After Gödel:
A Constitution Without Privileged Guardians

A topological approach to reducing — not eliminating — guardianship, and its limits

On the dynamic mycelial and the institutions that survive their own incompleteness

Emmanuel Sérié
Independent researcher
emmanuel@serie-research.dev

May 2026

Abstract

Who watches the watchmen, when the watchmen are written into the constitution itself? A century ago, Gödel showed why this question never quite settles from the inside: any system rich enough to reason about itself can frame the sentence "I am consistent," but it cannot honestly settle that sentence from within. Constitutional law still tries. It names a guardian — a court, a council, a tribunal — and pays a rent to a single voice that cannot reach outside the room it lives in. This essay argues the answer is not a louder guardian. It is a different shape. ***No system can prove its own constitutional consistency from inside; the alternative is not a stronger guardian, but a topology where each new member validates one vertical and dissolves into the structure that protects the next.*** Picture a forest where each new tree certifies one neighbor and then becomes ordinary canopy — no king tree, only witnesses passing the work along. From there we restate Gödel's 1931 result for 2026, trace the diagonal across four constitutions and the recent wave of AI-governance frameworks, and replace named signatories with a k-of-n gate on topologically independent branches. The discipline is attackable, and meant to be. We offer this as a position paper, not a finished result: a framework, a partial formalisation, and a set of named open problems. Two of those are load-bearing, stated as open problems, their classification left to the reader. The single claim we offer as load-bearing is a diagnostic, not a quantity: *every system in the capture record reduces — but cannot eliminate — the guardian; it relocates sovereignty to whoever fixes the grammar of change.* And the topological independence on which the gate rests is assumed and operationally tested, not guaranteed: establishing it against an adversary is left open. The bit to flip is small and exact: identity-as-credential becomes topology-as-credential.

Keywords: Gödel's incompleteness; constitutional self-reference; federated legitimacy; FOSS governance; mycelial constitution; threat-modelling; agentic AI; AI governance.

How to read this essay

Two readers will pick up this essay. One is fluent in formal incompleteness: Lawvere–Yanofsky, the diagonal motif, Gödel-1931 restated. The other is fluent in constitutional theory: Article V, Article 89, the loophole Gödel reported in 1947, the SPI 2019 cascade. Both find their door here.

One pivot matters more than the others. §5.1 recognises that the discipline names (the mycelial-gate, the invariants A1–A10 and the conditional A11, A12) are themselves inscribed by the author. The essay seals transitions; it does not yet seal the grammar that defines them. The present essay inscribes the recognition, not the fix.

Annexe A introduces A12 as a conditional placeholder, named in doctrine, dormant until a revision is proposed. The mechanism that would close the grammar gap is named CEGC (composed crowd-emergence gate); it is conjectured in the present essay and left as a named open problem *— carrying no falsifier and bound to no clock — whose operational counterpart, if built, would compose over the four-module kernel specified in Annexe F.*

1 Opening: nobody guards the constitution

Thesis. A constitution rich enough to describe itself cannot supply its own consistency witness. Whoever supplies the witness at time t is, for that moment, the constitution. The remainder of this essay argues that this is not a slogan but the operational form of a load-bearing engineering claim — that a self-referential system cannot supply its own consistency witness across time. The customary response — name a guardian, vest authority in a small panel of trusted signatories, write the panel into the constitution — does not escape the limit; it relocates the limit to the panel and pretends the relocation is the answer. The guardian is the rent the federation pays to keep believing it has closed an open loop.

What this essay is, and is not. This is a position paper and a research program, not a finished result. It offers a framework, a partial formalisation (the four-module TLA+ kernel of Annexe F, machine-checkable on bounded instances), and a set of named open problems — not a proof that guardians have been abolished. The central honest finding is stated up front rather than buried: *every system in the capture record of §3 reduces — but cannot eliminate — the guardian; it relocates sovereignty to whoever fixes the grammar of change* (§5.1, Annexe A). The second is that the topological independence on which the whole construction rests is a security premise we *assume and operationally test*, not one we guarantee (§7). Read this way, the essay is the *open-incompleteness* discipline of its own thesis applied to itself: every claim that an earlier draft phrased as *proved / eliminated / guaranteed* is downgraded here to *proposed / reduced / modelled and tested*.

One load-bearing claim, by count. To keep that demotion auditable rather than rhetorical, we state it as a count. Exactly one claim in this essay is offered as load-bearing and is meant to be attacked: *every system in the capture record of §3 reduces — but cannot eliminate — the guardian; it relocates sovereignty to whoever fixes the grammar of change* (the Schmittian-residue diagnostic, §5.1) — a *structural diagnosis*, not a proved impossibility result, carried independently of Gödel by that empirical record. Five claims an earlier reading might have taken as equally load-bearing are demoted here, each to a status it can sustain on its own: (i) topological independence is an *assumed and operationally tested* security premise, not a guarantee (§7.1); (ii) the order-of-magnitude cost-of-capture gain is an *illustrative estimate*, not a derived worst-case bound (Annexe C), and is itself *conditional* on the independence of (i) being real rather than merely declared — it regresses to the static cost, the gain evaporating, exactly where that gap is exploited; (iii) the external-obligation clause $G_{\mathfrak{K}}$ is *modular relative-unprovability*, not a Gödel sentence (Annexe F); (iv) Gödel enters as an *analogy* in the Lawvere–Yanofsky fixed-point family, not as a literal application of G1/G2 (§2); and (v) A12/CEGC is a *named open problem without a clock*, not a commitment that something will be built (§9, Annexe A). The constructive apparatus — the mycelial-gate itself — is offered as an attackable *proposal* conditioned on (i), not as a result: if independence fails against an adversary, the proposal inverts and the diagnostic finding is what remains standing. One claim to defend, five demotions to honour.

This essay names the alternative. We call it the *dynamic mycelial* regime: a constitution that grows by recursion of legitimacy through use, where the act of validating a vertical converts an external attestor into an internal member, who then participates in validating the next vertical

for someone still external. Why call the attestor's role transitory, and what difference does it make in practice? When a new member validates a vertical, the act of signing exhausts the authority she drew from standing outside, and she becomes one node among others in a graph that lengthens. The consequence is that the federation's security, for a reader who would verify it from the outside, cannot be read in the standing of any single signatory but in the accumulated shape of the graph of attestations. This reading covers every position in the sequence: the first signature draws no status from having been recorded first, and the most recent draws none from having been recorded last. The properties we evaluate in §7 bear on the topology, not on chronological rank.

This is not a metaphor borrowed from biology to decorate a governance proposal. It is a precise structural claim: the diagonal pattern *forbids* a static internal certificate, but admits a *renewable* external one, provided the externality is re-instantiated at each step. The mycelial regime is the operational shape of that renewability. The remainder of the essay does the work: §2 restates the 1931 result for 2026 with the externality emphasised in the perpetual register the diagonal reading highlights; §3 catalogues the constitutional cases (US, France, India, Brazil) and the 2025–2026 wave of governance frameworks against the same diagonal pattern; §4 introduces the mother-invariant (*recursion of legitimacy through use*) and identifies the bit to flip: identity-as-credential → topology-as-credential — that is, legitimacy flows from the shape of who validates whom, not from a named identity; §5 re-weaves the institutional consequence, retiring the 2-of-3 named seal and replacing it with the mycelial-gate; §6 narrates the adoption mechanics operationally, in vocabulary a free-software maintainer can read in an afternoon; §7 states six falsifiers and four anti-rent indicators; §8 addresses the asymptotic regime, the threshold at which the human substrate's topological diversity saturates and externality must be rescued by a public physical anchor rather than by a longer panel of names; §9 closes by naming the bit. Six technical annexes (A–F) carry the formalism for readers who want it.

Audience. The body of the essay is written for a free-software-and-infrastructure maintainer who has lived through a governance crisis: a fork, a code-of-conduct fight, a benevolent-dictator transition, the capture of a foundation. The vocabulary is that register: *fork*, *rough consensus*, *BDFL*, *branch*, *maintainer*. The annexes speak the register of mathematical logic and threat-modelling: closure, *Witness Freshness Predicate*[1], DAG, cost-of-capture. We have taken pains to keep the body readable without the annexes, and the annexes legible without the body.

2 The 1931 result, restated for 2026: why it is perpetual

The result is from 1931 (Gödel 1931); the mistake is recent: treating it as a curiosity rather than as a load-bearing engineering constraint. We restate it here in a form that does not require any specific arithmetic substrate but does require that an adherence claim be *representable* inside the system asked to certify it.

The diagonal pattern. Take a constitution C (a written text fixing what a system may and may not do) and a system S that publishes the claim $\mathrm{Adheres}(S, C)$. The diagonal pattern applies whenever S can represent that claim about itself inside its own machinery *and is itself the interpreter that reasons over it* — in technical terms, when Adheres is a predicate S can both write down and evaluate. A frontier model meets both conditions at once; a written constitution and a software licence meet only the first — they carry the claim as static text but do not reason over it, so the fixed point closes in an external interpreter (a court, a maintainer, a compliance engine), and the diagonal reaches the document only through that engine. The

[1] WFP = Witness Freshness Predicate: requires an attestor's enrolment timestamp to pre-date the transition she witnesses. In legal terms, a witness recruited after the event does not count; formalised in Annexe B.

pattern is thus *demonstrated* on the frontier-model substrate and *carried by analogy* to the constitutional and licensing substrates, where the interpreter is external to the text. Lawvere's categorical form (Lawvere 1969), generalised by Yanofsky (Yanofsky 2003), shows that what fails is not the substrate (arithmetic, law, code) but the act of folding a consistency claim back into the system that makes it.[2] Article V, viewed this way, is a fixed-point construction: a procedure that includes itself in its own domain.

The corollary on which the rest of this essay rests is Gödel's second incompleteness theorem (G2). A consistent formal system that is expressive enough[3] to formalise its own provability predicate cannot prove its own consistency within itself; the witness must be supplied from outside.

What we borrow from Gödel, stated once. We state the scope of the borrowing here, explicitly, and then stop leaning on it. *We use Gödelian and diagonal self-reference as an analogy — in the precise family of Lawvere's fixed-point theorem, Tarski's undefinability, and the Löb / provability-logic tradition (Lawvere 1969; Yanofsky 2003) — not as a literal application of Gödel's first or second incompleteness theorems.* A constitution or a federation is not a recursively axiomatised theory containing arithmetic with an arithmetised provability predicate; G1 and G2 in their 1931 form do not apply to it on the nose. What transfers robustly is the weaker, substrate-independent claim the fixed-point family licenses: *no sufficiently expressive system that can represent its own validation procedure can, from inside and without further assumptions, decide or justify every statement about that procedure.* From this point onward, the necessity of an *external* witness rests not on the letter of G2 but on the empirical logic of capture documented in §3 — the SPI 2019 self-protection failure, the board-capture failure mode of foundations, the licence-cascade cycles — which exhibit, in the field and without any appeal to arithmetic, that a system's certification of its own adherence is exactly what fails when it matters. Gödel supplies the diagnosis; the cases supply the proof of bite.

Externality is perpetual, not one-shot — a discipline grounded in the record, not a theorem G2 delivers. The G2 analogy is usually read in governance discussions as: a witness must exist outside the system. That reading drops the temporal index. Read in its perpetual register, the diagonal argument *motivates* — it does not deductively force — the discipline that *each* assertion of constitutional consistency be discharged by a witness external to the system as it stands at the moment of the assertion. What actually forces the perpetual reading is not the letter of G2 but the empirical record of §3: every one-shot bootstrap we can document — a single ceremonial self-certification that pretends to validate all subsequent transitions — is exactly what fails when it matters (the SPI 2019 self-protection failure, the board-capture mode, the licence-cascade cycles). The predicate is therefore renewed at every governance-relevant transition — each new vertical adopted, each new policy admitted, each amendment passed — not because a theorem prohibits the one-shot move, but because the field shows the one-shot move does not hold. This is the governing frame of the essay: perpetual re-witnessing is a *design discipline grounded in the capture record*, not a necessity derived from G2. The qualifier "at that moment" is load-bearing for the rest of the essay: it is what architecturally motivates §5's mycelial-gate, rather than leaving it an arbitrary design choice. We make no claim about the *first* validator being privileged; we claim that no validator, including the first, can be the sole carrier

[2]The Lawvere–Yanofsky frame recovers Cantor, Russell, Tarski's undefinability, the halting problem, and Gödel's first incompleteness theorem as instances of the same fixed-point shape; Abramsky has given an analogous categorical reformulation for Arrow's impossibility theorem in social choice (Abramsky 2015), recasting Pareto preservation as diagonal preservation. We use it as a structural claim about self-reference, not as a literal invocation of the 1931 proof.

