does not have; where Europe would fail it, the failure is informative about a different layer for each buyer class.

Resolution, measured. The category-level numbers above are aggregation-sensitive: a category is cut when its aggregate is, and majority-coded aggregates hide intra-category options. The applied companion's three-layer pilot measures option sets at supplier resolution (28 options, adversarially verified), and recomputing there splits the picture in two directions at once. For a productivity workload (a suite class over a hosting class, above the floor, measured classes only), supplier resolution reverses the category verdict: a fully self-controlled realization exists, a European suite on European-controlled hosting, so no coalition cuts the workload over the measured classes, where the category graph reported $\kappa = 1$ through the US bloc; aggregation had overstated coercion. For an AI workload the verdict survives every refinement: the accelerator class offers seven exercisable suppliers and exactly one controller (the second, third, and seventh suppliers are all US firms), so the certificate reads one, $\kappa = 1$, and the United States is the solo cutter at supplier resolution just as at category resolution; here a supplier-plurality test overstates diversification by a factor of seven. Both directions are identical under the strict and lenient codings of the two contested rows. The scope is the measured classes (identity and network hardware await supplier-resolution measurement), and the theory says exactly which quantity each added class estimates.

8. Related literature

Clayton, Maggiori and Schreger [2026] develop the general-equilibrium theory of geoeconomic power, coercion exerted through world equilibrium; this paper is its combinatorial micro-structure, answering which coalitions can coerce which buyers as a cut problem and giving diversification its dual. Elliott, Golub and Leduc [2022] study supply networks whose fragility comes from random disruption of thin sourcing; the option number is the strategic counterpart of their sourcing multiplicity, with disruption chosen by an adversary rather than drawn by nature. Liu [2019] locates the sectors where industrial policy has the most leverage through input-output linkages; the coercion number and the floor are the discrete, control-labeled analogue. The switching-cost and lock-in literature [Farrell and Klemperer 2007; Arthur 1989; David 1985] prices the edges; second sourcing [Farrell and Gallini 1988; Shepard 1987] is the two-agent ancestor of the public-good dynamics; weaponized interdependence [Farrell and Newman 2019; Hirschman 1945] supplies the coercion semantics that the cut calculus formalizes. The AND/OR structure itself is standard in other fields, and two non-economic literatures are the calculus's nearest formal kin, acknowledged here so the novelty claim is scoped precisely. Reliability engineering computes minimal cut sets on AND/OR fault trees (which component sets disable a system) and prices, as common-cause failure, exactly the correlation that defeats per-component redundancy counts [Barlow and Proschan 1975]; network interdiction optimizes an attacker's cut under a budget [Wood 1993]. The delta on both: controllers with identity and coalitions replace anonymous failure events, so the object is who can cut and who can diversify, with N as the dual quantity of interest; the certification corollary aims the mathematics at regulatory texts; and the support correspondence connects it to production-network economics. The contribution is therefore the control overlay, the duality pair (κ, N), the discount theorem with its certification corollary, the standing-option ceiling, and the exact computation of all of these on a measured national-scale structure.

9. Conclusion

Strategic autonomy in production networks has an exact combinatorics. Coercion composes as a weakest link and is priced by the smallest cutting coalition; diversification is priced by the largest family of controller-disjoint ways of operating; the two are dual, and the gap between what per-layer certificates count and what the system actually holds is a theorem, with a measured positive instance: the European public-sector buyer passes a two-suppliers test at every layer while one bloc, on either channel, can cut everything. Forkable substrates are the instrument the theory singles out, standing options that cap exit costs by construction; the floor is where no instrument but capacity works; and the open components, interior substitution of levers and the public-good economics of option maintenance, are stated with the evidence that motivates them and the status they currently have. The framework's practical reading is addressed to the texts Europe is now writing into procurement law: their criteria name the right objects, and the objects live one level up from where the criteria measure.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., & Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica* 80(5), 1977–2016.

ANSSI (2022). *Prestataires de services d'informatique en nuage (SecNumCloud): référentiel d'exigences*, version 3.2, 8 March 2022.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99(394), 116–131.

Baqaee, D., & Farhi, E. (2024). Networks, barriers, and trade. *Econometrica* 92(2), 505–541.

Clayton, C., Maggiori, M., & Schreger, J. (2026). A framework for geoeconomics. *Econometrica* 94(1), 105–136.

David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review* 75(2), 332–337.

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press.

Elliott, M., Golub, B., & Leduc, M. V. (2022). Supply network formation and fragility. *American Economic Review* 112(8), 2701–2747.

European Commission (2025). *Cloud Sovereignty Framework*, version 1.2.1, DG Digital Services, October 2025.

Barlow, R. E., & Proschan, F. (1975). *Statistical Theory of Reliability and Life Testing*. Holt, Rinehart and Winston.

Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: how global economic networks shape state coercion. *International Security* 44(1), 42–79.

Farrell, J., & Gallini, N. T. (1988). Second-sourcing as a commitment: monopoly incentives to attract competition. *Quarterly Journal of Economics* 103(4), 673–694.

Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. In M. Armstrong & R. Porter (eds.), *Handbook of Industrial Organization*, vol. 3, 1967–2072.

Fermigier, S. (2026). Path sovereignty: dependence, control, and options in the European digital stack. Working paper.

Hirschman, A. O. (1945). *National Power and the Structure of Foreign Trade*. University of California Press.

Liu, E. (2019). Industrial policies in production networks. *Quarterly Journal of Economics* 134(4), 1883–1948.

Shepard, A. (1987). Licensing to enhance demand for new technologies. *RAND Journal of Economics* 18(3), 360–368.

Wood, R. K. (1993). Deterministic network interdiction. *Mathematical and Computer Modelling* 17(2), 1–18.

Data and reproducibility. Every proved statement is verified by brute-force enumeration on random AND/OR structures in `verify/theory_core.py` (canonical example exact; Theorems 1–3 checked on 399 seeded random instances; Theorem 4(b) checked, including its closed form, on 500 dedicated instances); the application of Section 7 is computed by `verify/cuts_options_model.py` from `data/edges/compute_nodes.csv` and `compute_edges.csv`, with the class map, the floor definition, and the strict and lenient controller codings for mixed categories recorded, with reasons, in `verify/results/cuts_options_v1.json`; the supplier-resolution recomputation is `verify/pilot_discount.py` on `data/pilot/options.csv`, its 28-row coding pinned to the data and its self-checks asserting the headline facts, output in `verify/results/pilot_discount_v1.json`. The measured graph itself, per-edge sources, and the three-layer option-set evidence are documented in the applied companion.