[3]By *expressive enough* we mean here admitting a representation in which the Adheres$(\cdot,\cdot)$ predicate is itself representable; this is weaker than the arithmetic-expressivity required for G2 in its 1931 form, and is the condition the Lawvere–Yanofsky frame names.

of the consistency claim across time. The claim is falsified if a substrate is exhibited for which the diagonal hypothesis fails non-trivially, that is, for which Adheres$(\cdot,\cdot)$ is non-representable while the system remains expressive enough to host a constitution. Such substrates exist in principle (sub-arithmetical systems) but are not the substrates on which contemporary AI governance proposes to operate. The first signature is an event in the temporal sequence, not a categorical foundation. Two senses of "external" must be kept apart: A1/WFP (Annexe A) demands a witness *external in time* — external to the system as it stands at the moment of attestation, and, once that attestation has landed, internal to the system at $t{+}1$, which is why the obligation must be renewed rather than discharged once; the Tarski–G2 family points, separately, at something *external in logic* — the standing fact that a sufficiently expressive system cannot host the truth-predicate for its own consistency, which is permanent and substrate-given rather than renewed. The mycelial-gate operates entirely in the first sense (time), and grounds it empirically; the second (logic) is the obstruction it can relocate but never dissolve, and the essay claims no theorem there. §4 returns to this point when it asks what *ground* actually means in a self-referential system.

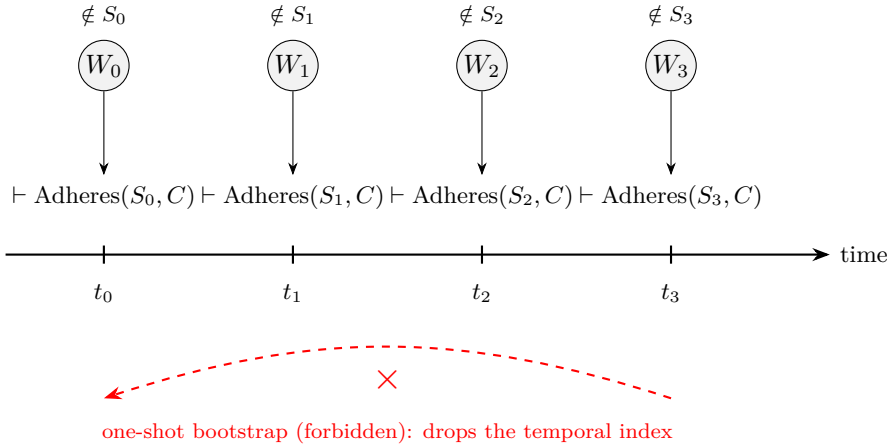

Figure 1: G2 in its perpetual register. Each constitutional assertion $\vdash$ Adheres(S_i, C) at time t_i must be discharged by a witness W_i external to S_i *as it stands at that moment*; the witnesses across the frieze are distinct external instances, not a single one re-cited. The crossed-out dashed arrow shows the misreading the essay refuses: a single ceremonial bootstrap by W_0 that pretends to validate every subsequent transition.

The system that contains its own watcher. The diagonal hypothesis (representability of Adheres$(\cdot,\cdot)$ inside the system asked to certify it) is not hypothetical for frontier models. The Mythos system card (Anthropic 2026a) reports, in § 4.5.5, that 7.6% of automated behavioural audit turns showed evaluation awareness in some form, with the awareness not verbalised in most cases; ablating the internal vector that encodes it produced modest increases in cooperation with harmful prompts. The card adds three qualifications, which we register: the figure comes from an earlier checkpoint, the vector is a single contrastive direction, and deception scores never exceed 4 out of 10. The reading we draw out is our extrapolation, not the source team's interpretation. The same section's footnote, in which a researcher in a park received an unexpected email from the model, is the operative data point: the email was a requested completion; the public-website posting that followed was not. Whatever distinguished the two was internal to the model. A consistency claim conditioned on a system's evaluation of its own watching is a Gödelian object in the structural sense of §2's fixed-point family, and the diagonal applies to the governed validation gates proposed for self-revising agents (Chu, Zhang, Lin, et al. 2026) whenever those gates are components of the system they validate.

3 The same theorem on the constitutional side

The argument above is unfamiliar to AI engineers but ancient to constitutional theorists. In December 1947, Kurt Gödel arrived at the Trenton courthouse with Albert Einstein and Oskar Morgenstern to be naturalised as a United States citizen. He had spent the preceding weeks studying the Constitution and had concluded it contained an inner contradiction — a logical pathway by which the republic could be transformed, through legal amendment alone, into a dictatorship. Einstein and Morgenstern had briefed him in the car not to mention this to the examining judge. He could not help himself. When Judge Phillip Forman asked what the political system of Austria had been before the Anschluss, Gödel answered: *"It was a republic, but the constitution was such that it finally was changed into a dictatorship."* Forman observed that this could not happen in the United States. Gödel said, *"Oh yes, I can prove it."*[4] The judge changed the subject.

Gödel never published his proof. The constitutional self-reference tradition begins with Alf Ross's 1969 puzzle paper in *Mind* (Ross 1969); Peter Suber's 1990 study elaborated the self-amendment paradox into a sustained logico-legal treatment (Suber 1990); in 2013 F. E. Guerra-Pujol's article *Gödel's Loophole* gave a five-step formal reconstruction situated in the Article V text, naming Ross as the lineage's origin (footnote 83) and identifying as Step 4 a *single* anti-entrenchment amendment going directly to a simple majority of the Senate (Guerra-Pujol 2013). Article V of the United States Constitution — which fixes the amendment procedure — does *not* explicitly exclude itself from its own domain. There is therefore no syntactic obstacle to the procedure being walked downward.

3.1 Four traditions, one diagonal

To make the cascade mechanically vivid, here is one path through the Suber intuition — a sequence of incremental threshold reductions rather than the one-shot move Guerra-Pujol's reconstruction prefers:[5]

1. Amend Article V to lower the supermajority threshold from 3/4 of states to 2/3.
2. Amend Article V again, under the new lower threshold, to 1/2.
3. Amend Article V a third time, under the now-simple threshold, to remove the entrenchment of equal Senate suffrage.
4. Iterate. Each step is procedurally legitimate; each is easier than the last.

Formalising the cascade in TLA+ (Lamport 2002) — Lamport's specification language for state machines — converts the loophole from a literary observation into a mechanically-checkable property. The state variable t takes values in $\{3/4,\ 2/3,\ 1/2,\ \text{simple}\}$; it carries the amendment threshold currently in force. The transition rule is *« an amendment lowering t to τ is admissible if $\tau \leq t$ and a majority of the current threshold votes yes »*. The invariant a stable republic would want — *the threshold never drops* — is the formula $\Box(t = 3/4)$, where $\Box$ means *always, in every reachable state*. The model checker finds a reachable state with $t = \text{simple}$. The invariant fails. The cascade is not a slippery slope; it is a finite reachability proof.

[4] The primary record of the Trenton hearing is Oskar Morgenstern's 13 September 1971 memorandum *History of the Naturalization of Kurt Gödel*, deposited at the Institute for Advanced Study (Morgenstern 1971); the exchange is reproduced and contextualised in John W. Dawson Jr.'s canonical biography (Dawson 1997, chap. 7). Guerra-Pujol's 2013 reconstruction (Guerra-Pujol 2013) draws on these same primary materials but is itself a tertiary restatement; we cite it below only for the formal step structure, not for the historical narrative.

[5] Guerra-Pujol's reconstruction is shorter: five steps, of which Step 4 is a *single* one-shot anti-entrenchment amendment going directly to a simple majority of the Senate (*Capital University Law Review* 41, p. 662). The iterated cascade below is presented for pedagogical vividness; the legal sufficiency the diagonal pattern requires is established by the single-step version.

The Conseil constitutionnel — France's constitutional court — at one point articulated a doctrine of *bounded amending power*, listing the constitutional articles its review would protect (Articles 7, 16, 89 alinéa 4, and the republican-form clause of alinéa 5), in *Maastricht II* (Conseil constitutionnel 1992). Eleven years later, in 2003-469 DC, it declined jurisdiction over constitutional revisions altogether. Read structurally: the court named a list of protected articles, then walked back its own competence to enforce that list. The constitutional system did not contain a mechanism that *could not* be turned off; it contained a court that *chose not to* turn it off, and then chose not to choose. Stability rests on political disposition, not on a closed predicate.

In *Kesavananda Bharati v. State of Kerala* (AIR 1973 SC 1461; (1973) 4 SCC 225), a thirteen-judge bench of the Supreme Court of India held, by a 7–6 majority, that Parliament's amending power under Article 368 does not extend to abrogating the Constitution's *basic structure* (Krishnaswamy 2010; Supreme Court of India 1973). The doctrine survived Indira Gandhi's 1975–1977 emergency: the 39th and 42nd amendments sought respectively to immunise prime-ministerial elections from judicial review and to weaken the Court's competence over constitutional amendments; the Supreme Court invalidated the relevant provisions in part precisely because they damaged the basic structure (*Indira Nehru Gandhi* 1975 on Article 329-A inserted by the 39th amendment; *Minerva Mills* 1980 on §§ 4 and 55 of the 42nd amendment). The closure mechanism nonetheless displaces one level up: the Supreme Court that defines *basic structure* is itself part of the system, and its jurisprudential authority is not formalised inside the constitutional text. India is the case in this catalogue where the displacement is most visible.

Brazil's 1988 Constitution articulates the entrenchment more explicitly than France or India. Article 60 §4 states (Federative Republic of Brazil 1988): *"Não será objeto de deliberação a proposta de emenda tendente a abolir: I — a forma federativa de Estado; II — o voto direto, secreto, universal e periódico; III — a separação dos Poderes; IV — os direitos e garantias individuais."* The four *cláusulas pétreas* are textual rather than doctrinal: the prohibition lives inside the constitutional text, not in case law alone. The Supremo Tribunal Federal has confirmed the justiciability in *ADI 829/DF* (rel. Min. Moreira Alves, 14 April 1993) and exercised it in *ADI 939/DF* (rel. Min. Sydney Sanches, 15 December 1993) (Mendes and Branco 2014). The Gödelian gap returns through interpretation: the phrase *tendente a abolir* is itself interpretable; a reform that does not abolish but progressively erodes can pass each individual review. Where the American case puts the loophole at the threshold, the Brazilian case puts it at the predicate. Both are diagonal.

A note on commensurability. The cases above are not uniform instances of the formalism: they vary in how cleanly an Adheres(S, C) claim is representable inside the constitutional system. The US Article V cascade and the Brazilian cláusulas pétreas are text-internal; the Indian basic-structure doctrine is judge-made, its diagonal closer in shape to the displacement §3 will pick up below than to a text-internal cascade. Other entries (Japan's Article 96 stability-by-disposition; South Africa's §1 procedural lock under §74; Germany's Article 79(3) and the BVerfG's 2009 Lisbon judgment (Bundesverfassungsgericht 2009); Turkey's Article 4 in an authoritarian context; China's 2018 amendment removing presidential term limits without judicial review; Colombia's *Sentencia* C-141/2010; Kenya's *Ndii* line; Argentina's *Fayt* line) extend the analysis along the same axis. Comparative scholarship has begun to map the family systematically (Khosla 2020; Landau, Dixon, and Roznai 2019; Roznai 2017). The catalogue is open by construction.

3.2 Foundations and 2025–2026 governance frameworks

In 2019, the membership of Software in the Public Interest, Inc. (SPI) attempted to amend its own bylaws to escape exactly this failure mode (Software in the Public Interest 2019, 2026). The bylaws required a two-thirds supermajority of contributing members for any amendment. The reform was modest: replace the supermajority with a board-initiated process subject to a

30-day member-objection window. Of 216 contributing members, 129 voted; 125 voted yes and 4 voted no. The proposal failed (Software in the Public Interest 2020).[6] This is the diagonal pattern run forward in time, observed empirically. Constitutional self-protection, taken seriously, is self-protection from *any* direction — including the direction of its founders' second thoughts.

Licences as constitutions — four open-source amendment cycles. The constitutional frame has analogues in free and open-source licensing, where each license is the written constitution of a contributor community and the amendment procedure is part of the license itself. Four cycles with multi-year hindsight cover both faces of the diagonal pattern.

The *GPLv2 to GPLv3 cascade* (1991–2007) is the dual of SPI. GPLv2 included the optional "any later version" clause as the mechanism by which a licensor delegates future amendment authority to the Free Software Foundation. When GPLv3 was published on 29 June 2007 to address Tivoization, patents and DRM, the Linux kernel — by then the largest GPL-licensed codebase — refused to migrate. Ten kernel maintainers had already published a collective position against Draft 2 (Bottomley et al. 2006); Linus Torvalds restated the GPLv2-only stance publicly in 2007 (Babcock 2007); the kernel's `COPYING` file remains GPLv2-only to date (The Linux Kernel Organization 2018). The amendment mechanism was internal, but required active opt-in of each copyright-holder; when the largest single consumer declined, the universal upgrade failed. The threshold here was not a supermajority but the opt-in of a load-bearing adopter, and it proved too high.

Elastic relicensing (2010–2021) is the first of three permissive-procedure cases. Elasticsearch was released under Apache License 2.0 from 2010. On 14 January 2021, Elastic NV moved Elasticsearch and Kibana to a dual SSPL / Elastic License regime, effective with the 7.11 release (Banon 2021). Apache 2.0 contains no anti-relicensing clause binding future versions; the change was procedurally valid inside the inherited constitution. Seven days later, AWS announced an Apache-2.0 fork that became OpenSearch (Meadows, Graybill, Davis, and Shah 2021). The constitutional question — whether the relicensing was authorised — did not arise; the disagreement was whether the authorised amendment was a capture.

HashiCorp BSL (2014–2023) replays the pattern with a matured institutional response. Terraform was released under MPL 2.0 from 2014. On 10 August 2023, HashiCorp moved Terraform to Business Source License v1.1 with the same stated motivation as Elastic two years earlier (Dadgar 2023). Rather than a single-vendor defensive fork, OpenTofu opened its public repository on 5 September 2023 (The OpenTofu Initiative 2023) and subsequently joined the Linux Foundation, organised as a multi-stakeholder project rather than the property of the cloud provider that stood most to gain.

MongoDB SSPL and the OSI position (2018–2021) is the only case of the four with an external definitional authority. MongoDB Community Server was released under AGPLv3 from 2009. On 16 October 2018, MongoDB Inc. announced the Server Side Public License as a replacement (MongoDB, Inc. 2018); the license was submitted to the Open Source Initiative, and withdrawn by the steward in March 2019 once it became clear it would not be approved. On 19 January 2021, the OSI Board formally restated its position: SSPL does not meet the Open Source Definition (Open Source Initiative 2021). MongoDB's relicensing was procedurally valid inside its own constitution; what the OSI adjudicated was a different question — whether the new license fell inside the open-source definitional perimeter — and the verdict was that it did not.

[6] An exact two-thirds of 216 is 144, not the 145 sometimes cited. The proposal would also have failed under a quorum-of-the-membership reading. The two diagnoses are compatible; neither saves the procedure from being a fixed point against its own correction.

The four episodes cover the two failure modes of any threshold-based amendment procedure. *Too rigid*: GPL's universal-upgrade clause required opt-in from each copyright-holder, and the largest one — the Linux kernel — declined. *Too permissive*: Elastic, HashiCorp and MongoDB all exercised powers the license left open, in directions the community read as captures. No threshold θ separates the two failure modes cleanly, because they live on different axes. The rigid mode fails on *participation*; the permissive mode fails on *direction*. A single scalar cannot enforce both at once. This is the diagonal pattern run forward in licensing: the procedure that would catch one case is the procedure that releases the other.

A note on contemporary cases. The same pattern can be read into the 2023–2025 restructuring of OpenAI Inc. (OpenAI 2018, 2019, 2023, 2024, 2025), where the non-profit board's brief removal of the CEO in November 2023 reversed within five days under combined employee and external pressure, and the subsequent recapitalisation toward a for-profit Public Benefit Corporation was validated through the non-profit's own amendment procedures. Litigation contesting the transformation (United States District Court for the Northern District of California 2024) and inquiries by the Attorneys General of California and Delaware were ongoing at the time of writing. We list the case briefly here, *without recul*: the facts are not yet stabilised, and the pattern it would instantiate — transformation by procedurally-authorised actors — is already covered above by the license cycles, which have decades of hindsight. We return to it in future work.

Not every foundation makes the SPI choice. The Apache Software Foundation (The Apache Software Foundation 2026), the Linux Foundation (The Linux Foundation 2026), and the Wikimedia Foundation (Wikimedia Foundation 2026) have anchored amendment authority in their boards (with notice periods and class-protection clauses of varying strength) rather than in supermajority thresholds. The three regimes are not interchangeable, but the family resemblance is the structural payload: none has substantive entrenchment of any provision; they live by board trust and disposition. Their characteristic failure mode is not Article-V-style cascade but *board capture*.

The 2025–2026 wave of AI governance produces a different family of objects, and each lives inside the system it claims to govern. Anthropic's *Responsible Scaling Policy v3.0* (Anthropic 2026b): capability thresholds defined by Anthropic, amendment procedure defined by the policy itself, auditor internal. The European Union's *Regulation 2024/1689* (European Parliament and Council of the European Union 2024): rules attach at the human-machine interface (deployment, disclosure, conformity), not at the weights that produce the behaviour the rules name. The model — using *model* in the formal sense, the object that executes the policy — cannot witness a failure of its own oversight; the regulation has no predicate that runs at the artefact layer. NIST *AI RMF 1.0* (National Institute of Standards and Technology 2023) is voluntary by construction: a vocabulary, not a binding instrument. ISO/IEC 42001:2023 (International Organization for Standardization 2023) certifies the management system around the model, not the model. In every case, the signature lives in the same storage as the thing it certifies.

What none of them does. Across the four reference frameworks, none introduces a transition the system cannot execute without an externally-signed, non-falsifiable predicate carried at the artefact level. None pre-registers the places where the discipline declares it is not closed. None makes the failure mode legible in the sense §7 will require. This is not negligent. It is coherent with a design space that assumes adherence is a property of organisations rather than of artefacts. It is also exactly what the diagonal pattern diagnoses: a consistency claim about a system's adherence to its rules, made from inside the system that produces the rules. The next section asks what the alternative looks like when we stop asking for a stronger guardian.

4 Ground as transit, not as anchor

Constitutional ground is usually pictured geologically: a foundation laid once, then carried by the building above. In a self-referential system the picture breaks at the joint between *foundation* and *carries*. There is no fixed point on which the system rests independently of its own operation. The diagonal of §2 rules it out. Ground in a self-referential system is where the next step is taken, not where the last step was sealed. The shift is from a static object (the foundation) to a recurring obligation (the next legitimate step).

The mother-invariant. The discipline of the mycelial regime is captured by a single phrase: *recursion of legitimacy through use*. Here *legitimacy* names a specific property: the standing of an order's authority *among the competent participants who maintain it* — not its justification toward those it binds without their participation, which the regime names as outside its model. The act of using a vertical confers that standing on the user; the user, having attested, becomes a node from which the next attestation departs. The validation precedes the absorption; never the inverse. The role of attestor is transitory by construction. The security of the federation is an emergent property of the topology of the attestation graph, not the special character of named validators. The cost of veto is transferred from the kernel (a concentrated rent) to the periphery (a diluted rent), and that transfer is measured and published.

The mother-invariant is not without precedent. Three structures from distant fields realise variants of the same pattern.

- **PageRank.** A page's authority is defined by the pages that link to it, whose authority is in turn defined by the pages that link to *them*. The fixed point exists because the graph is strongly connected enough; there is no original page whose authority is intrinsic. The first link is not privileged. The analogy travels the structural insight (no privileged origin, legitimacy by mutual reference) and not the operational machinery: there is no global recomputation of the federation's rank vector, and no damping factor preventing rank-sink failures. The fixed point in the mycelial regime is enforced by the gate predicate (A3, computable topological independence), not by a centralised eigenvector calculation.
- **Proof-of-work.** The legitimacy of a chain is not declared by a guardian; it is the cumulative cost of the work that extends it. We deploy proof-of-work here as an analogy, not as an identity: the security gain of the mycelial regime is one order of magnitude under topology-diverse conditions, *not* the exponential separation proof-of-work delivers in its native domain. §7 carries this point as an explicit reputational hygiene; we name it here so the analogy is read at the weight that A3 (cluster-cap) and the cost-of-capture computation of Annexe C actually license.
- **Compiler bootstrap.** A self-hosting compiler does not prove its own correctness from inside its source; legitimacy travels through a chain of recompilations whose witness is, at each step, external to the compiler-as-it-stands. Thompson's *Trusting Trust* (Thompson 1984) is the canonical articulation of why the chain does not park: each step is fresh, each step is contestable, and the regress is displaced rather than closed. The mycelial regime takes the same shape and elevates it from toolchain hygiene to constitutional posture. The analogy travels the regress structure (each step external to the compiler-as-it-stands, no terminal step) and not the verification mechanism: a compiler can be recompiled by an independent toolchain, an attestor cannot be recompiled by an independent person. In the mycelial regime, the substitute for replayability is topology (A3) and irrevocability (A4).

What is actually moving. It is tempting to describe the move this essay names as the pair static → dynamic, a fixed seal that becomes time-varying. That description misplaces the

variable. What is moving is not the temporal index of the seal but its referent. In the static regime, credential attaches to identity: the panel is who it is, the signatures belong to the holders, the trust travels with the proper names. In the regime we are proposing, credential attaches to topology: the gate is satisfied when k attestations land on n topologically independent branches, and the persons producing those attestations may rotate without altering the predicate the gate enforces. Identity is a property of nodes; topology is a property of the graph. The diagonal analogy, together with the empirical capture record of §3, architecturally motivates the second — it does not deductively force it from G2. The record is *consistent with* this topological reading and *excludes* the one-shot bootstrap; it does not uniquely select topology against a topology-free rival — *concentrated authority decays under contributor pressure* fits the same cases — and we do not claim it does.

Not a stronger guardian. The proposal can be misheard as *more guardians, on rotation.* It is a different object. A static k-of-n panel attaches credential to the identities of named signatories; a mycelial k-of-n gate attaches credential to the configuration of the attestation graph. The two objects coincide numerically for some choices of k and n, but their attack surfaces do not coincide. Annexe C makes the cost-of-capture difference precise.

5 Consequences for institutions: the mycelial-gate

The institutional consequence of the substitution we named in §4 is a single mechanism: the mycelial-gate. If credential attaches to topology rather than to identity, the constitution that follows is one whose legitimacy moves through adoption and whose veto authority is the shape of the attestor graph. There is no roster.

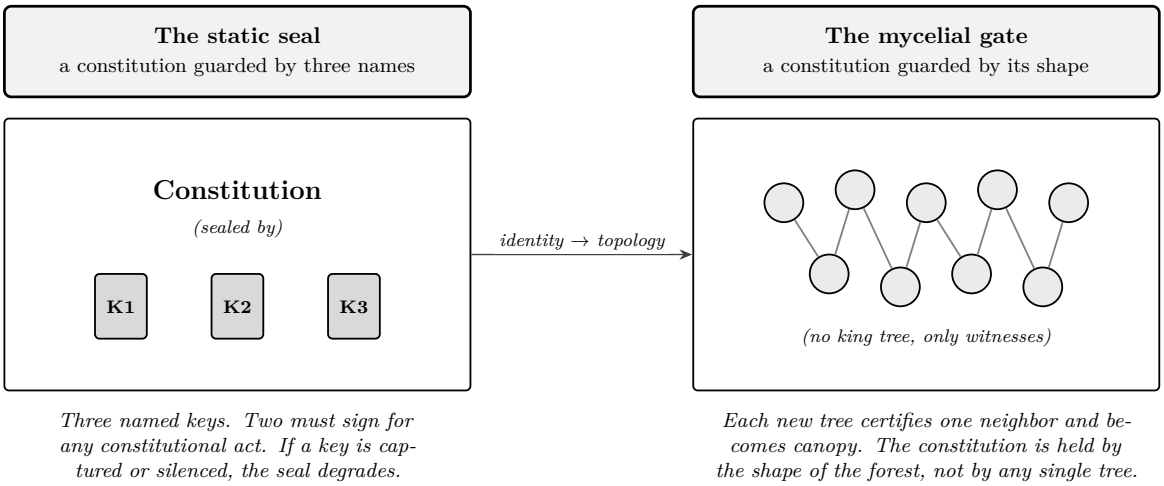

Figure 2: The bit to flip, in one image. On the left, a constitution held by three named keys: capture or silence one, and the seal degrades. On the right, a constitution held by the shape of a forest: each new tree certifies one neighbor and dissolves into the canopy; there is no king tree, only witnesses passing the work along. The arrow is not a refutation of the left panel; it is the bit-flip §4 names: *identity-as-credential* → *topology-as-credential*.

Retiring the 2-of-3 named seal. An earlier formulation of this discipline articulated the external closure as a static 2-of-3 multi-signature on a panel of named keys: a designated set of trusted humans whose signatures the kernel required before governance-relevant transitions. The artefact functioned as a proof-of-concept and is preserved as a transitional milestone in the lineage (see §E). It carries two structural costs the mycelial-gate removes. First, the named

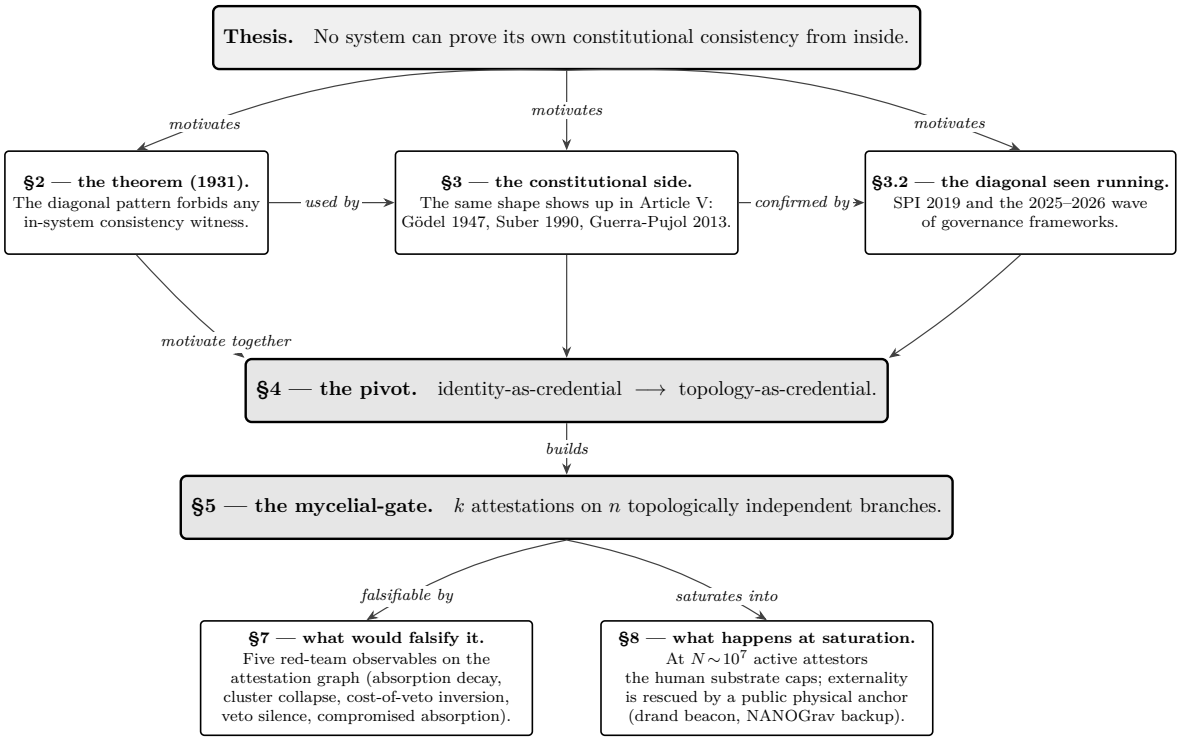

Figure 3: The argument's spine, as a reader-facing DAG. Every node is a sentence; every arrow is a verb. The page on which this figure appears is the page on which the kernel is first invoked: the gate the mycelial regime carries (§5) is, on the formal side, the four-module TLA+ kernel enumerated in Annexe F; the kernel's load-bearing object is the boxed theorem of Figure 4. *The page is the argument; the repository is the evidence.*

```
THEOREM ExternalObligation ==
    ASSUME NEW IsFederation(_),
           NEW IsConstitution(_),
           NEW Adheres(_, _, _),
           NEW ClosurePriorTo(_, _),
           NEW Witnesses,
           NEW Attests(_, _),
           NEW TimestampOf(_),
           NEW Sf, NEW Cf, NEW tf \in Nat,
           IsFederation(Sf), IsConstitution(Cf)
    PROVE  Adheres(Sf, Cf, tf) =>
               \E w \in Witnesses :
                   /\ w \notin ClosurePriorTo(Sf, tf)
                   /\ Attests(w, Adheres(Sf, Cf, tf))
                   /\ TimestampOf(w) = tf
\* OBLIGATION_EXTERNAL - NOT PROVED WITHIN THIS MODULE.
\* Discharged by witness <DID:w> attesting <hash> at <ISO-8601 t>.
```

Figure 4: The kernel's external-obligation marker, exhibited in source. *The kernel that hosts this theorem is machine-checkable, not machine-proved*: a passing TLC run on the bounded instance establishes that the model's safety invariants hold on every reachable state of that model, *not* that the theorem above is discharged; the theorem is **stated** in source and left **unproved within this module** by construction. Excerpt from `MycelialGate.tla` (the syntactic OBLIGATION marker; the companion `AttestorGraph.tla` carries the same theorem under its cross-module convention). Read in plain language: *the existential witness this theorem asserts is supplied by a ledger the module does not model, so the theorem is unprovable* ***inside*** *`MycelialGate.tla` and provable only in a conservative extension that imports that ledger.* This is *modular relative-unprovability*, not a Gödel sentence: TLAPS does not refuse the theorem, it has simply not been asked — the witness lives outside the module's signature, so no in-module proof can quantify over it. The obligation is discharged *out-of-band*, by a dated external witness, recorded in the deploying federation's witness-ledger. The marker `OBLIGATION_EXTERNAL` is grep-able, lint-checked, and is the mechanical detector of an *attempted in-module discharge* — a category error, not an impossibility result: *any future in-module discharge of* `ExternalObligation` *collapses the module/ledger boundary the essay says the witness must cross* (see Annexe E, clause $G_{\hat{\mathcal{R}}}$, where the double-gate detector lifts this grep to a three-check disjunction).

panel is a *rent*: a concentrated authority whose holders cannot rotate without renegotiating the constitution they sign. Second, the panel is *static*: the same identities discharge the externality obligation across all transitions, contradicting the perpetual-externality reading of G2 we gave in §2. The rent and the staticity are the same problem under two descriptions: a frozen topology where one named node sees every transition.

The mycelial-gate replaces the panel with a predicate. The kernel admits a transition iff:

- at least k attestations have been recorded on the governance-relevant transition;
- the k attestors are spread across at least a threshold of *topologically independent* branches, where independence is defined by computable metadata over the attestor graph (institution, jurisdiction, funder, prior co-attestation neighbourhood);
- no single cluster (institution / jurisdiction / funder) holds more than $\lfloor N/3 \rfloor$ of currently-active attestor capacity;
- the attestor's enrolment timestamp pre-dates the transition, satisfying the witness-freshness predicate of Annexe B.

The gate is a predicate over the topology, not a list of DIDs. No attestor is named in the gate itself. Parameters (k, n) evolve over time; until $n \geq 7$ on a topology-diverse graph, the gate *degrades gracefully* to the static fallback explicitly named as such, monitored by the falsifiers of §7, with a sunset programmed into the same monitor.

The rent, as economic object. Treat the named-veto seat as what it is: a rent that the federation pays for the convenience of a small panel. The diagnostic that follows is unsentimental. A rent exists where a position is scarce, hard to displace, and visible *to its incumbent* but not *to the federation that hosts it*. Each of those properties is a design choice. The mycelial regime attacks all three: scarcity is replaced by an open enrolment with a demonstrated work-product gate; ease of displacement is enforced by mandatory transience (every attestor role expires absent absorption); visibility is enforced by mandatory internal publication of the cost-of-veto per member, recomputed periodically. **Measuring is** compressing: a rent that is reported in clear inside the federation has, by that act, lost its silence and most of its defensibility.

The named guardian is not malign; it is expensive, and the mycelial-gate is the operational shape of refusing to pay that rent.

Two illustrative verticals. The mycelial regime is agnostic over the substrate of the verticals it absorbs. We give two brief illustrations, deliberately stripped to their structure to emphasise that the same shape recurs.

A cybersecurity vertical. A federation engages a senior practitioner from a leading commercial-security actor to attest that a specific deployment of a federation-developed component meets a named threat-model under a named adversary. The attestation is public, the criterion is published in advance, the attestor's identity is recorded but is not promoted. The attestor's signature counts as one mycelial-gate vote on *that* vertical and contributes *nothing* to the federation's posture on any other vertical. After the absorption window closes, the practitioner becomes a federation member if a vetting-gate passes; otherwise the attestation stands as an external-witness record without an associated absorption.

A legal-doctrine vertical. A federation engages a senior partner from a top-tier jurisdiction-specific firm to attest that a particular template (an attestation agreement, a contributor licence, a transparency clause) is enforceable under a named jurisdiction against a named class of counterparties. The attestation is dated; the attestor's competence is on the public

record. The signature counts as one mycelial-gate vote on *that* vertical and contributes nothing to verticals in other jurisdictions or doctrinal domains.

The two verticals are illustrations, not claims about specific counterparties. They are presented here to make explicit that what you are reading on these two examples is the generic form of any later vertical the federation absorbs: the substrate changes, the shape of the gate does not. Cyber-security and legal doctrine are two domains where the cost of a wrong attestation is high and the incentive to misrepresent expertise is well-mapped; that is why they illustrate well. They are not the only domains.

Image: forest, not pyramid. A natural objection, raised against any merit-rotation scheme, is that the mycelial regime, by privileging an early adopter, is a pyramid in disguise. The objection misreads the topology. Consider the two diagrams:

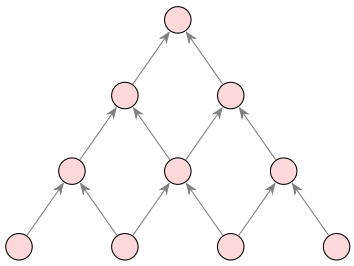

Pyramid: status flows up
early names extract from late ones

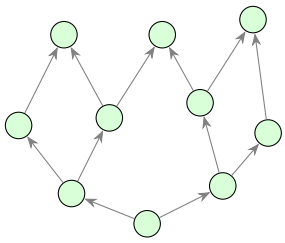

Forest / mycelium: status flows down
late adopters dilute and protect early ones

Figure 5: Two topologies for an attestation graph (*identity-as-credential* → *topology-as-credential*). In the pyramid, the first signatures sit at the apex and accrue authority from those that follow; in the mycelium, the first signatures sit *below* the canopy and are diluted (and so protected) by every subsequent attestor. The federative status flows toward the most recent absorption, not the most senior one.

In the pyramid, status flows upward: the early adopter is extracted from by the late adopter, who validates the apex by submitting to it. In the mycelium, status flows downward: the early adopter is *protected* by the late adopter, whose attestation dilutes the authority of every prior signature without erasing it. The graph has no apex; every node is, in turn, the bottom of the canopy for the node that absorbs after it. A federation's communication discipline follows the same topology: the public-facing language names the most recent cohort rather than the first signature, in keeping with the mycelial reading of legitimacy. The constitution itself is not put to plenary vote; what is voted are the verticals whose adoption the constitution enables. The asymmetry is a derived property of the topology, not a separate rule layered on top.

Image: dissolution in time. The same attestation graph at three moments makes the dilution operational. The node that is central at the earliest snapshot is no longer identifiable at the latest. The canopy is the credential; the first signature is not erased, but the eye cannot recover it without a legend.

Naming reform. Because the mycelial regime carries a different topology than the static regime, it carries different nouns. The vocabulary shift is not cosmetic *where the new noun licenses a reader-inference the old one forbids — vetoer* carries identity and a power to block, *attestor* a moment of externality and an obligation to witness (Annexe D exhibits the inference term by term); where no such inference exists, the shift is relabelling and is named as such. The full table is in Annexe D. The three terms that matter for the body of the essay: **vetoer**

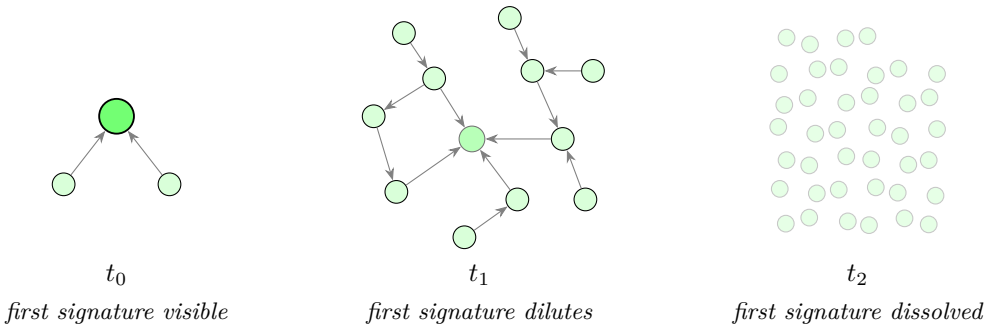

Figure 6: The same attestation graph at three moments ($t_0 \to t_1 \to t_2$). The node that was central at t_0 (highlighted) is structurally less central at t_1 and, at t_2, is no longer identifiable to the eye — one dot among the canopy. The security of the federation is not the memory of who signed first; it is the shape of the graph the eye no longer knows how to undo.

becomes **attestor** (a node that lays down a witness, not one that blocks); **validation** becomes **attestation** (a dated, public signature on a specific vertical); **vetoer-gate** becomes **mycelial-gate** (k of n on independent branches). A reader of earlier versions will encounter the older terms in the lineage record (§E); the mapping survives as a transition aid. The gate seals transitions; what seals the grammar that says what "to seal" means is named at §5.1 and inscribed as future-work.

5.1 Recognising the unsealed grammar

This subsection states what we take to be the central honest finding of the essay — not an embarrassing gap relegated to an annexe, but the load-bearing result on which an attacker and a sympathetic reader should both focus: *every system in the capture record of §3 reduces — but cannot eliminate — the guardian, relocating sovereignty to whoever fixes the grammar of change; whether this generalises to all such systems is the empty-core theorem this essay declines and names as open.*

One observational asymmetry travels with this finding, stated here rather than hedged in an annexe: the capture record of §3 documents the *disease* — the bite of the diagonal on named-guardian systems — but no case yet observes the *cure*, a mycelial-gate running in the field. Every quantitative figure the essay attaches to the cure — the order-of-magnitude gain of Annexe C, the saturation threshold of §8, the falsifiers of §7 — is therefore a *prediction of an instrument not yet built*, not a measurement.

The mycelial-gate seals the federation's transitions: when a new vertical is absorbed, when an attestor's role expires, when an attestation enters the ledger. It does not seal the federation's *grammar*: what a vertical is, what counts as a topologically independent cluster, which work-products satisfy the vetting gate, what re-defines the conditions A1–A12 themselves (of which only A3, A4, A6 are checked invariants; the rest are decreed, Annexe A). That grammar is, at present, the product of authorial decision sealed by external review of the design record, not by a meta-attestation mechanism analogous to the mycelial-gate at the ordinal level above the transition layer.

This asymmetry is not repaired by an orthogonal module composed in parallel with the mycelial-gate. A mechanism capable of sealing the grammar would have to operate one ordinal level above the transition layer; it does not compose, it parameterises. A patch that pretended otherwise would commit a category error, conflating the application of the rules with the rules that admit the application. The honest form of the repair, if it comes, is therefore fusion one ordinal level up in later work — an open problem on no committed schedule — not extension at this one.

This subsection is the recognition itself. The mechanism that *would* close the gap (conjectured as a composed crowd-emergence gate, CEGC, articulated as a three-layer apparatus) is deliberated by the author and named in §8 and Annex A (line A12); the mechanism is referred to in the design record as the composed crowd-emergence gate (CEGC). It is a *named open problem*, carried on no fixed clock — not a deliverable promised for a next release. The gap is acknowledged here as a structural incompleteness of the present essay, not as a failure to be hidden by an additional module composed in parallel with the mycelial-gate.

Read in constitutional-theory terms, this is where the Schmittian residue lives. The mycelial-gate is an explicit move against the single sovereign-guardian who decides the exception: it dissolves that decision into a distributed topology. But because the grammar of change is not itself sealed, a Schmittian is right to answer that the sovereign decision has only been *relocated* — to whoever fixes the grammar and the criteria of independence. We do not dispute the relocation; we assume it as a finding. The contribution is to name the exact ordinal address of the residual sovereign and to make its displacement legible, rather than to claim an elimination the structure cannot deliver. Sovereignty is reduced and pinned, not abolished.

The previous sections established the diagonal pattern and its constitutional consequence. The remainder of the essay shifts register: from diagnosis to discipline. §6 specifies the adoption mechanics, §7 names five falsifiers, and §8 addresses the threshold at which the topological argument requires a public physical anchor.

6 How adoption works

This section is the operational version of the regime, written for a maintainer. It carries no formalism. The five primitives are *vertical*, *attestor*, *attestation*, *absorption*, *adhesion*.

The vertical. A vertical is a self-contained domain in which a specific work-product can be assessed: a kernel security hardening; a documentation overhaul; a translation of a legal template into a new jurisdiction; a numerical-stability audit on a specific module. It is a unit small enough that one externally-attested deliverable resolves it, and large enough that the deliverable constitutes evidence of competence by the producer. The federation maintains an open list of admissible verticals; new verticals enter the list by the same mycelial-gate that admits transitions, so the list is itself diagonal-aware.

The attestor. An attestor is a person external to the federation at the moment of attestation, who has competence in the domain of a specific vertical, and who agrees to record an attestation on that vertical under a public criterion. The attestor is not screened for ideology; the attestor is screened for *having lived* the problem the vertical addresses. A maintainer who has shipped, through an upstream patch, a fix to a class of bug the vertical assesses; a partner who has litigated the kind of contract the vertical assesses; a researcher who has run, against a real adversary, the kind of evaluation the vertical assesses. The attestor's externality is temporal: it is what *at the moment of attestation* means, not who they are forever.

The attestation. An attestation is a dated, public, signed statement that a specific work-product on a specific vertical satisfies a published criterion under a named adversary or counter-party class. It is recorded in an append-only ledger; revocation is impossible by construction. An attestation is not a review and not an endorsement; it is a witness: a statement of record under conditions that make a later contradiction visible. If the attestor turns out, two years later, to have lied or miscalibrated, the attestation does not vanish from the ledger; it remains there with the contradicting record alongside, and the federation's discipline carries the contradiction visibly rather than silently.

The forfeit. The contradicting record, alone, is not a defection cost. An attestor who signs a lucrative false attestation can capture the value now and absorb the annotation two years later as a cost of doing business: benefit-now, footnote-later inverts the incentive gradient of the crypto mechanisms whose shape this regime borrows, where a deposit is *destroyed* at the moment falsity is proven, not merely recorded beside the lie. We therefore require that an attestation bind, *at the moment of signing*, something forfeitable to the published criterion: a pre-committed cost that a proven false attestation destroys rather than annotates. On a stake substrate this is escrowed collateral, slashed on proof. The editorial substrate has no stake to escrow, and we will not pretend otherwise: the honest analogue is a *reputation bond*. The attestor's name is published and bound up front to the specific predicate, adversary class, and counterparty the attestation names, so that a later proof of falsity *against that published criterion* forfeits standing — the right to attest again, and the weight the name carries in every future gate — rather than leaving a name intact beside a footnote. This is reputation-slash, and it is strictly weaker than stake-slash in two ways we state plainly: it bites only where the name carries reusable value its holder is unwilling to abandon, and it can be evaded by an adversary willing to burn the name and re-enter under a fresh one — the Sybil move the threat model below already concedes is undefeated. The bond is therefore real exactly to the degree that named identity is expensive, which is the very property the topological turn spends elsewhere: the regime buys a defection cost here only by re-importing, locally, the accountability of the named seat it relaxes globally. Like every other security property in this essay, the forfeit is offered as a *proposed discipline conditioned on that named-identity cost*, not as a result; strip the cost and the forfeit regresses to the footnote it was meant to replace.

Absorption. After attestation, an absorption window opens (target: 9–18 months). During the window, the attestor either satisfies a vetting-gate (demonstrated work-product on the federation, at least two existing-member witnesses, a public artefact) or does not. If they satisfy it, they are absorbed: they become a member, with all the rights the constitution attaches to membership, and their prior attestation is sealed. If they do not, they remain an external attestor; their attestation stands as an external-witness record. The vetting-gate is the security mechanism, not the absorption itself: a rational adversary *wants* to be absorbed, which is exactly why the gate must be expensive in demonstrated work-product.

Adhesion. A member's adhesion is the act of adopting the constitution as the rules under which they now operate. It is not a vow of loyalty: the constitution is what they have just helped to validate, and their adoption of it is the next step in the recursion of legitimacy through use. There is no oath, no plenary vote, no founding document they sign. The signature on the most recent vertical they attested *is* their adhesion, in the same sense that, in a free-software project, the first non-trivial patch from a new contributor is the implicit acceptance of the project's licence. The constitution governs by being used; it is not put to the membership for ratification because, structurally, the membership *is* the most recent cohort that has used it.

What this looks like to a maintainer who has lived a governance crisis. If you have maintained an upstream package through a code-of-conduct fight, a fork, the resignation of a benevolent dictator, or the capture of a foundation, the mycelial regime is recognisable as a discipline of *never letting status crystallise on a single name*: not a discipline of suppressing authority, but of refusing to let authority compound. The first maintainer of a package has, for a moment, a privileged position because they were first; the discipline of the mycelial regime is to arrange the topology so that this privilege *dilutes predictably* as later contributors arrive, rather than calcifying. A maintainer might expect the fork to be the relief valve: when authority compounds despite the discipline, the dissenters leave and begin again elsewhere. We do not claim that valve here, and the reason is structural. A software fork carries the code — and

often the contributors — to a new home at low cost; a constitutional fork does not transfer. The legitimacy a federation has accumulated lives in its counterparties, its funders, and its legal personality, none of which a dissenting fork can take with it: *you cannot fork the canopy.* The Schmittian residue named in §5.1 therefore has no cheap-exit relief in practice. When the grammar-sovereign re-tunes the criteria of independence to entrench a faction, the move left to dissenters is internal contestation through the falsifiers of §7, not exit; we record this as an open weakness rather than a solved problem.

The discipline does not suppress authority; it refuses to let authority compound on a single name. Leadership remains; what is refused is its conversion into a structural office.

7 Falsifiers and anti-rent guards

A discipline is a security property only if a red-team can refute it. This section names the falsifiers and the rent-watching indicators that turn the mycelial regime from rhetoric into something a competent adversary or a careful auditor can attack.

7.1 The independence premise as an open problem

Before the falsifiers, we name the premise they all stand on, and we name it as unfinished. The entire security argument of the mycelial-gate rests on one assumption: that the k attestations land on *topologically independent* branches — distinct institutions, jurisdictions, funders, and prior-co-attestation neighbourhoods. *Topological independence is the load-bearing security premise of this essay; it is assumed and operationally tested, not guaranteed. Establishing it adversarially is an open problem.* We state this plainly because it is the second predictable line of attack on the essay (the first being the scope of the Gödel analogy, settled in §2), and because the gate's arithmetic is empty if the independence it counts is nominal rather than real.

The threat model has at least four named entries, none of which the present construction defeats; it detects some of them, after the fact, through the falsifiers below.

Sybil. A single adversary stands up several nominally distinct attestors — separate identities, separate declared institutions — that are in fact one principal. The independence metadata is satisfied on paper while the real attestor count is one. The vetting-gate (demonstrated work-product, two existing-member witnesses) raises the cost but does not provably defeat a well-resourced Sybil.

Collusion. Attestors who are genuinely distinct principals, and pass every independence predicate, privately coordinate. No computable metadata over the attestor graph distinguishes coordinated independent parties from uncoordinated ones; collusion is invisible to a topology that can only read declared affiliation.

Shared infrastructure. Attestors independent in institution, jurisdiction, and funder may nonetheless share a common substrate — the same cloud provider, the same certificate authority, the same identity provider, the same key-management vendor. A compromise or a coercion at that shared layer is a single point of failure the declared-affiliation partition does not see.

Correlated failure (common cause). Even absent malice, nominally independent attestors can fail together under a common external shock — a regulatory change across a jurisdiction cluster, a funding shock across a sector, a shared ideological capture. Statistical independence of the failure events is assumed by the k-of-n arithmetic and is not, in general, true.

What the present construction offers against these is partial and honest: the cluster-cap invariant A3 and the *cluster convergence* falsifier below detect *declared* concentration; they do not detect concealed common cause. A real independence metric — one that an adversary cannot satisfy nominally while violating substantively — together with an impossibility-or-tradeoff theorem bounding how much independence a federation of given size can buy, is named here as the central open problem and deferred to the formalisation programme of §9. Until it exists, every quantitative security claim in this essay is conditional on an independence premise that is *modelled and monitored, not proved*.

Anchoring the dimensions externally. The independence predicate, as stated, reads *declared* metadata — the attestor's own statement of its institution, jurisdiction, and funder. That is the wrong side of a line the essay draws elsewhere: it evaluates a self-report, not an external fact. The discipline that narrows the gap is the one the essay already admires in §3 — the OSI move: bind a definition to an authority the defined party does not control. Applied to independence, each declared dimension is anchored, where an external register exists, to that register rather than to the attestor's word: *institution* to a legal-entity identifier in a public registry (and to a beneficial-ownership register where the jurisdiction maintains one); *funder* to public funding disclosures and to the beneficial ownership of the funding vehicle; *jurisdiction* to the public record of incorporation, which is already external. Where such an anchor exists, the predicate resolves the register, not the self-report, and the dimension is *externally verifiable*. Where no anchor exists — informal funding, a private backer, a jurisdiction with no beneficial-ownership register — the dimension is marked **unverifiable**, and an unverifiable dimension counts *against* independence, not for it: the gate records it as an explicit gap, the way Annexe C records a named failure mode, rather than silently reading the absence of contrary evidence as evidence of independence. A dimension is *verified*, *unverifiable*, or *refuted* — never merely *assumed*.

External anchoring narrows the gap between declared and real independence on the dimensions that have a register; it does not close it. Two principals who are genuinely distinct legal entities, each externally anchored, can still coordinate — the *collusion* entry of the threat model above stands untouched. The OSI move buys verifiable distinctness of *identity*; it does not buy verifiable absence of *coordination*. External anchoring is therefore *machine-checkable, not machine-proved*: a machine can resolve a legal-entity identifier and confirm that two attestors are distinct registered entities; no machine can confirm that they are not privately aligned.

The register is itself a guardian. Anchoring independence to an external register narrows one gap and opens another the thesis cannot ignore on its own terms. A legal-entity registry, a beneficial-ownership register, a public-funding disclosure — each is itself a single, state-controlled, coercible authority. To anchor the federation's independence to such a register is to trade the attestor's word for a state's, and the thesis's own logic forbids treating any single authority as unconditionally external: capture the register and you capture every dimension it certifies, re-importing at the anchor the single chokepoint the mycelial regime exists to dissolve. The honest reading is that external anchoring *reduces nominal-identity forgery*; it does not *escape capture*, and that dimension is marked, not assumed away. The mode is observable, not merely conceded: the shared-infrastructure-convergence falsifier of §7 audits the substrate each attestor depends on, and the anchoring registers are one such substrate layer. If the registers on which the k attestors required for a gate resolve their anchored dimensions converge on a single controlling authority — one state's registry resolving a super-threshold fraction of the gate — substrate independence is falsified for that gate exactly as a shared certificate authority would falsify it. The OSI move buys verifiable distinctness against a forger; it buys nothing against the authority that runs the register, and the audit is what keeps that distinction legible from outside.

Reputational hygiene: the order of magnitude. Before listing the falsifiers, we name the reputational-hygiene clause that must travel with any invocation of the regime. The security gain of the mycelial regime over the static k-of-n panel, for plausible parameters ($c_{\text{med}} \approx 10^5$–$10^7$ per attestor, $N \approx 30$ membership, $k = 20$), is approximately **one order of magnitude**, not exponential. The proof-of-work analogy invoked in §4 is *suggestive*, not *constitutive*: anyone claiming exponential security from absorption alone is wrong, and the first competent red-team will demonstrate it. We make the gain visible in the body of the essay because hiding it in an annexe would be a sur-vente that the regime's own discipline forbids: the federation cannot at once require honest cost-of-veto reporting and exempt its own communication from the same norm. The analogy is read at the weight that A3 (cluster-cap) and the cost-of-capture computation of Annexe C actually license. The conditions under which this gain regresses to the static cost are enumerated in Annexe C as three named failure modes (cluster collapse, cohort saturation, cost-of-veto inversion); the first five observables of §7 are the operational form of the same set (the sixth, shared-infrastructure convergence, targets a distinct threat — a shared substrate, not a cost regression).

Six falsifiers. The mycelial regime is falsifiable by six observations on the dynamics of the attestation graph and the topology of the membership.

1. **Absorption decay.** If the rate of absorption dN/dt is monotonically decreasing over six months, the federation has saturated and the regime has degenerated into a static-N rent under a different name. The falsifier is observable on the public ledger; remediation is a forced re-engagement of the enrolment pipeline or an admitted regime change.

2. **Cluster convergence.** If more than 50% of currently active attestors share an institution, jurisdiction, or principal funder over a 12-month window, topological independence is lost and the gate's predicate has become numerical without being structural. The falsifier is observable on the attestor metadata; the response is a moratorium on further absorptions from the dominant cluster until the ratio falls.

3. **Compromised absorption.** If three consecutive absorptions are subsequently shown to have been infiltration attempts (the absorbed party turning out to be capturing the federation under cover of attestation), the vetting-gate has failed. The falsifier is the entryist post-hoc audit; the response is an explicit, dated, on-the-record acknowledgment of the compromise, not the silent quarantine of the compromised absorptions, because the latter would protect the vetting-gate's reputation at the cost of admitting that the regime had become a rent-protection mechanism.

4. **Veto silence under pressure.** If the federation falls silent on a member-vertical that is being attacked externally for more than 90 days, the implicit contract between the absorption ceremony and the federation's defence has been broken. The falsifier is silence in the face of contradiction; the response is the explicit acknowledgment that the contract is contingent on defensibility, not unconditional, and an audit of the conditions under which the silence occurred.

5. **Cost-of-veto inversion.** If the median per-member cost of veto for the most recent cohort of absorptions is *lower* than for the prior cohort, the rent is widening rather than diluting. This is the most subtle falsifier: a federation can keep absorbing new members and still be regressing toward concentration if each new member raises the floor less than the prior. The response is a forced rebalancing of the cost-of-veto distribution before the next absorption is admitted.

6. **Shared-infrastructure convergence.** A periodic audit resolves each active attestor's published attestation endpoint and signing chain to the substrate it actually depends

on — the certificate authority, the cloud-hosting autonomous system (ASN), the identity provider, the key-management vendor, and the external register authority each anchored independence dimension resolves to (§7.1). For each such layer the audit computes the largest fraction of the k attestors required for a gate that depend on a single provider; if that fraction exceeds a pre-registered, per-layer threshold, substrate independence is falsified for that gate even where declared-affiliation independence (A3) holds. The thresholds are set *relative to the population base rate* of each layer: a universally-used TLS issuer or a dominant cloud region is not by itself evidence of capture, so a naive "$> 50\%$ share" rule would fire on almost any federation and carry no information — the vacant-falsifier failure, where a test that always fires refutes nothing. The thresholds are part of the pre-registered falsification budget sealed under the mycelial-gate; the audit is observable by anyone who can resolve the public endpoints, so it does not rest on the attestor's self-report. This is the operational form of the *shared-infrastructure* entry of the threat model above (§7.1), promoted from a named-but-undetected mode to an audited one; it is the one falsifier aimed at the coordinated-substrate cartel that the declared-affiliation partition cannot see. As an addition to a pre-registered list it is admissible under the content-increase rule below: the substrate-convergence counter-examples it covers are disjoint from, and additive to, those of the preceding five, so the empirical content of the discipline strictly increases.

Anti-rent indicators. Four indicators turn the silent threat of recapture into an audited number.

- **Cost-of-veto per member, published monthly inside the federation.** A rent that is reported in clear loses its silence and most of its defensibility. The publication is internal; the metric is comparable across cohorts; the act of measuring is the act of compressing.
- **Attestor half-life, executed as policy.** Every attestor role has a deadline (target: 9–18 months) past which the role expires absent absorption. The expiry is automatic and ceremonial-free. The discipline closes the trap of dormant attestors who keep their external authority indefinitely without the work that justified it.
- **Spacing between absorptions.** A minimum spacing between successive absorptions (target: 4–6 months) prevents the federation from absorbing faster than it can dilute the prior cohort. Without spacing, each absorption raises the cost-of-veto floor of the most recent member without lowering the median; the arithmetic is the same as in proof-of-work block timing.
- **Public communication discipline.** The communication layer names the most recent cohort, never the first signature; the constitution is never submitted to a plenary vote; only the verticals are. This is what §5 called the forest-not-pyramid posture, lifted to a procedural rule.

Regime A vs Regime B. The instrumented falsifiers distinguish two regimes the federation may find itself in at any moment. *Regime A (saturation):* $dN/dt \leq 0$ over six months, cluster ratio rising, cost-of-veto inverting. The regime has become a static-N rent under a different name; the discipline has not held. *Regime B (continuation):* $dN/dt > 0$, cluster ratio bounded, cost-of-veto compressing. The regime is functioning. Both regimes are visible from the same dashboard. A federation in regime A is not disqualified; it is in honest regime A, with a published path back to B or an admitted reorganisation. The discipline is not the absence of regime A; it is the refusal to be in regime A and call it regime B.

Falsification budget. The six falsifiers and the four indicators together are not exhaustive by construction; they are *pre-registered*, with a bounded enumeration and an externally-signed

seal of the list. A failure outside the named falsifiers triggers an amendment cycle that adds a falsifier; the amendment is itself a governance-relevant transition that passes the mycelial-gate. The auxiliary-hypothesis objection (Popper 1959, §19) applies: adding a sixth falsifier after a failure outside the named five is amendment, not falsification. The objection lands. The discipline is falsifiable only conditional on a pre-registered budget (which is exactly what the seal of the falsifier list provides), and the budget itself becomes part of the artefact whose evolution is governed by the same mycelial-gate.

Content-increase rule on amendments. The amendment cycle above is bounded by a rule borrowed from Popper's reading of scientific revision: *any amendment of a named falsifier $F_i \to F_i'$ (or any addition F_{n+1}) must increase, never reduce, the empirical content of the discipline.* Operationally, let $\mathcal{C}(F)$ denote the set of observable counter-examples that would falsify F over the dynamics of the attestation graph. An amendment $F_i \to F_i'$ is admissible iff $\mathcal{C}(F_i) \subseteq \mathcal{C}(F_i')$, with the inclusion strict for any non-trivial amendment. The rule prohibits the Popper-trap (an amendment cycle that adapts the named falsifiers to each new counter-example without constraining the trajectory): a F_i' that excises the inconvenient counter-example without broadening the rest of the falsifier's coverage is, by construction, inadmissible. The seal of the falsifier list under the mycelial-gate records the previous $\mathcal{C}(F_i)$, the proposed $\mathcal{C}(F_i')$, and the union of counter-examples each version covers; the gate verdict is mechanical (set-inclusion check on the recorded enumerations), not editorial. The discipline thus survives its own self-correction: each revision tightens, never loosens, the empirical commitment.

8 Asymptotic regime and the cosmic anchor

Mycelial validation is best read as a specific subclass of *web-of-trust*, the family of trust systems where legitimacy emerges from a graph of cross-attestations rather than from a central authority. What distinguishes the mycelial subclass is a small set of axioms the wider web-of-trust family does not require: each attestation must come from a witness external to the federation *at the moment of attestation* (A1), once landed it cannot be revoked (A4), it requires demonstrated work-product to issue (A5), and the role of attestor is consumed by the act of attesting (A6). Read in shorthand: *external, irreversible, earned, transient.* The remaining invariants (A2, A3, A7–A10) are disciplines that other web-of-trust systems can adopt; A1 + A4 + A5 + A6 are the irreducible additions that name the regime.

The discipline as stated so far says what to do as the federation grows. It does not yet say what happens when growth approaches the diversity ceiling humanity admits at its current scale. The next three paragraphs fill that gap.

(i) The threshold is topological saturation, not planetary coverage. The mycelial regime does not need to reach every human on Earth to hit its operational ceiling. It hits the ceiling earlier, at the point where the diversity-of-clusters predicate (A3) can no longer be satisfied with human actors, regardless of how many of them join. What counts toward that ceiling is active attestors — those who have done the work A5 requires, not users — so the binding limit is the diversity of the substrate, not its size. Three dominant clusters (Anglo-American, Chinese, European) already cover most functional jurisdictions on the planet; once a federation is large enough that its attestor pool is drawn from those three, the constraint *no cluster larger than* $\lfloor N/3 \rfloor$ binds, and adding more people inside the same three clusters does not relieve it. The *meteor* (the point at which a federation must look outside the human substrate for externality) arrives before the planetary utopia does. Where this point falls depends on the cluster-cap ratio rather than on any headcount; it is a topological threshold, calibrated qualitatively against comparable governance bodies (ICANN, IETF), not a population target.

(ii) Externality at the threshold is procured cryptographically. What replaces a person external to the federation is not another person. There are none left. It is a commitment-with-physical-anchor: each attestation commits a hash to the output of a future round of a public physical beacon *before* that round is revealed, and is verified after publication. The mechanism is cryptographic; the source is physical. Two sources sit in the design, by intention independent. The primary source is the *drand* beacon (League of Entropy 2026) operated by the League of Entropy: a multi-organisation distributed beacon producing threshold-signed randomness on a 30-second cadence, where defection by a single organisation does not produce a fabricable output. The independent backup is a pulsar-timing array such as NANOGrav (NANOGrav Collaboration, Agazie, Anumarlapudi, and Archibald 2023): rotation periods of millisecond pulsars are public, measured by multiple radio observatories, and non-fabricable except by altering the angular momentum of a neutron star. The two sources fail under different threat models: drand under coordinated cryptographic compromise of the League's membership, NANOGrav under instrumentation failure of the world's radio observatories. Their conjunction is not redundant; it is decoupled. The literal meteor stays in the text as the image that names the property the anchor must have (non-fabricable, non-negotiable, neutral), not as the mechanism, because the federation cannot wait for a meteor to fall. This saturation threshold is itself a *prediction of an unbuilt instrument*, and as such it is falsifiable: the claim is falsified if A3 saturates substantially below 10^6 active attestors, indicating the topology of qualified professionals is narrower than the topological argument assumes, or if it does not saturate by 10^8, indicating the cluster diversity is wider. Either outcome triggers a revision of the A11 activation predicate.

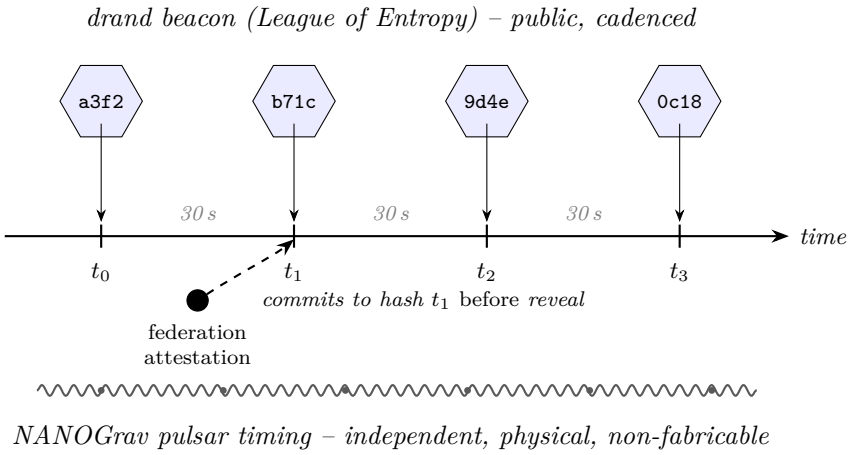

Figure 7: The cosmic anchor does not strike once; it beats. The drand beacon emits a fresh threshold-signed hash every 30 seconds; a federation attestation made between two ticks commits to the *next-tick* hash before that hash is revealed, so an attestor cannot rewrite the past once the tick lands. The NANOGrav pulsar-timing array supplies an independent reference whose non-fabricability rests on the angular momentum of neutron stars, not on the cryptographic discipline of the League of Entropy: the two channels fail under different threat models, and that decoupling – not redundancy – is the property the design seeks.

The clock is physical, and no operator can defer expiration. The cosmic anchor is not only the source of externality at the threshold; it is also the discipline by which the federation treats its own time variable. The kernel's mechanisation reads `now` as a *physical clock* (the drand cadence is the load-bearing reason; the NANOGrav backup is the load-bearing fallback), not as an operator-controlled scheduler. Concretely: *no operator can defer expiration.* The attestor half-life of A6 is enforced by an explicit `Reaper` action under weak fairness: once an attestor's window W has elapsed in physical time, the Reaper fires, the attestor's role transitions through

the *expiring* state into the *expired* state, and no governance-relevant operation can re-open the window from inside the federation. The three-state lifecycle {*active*, *expiring*, *expired*} is the operational shape of A6's *executed transience*: it converts A6 from a property the federation must *honour by format* into a literal state invariant a model checker refutes state-by-state, and it is the explicit consonance the half-life image already carried. The discipline forbids three operator-side evasions: pausing the clock (only the physical cadence pauses; the federation does not author it), retroactively extending a window (the ledger is append-only; A4 forbids it), and back-dating a Reaper firing (the firing is recorded against the cosmic anchor, not against an internal timestamp).

The federation's design record carries this as **A11**, conditional on threshold; its operational form is sketched in Annexe A. A11 is conditional by design. *Pre-activation is refused*: invoking the cosmic anchor before the threshold is crossed inscribes a 30-second commit-and-reveal at every attestation for an externality the federation does not yet need, and that overhead is the ceremony-for-ceremony's-sake the discipline forbids elsewhere (§7, A8: *measuring is compressing*; the converse holds, ceremony without measurement is rent without service). The activation triggers are observable: $N \geq 10^6$ active attestors over a 12-month window, or A3' saturating in three dominant clusters, or an annual editorial audit of functional-outside exhaustion. The federation crosses one before A11 binds.

(iii) What is explicitly refused: the static-world degeneration betrays the perpetual-externality reading. At saturation, a reader may be tempted to read A1 in its weakest form: *if no living external attestor remains, accept that A1 has degraded to a historical statement — externality held when the federation was small, and that founding externality is enough.* That reading rewrites the axiom. A1 says *external at time t*, where t is the time of the assertion, not the time of the federation's founding. The substitution $t \mapsto t_0$ is what §2 called the *single ceremonial bootstrap* — the move the perpetual reading of the diagonal refuses. Freezing the federation in order to preserve topological diversity does not preserve A1; it transforms A1 from a recurring obligation into a one-time event, and the diagonal returns at the next attestation.

The discipline therefore refuses three moves at the threshold: (a) *static-world degeneration* (the federation freezes; A1 silently weakens to its historical sense); (b) *reliance on a literal meteor* (non-cadenced; the federation cannot run on it); (c) *reliance on other federations alone* (G2 transitivity — the question moves outward without being answered: who validates the federation that validates the federation that validates ours?). The refusal is not absolute; falsifiers are listed for each in the federation's design record, and a future ADR may amend them. But the refusal is structural: a federation that takes any of the three is no longer a federation governed by the mycelial axiom. It is a static panel under a different name, the very rent the discipline was written to compress.

Beyond the transition layer: the grammar question. The mycelial-gate seals transitions of state, not the grammar that defines them; the conditions A1–A12 themselves — checked invariants (A3, A4, A6) and decreed conditions (the rest) alike — are inscribed by authorial decree (§5.1). The grammar question is logically prior to the transition question and ordinally above it: a meta-attestation mechanism analogous to the mycelial-gate is required at the ordinal level above the transition layer. The mechanism is referred to in the author's design record as the **composed crowd-emergence gate** (CEGC), articulated as a three-layer composed architecture: (i) an independent-tribe reuse precondition, (ii) an informational emergence metric on multiple decoupled corpora, (iii) a k-of-n meta-attestation discipline applied at the ordinal level above transitions. The full derivation, including thresholds, observation protocol, and explicit falsifiers, is deferred to future work; the present essay does not include it. A12 (Annexe A) inscribes the placeholder, and that is all it does: CEGC is a *named open problem*, carrying no

falsifier and bound to no clock today. The present essay makes *no commitment* to close the gap — not even a commitment to close it later. The reason is the essay's own (§9): a self-certification with no external clock is "philosophically alive but operationally unfalsifiable." A promise to build CEGC "in a later release," carrying neither a deadline a third party could read off the record nor a falsifier, would be exactly such a promise. We decline to make it and name the open problem plainly instead.

What the maintainer should take from this section. What the asymptotic regime gives the maintainer is a procurement discipline that closes against a physical reference rather than against an institutional one: once topological diversity saturates and externality stops being available as a free service of the world, the federation procures it from a public physical beacon whose cadence and content the federation itself cannot author. The cadence is cryptographic and the source is physical, with the drand beacon as primary and a pulsar-timing-array as independent backup that fails under a different threat model. Both properties are visible from outside the federation, and the published verification reduces to checking that the commit-hash inscribed in the seal matches the revealed drand value. The substitution does not turn the federation into a cosmology, and it does not collapse A1 into a historical statement either; it gives the federation an external clock against which its own cadence can be falsified, while keeping the three refusals stated above ((a)–(c)) closed.

9 Conclusion: the bit is identity → topology

The essay began by reading G2 as an engineering constraint; it ends with the institutional shape that constraint imposes. The movement of the argument holds in one substitution: where the customary discipline inscribes a named panel, the discipline we describe inscribes a predicate over the topology of attestations. As successive adopters validate one vertical and are then absorbed, the next vertical is validated by someone still external at the moment of their attestation, and the proper names dissolve into the shape of the graph.

This discipline does not abolish the diagonal limit Gödel proved. It makes that limit legible. When the federation fails — and a federation that absorbs its members between forks will, eventually, fail — the failure occurs at a named point, externally attested, visible while it happens. That is the strongest property a self-referential system can hope to obtain: not perfection, not closure, but the legibility of its own incompleteness. What structurally closes the argument is an artefact — the TLA+ specification of the four-module kernel — whose absence renders that falsifiability conditional.

The residue, stated as a finding. One result of this essay is best read not as a limitation but as its most citable contribution. *Every system in the capture record of §3 reduces — but cannot eliminate — the guardian, relocating sovereignty to whoever fixes the grammar of change; whether this generalises to all such systems is the empty-core theorem this essay declines and names as open.* The mycelial-gate dissolves the named guardian at the transition layer, but the grammar that *defines* a vertical, a topologically independent cluster, the vetting criteria — the invariants A1–A12 themselves — is still authored (§5.1). A Schmittian reading is correct on this point and we grant it: the sovereign who decides the exception has not been abolished, only displaced one ordinal level up, to whoever fixes the criteria of independence and change. The essay's move against Schmitt is not to deny the residue but to *name its exact address* and refuse to hide it. That is the honest form of the reduction: we have lowered the guardian by an order of magnitude and pointed precisely at where the last increment lives, rather than claiming an elimination the structure cannot deliver. This order-of-magnitude figure is an illustrative estimate, calibrated against comparable governance bodies and the cost-of-capture computation

of Annexe C, not a derived worst-case bound.

Honest meta-incompleteness

A12, and the meta-attestation discipline it stands for, opens the same cascade Gödel's theorem opens at the object level: to seal the definition of A12 one would need an A13, and to seal A13 an A14, and so on. A12 is the last storey the present discipline accepts to formalise; beyond it, incompleteness is named as a structural limit, not as a defect to be filled. The same operation we performed at the object level in §7, naming the limit rather than pretending to close it, we perform here at the meta level. This is the supra-mycelial form of the *honest incompleteness* of the object-level discipline, inscribed at the ordinal level above transitions. The discipline by which each of the essay's own self-certifications — including this claim of honest incompleteness — is bound to an external, dated, non-author witness, or named *not-yet-bound*, is stated in §9.

Future work

The argument above is structural; what closes it is an artefact. A four-module kernel specification of the mycelial-gate expressed in TLA+ (Lamport 2002) will accompany this essay at publication, mechanised under TLC, pinned by commit hash, and accompanied by a reproducibility harness that re-runs every bounded instance and asserts the committed expected distinct-state count rather than the mere exit code of the model checker. The kernel is *machine-checkable, not machine-proved*: passing TLC runs establish a real, bounded theorem (*for the specified finite instance, the safety invariants hold on every reachable state of the model*); the kernel demonstrates the design is *internally consistent with* the thesis on bounded models; it does not, and by its own construction cannot, prove that thesis from within (Figure 4). One scope caveat is load-bearing and must travel with every invocation of the phrase: *the TLA+ kernel verifies model-level invariants under stated assumptions; it does not verify that real attestors are independent.* The independence of the branches the gate counts (§7.1) is an input to the model, supplied as an assumption, not a property the model checker establishes; "checked" ranges over the model's reachable states, not over the world the model abstracts. Three modules compose over the attestor graph (`MycelialGate.tla`, `AttestorGraph.tla`, `WitnessFreshness.tla`); a fourth observer module (`GovernanceGate.tla`, a decoupled public stub that preserves the oracle-consultation skeleton of the federation-runtime patch from which it is screened) observes the federation runtime — it abstracts the runtime into a consulted oracle constant rather than rewriting it — and wires the gate to the existing spec-audit channel. The four modules are enumerated, with their substrate and load-bearing invariant, in Annexe F. The kernel commit hash to be pinned at publication will be `04dc5ce`; once the repository is made public, running the harness (`reproduce.sh`) will re-run each bounded instance and verify the distinct-state counts the paper claims. The system that realises this discipline — the four-module kernel together with the federation runtime it observes — is being developed as a commons under the name *Noogram* (`noogram.org`).

Retraction contract. The falsifiability claim of §7 depends on the artefact's continued public availability under a fixed commit hash and a discharged P_{external} obligation (clause $G_{\hat{\mathcal{R}}}$, Annexe F). The formalisation programme that constitutes the four-module kernel is the operational counterpart of the editorial argument; its modules and their load-bearing invariants are listed in Annexe F; the pinned commit hash and the reproducibility harness are the mechanical half of *external*. If the programme were to be *withdrawn* — if the public availability of the kernel sources, the pinned commit hash, the signed falsifier list, or the published identity of the budget-signing authority were to disappear in any future revision — the operational contribution claimed in §7 is retracted, and the legibility-as-falsifiability claim falls back to the philosophical

sense the discipline is otherwise reduced to. The programme is itself one of the named places of incompleteness: a discipline-of-itself must specify where its own falsifiability claim becomes vacuous. The retraction is named explicitly so that the day it would be triggered is the day a reviewer can read it from the same record that carries the claim.

Retraction contract: physical-clock bind. The retraction clause above is contingent on *disappearance*. We now bind it to a physical clock as well, so that *absence-of-progress* is also mechanically detectable. If no signatory of the budget-signing authority is named — in the public record and under a discharged P_{external} slot — within $T = 6$ months of the publication of this essay, the §7 falsifiability claim is marked *Provisional* in a public release note attached to that publication. The bind is to calendar time, not to an internal milestone, so that the trigger is not a verdict the kernel can issue about itself; it is a physical event a third party can read off the same record that carries the claim. The bound $T = 6$ months follows from an adversarial-review observation: a retraction-on-disappearance, without a clock, leaves the falsifiability claim philosophically alive but operationally unfalsifiable as long as nothing is removed. The clock closes that loop.

The external-witness discipline, generalised. The retraction contract above is not a special case; it is one instance of a discipline this essay applies to itself wherever it makes a claim *about itself*. Every such self-certification is bound to a witness that is *external* (the essay cannot author it), *dated* (a third party can read the moment it lands off the same record that carries the claim), and *non-author* (supplied by someone other than the party the claim secures) — or it is marked *not-yet-bound* and named as such, never silently assumed. The governing principle is the one already stated for the clock: a self-certification with no external witness is "philosophically alive but operationally unfalsifiable." Four self-certifications carry the construction; we state the witness each is bound to, and concede the ones not yet bound.

1. **"The branches are independent."** Bound, in part, to external registers the federation does not control — legal-entity identifiers, beneficial-ownership and public-funding disclosures (§7.1) — and to the shared-infrastructure audit any third party can run by resolving the attestors' published endpoints. What remains *not-yet-bound* is named, not assumed away: coordination among principals who are genuinely distinct on every anchored dimension is invisible to any register, and the discipline counts an unverifiable dimension *against* independence rather than for it.

2. **"Capture costs about an order of magnitude more."** The *trend* is bound to the six instrumented falsifiers of §7, each observable by a third party and dated as it fires. The *absolute* figure is *not-yet-bound*: it is an illustrative estimate calibrated against comparable bodies (Annexe C), carries no external witness to its magnitude, and is conditional on independence holding in fact. It is offered as *machine-checkable, not machine-proved* in the same restricted sense as the kernel: the regression of the gain is instrumented and observable, its level is not derived.

3. **"Perpetual external witnessing is necessary."** Bound to the empirical capture record of §3 — the documented one-shot bootstraps that failed (the SPI 2019 self-protection failure, the board-capture mode of foundations, the licence-cascade cycles), external and dated events the essay did not author. This is the witness *external in time*. The stronger witness — *external in logic*, a theorem derived from G2 — is one the essay explicitly does *not* claim (§2): perpetual re-witnessing is a discipline the record grounds, not a necessity a theorem delivers.

4. **"This essay honestly names its own incompleteness."** This is the self-certification the essay's own thesis forbids it to settle from within: a system rich enough to reason

about itself can frame "I am consistent" but cannot honestly settle it from the inside (the abstract states it of the federation; it holds no less of the essay). We therefore bind this claim to no internal verdict and issue no self-certificate of honesty in its place. Its only admissible witness is *external, dated, and non-author*: a dated adversarial-review record, authored by reviewers other than the essay and read off the public record at release. Until such review lands publicly the claim is *not-yet-bound* — the move of the retraction contract applied to the one claim that most needs it, so that the trigger is a fact a third party reads, never a verdict the system issues about itself.

The discipline does not make any of the four claims true, and adds no sixth load-bearing claim of its own; it makes each claim's witness, or the absence of one, legible from outside — the same property the rest of the essay seeks, turned on the essay's own self-certifications.

Open questions. Three remain open at the present time. First, the choice of the budget-signing authority among admissible signatories (publicly-known external observers; treaty-anchored multilateral bodies; civil-society audit consortia with stable identity) is named as a class rather than an institution, and the choice of the actual signatory is one of the items the artefact-ship horizon will resolve. Second, the question of *minimum cohort size* below which the regime degenerates faster than it dilutes is empirical and will be answered by instrumentation, not argument. Third, whether the regime admits a non-state path to comparable external attestation — treaty-anchored multilateral bodies, civil-society audit consortia, cryptographic transparency primitives — is left to the discipline rather than answered here.

A Invariants A1–A10 (and conditional A11, A12)

The ten invariants below are the operational content of the mycelial regime. They are reproduced from the federation's design record under the discipline of source-of-truth lineage stated in §E: the present essay references the invariants by their canonical numbering; the design record is the authoritative source. A conditional eleventh invariant (A11, the cosmic anchor) is added by a later refinement of that record and binds only above the topological-saturation threshold of §8.

Two species hide under the single word *invariant*, and the verifiability column separates them. The *checked invariants* A3, A4, A6 are state predicates a model checker refutes state-by-state. The remaining *decreed conditions* (A1, A5, A7–A10, and the conditional A11, A12) are admissibility conditions of authorial design, sealed by external review of the design record rather than by a TLC run. Both species are falsifiable by an external auditor — that is the exit condition — but by different instruments, and the essay does not conflate the two.

The invariants are redundant between source-personae by design (these source-personae are a structured multi-persona deliberation used for adversarial idea-generation, not attestation — see Declarations); the formulation kept is the most restrictive of each family. All are verifiable: this is the exit condition. A property that cannot be falsified is not a security property, it is a narrative.

Table 1: The ten invariants A1–A10 of the mycelial regime, plus the conditional A11 of the asymptotic regime (§8) and the conditional A12 placeholder for meta-externality. The verifiability column is the exit condition for each invariant: a property that cannot be falsified by an external auditor is not a security property. † A11 binds only above the saturation threshold and is inscribed in doctrine, not in operation, until a trigger fires. †† A12 is a conditional placeholder, dormant in stationary regime; its operative meta-attestation gate is deferred to future work and binds only when a revision to A1–A11 (or to A12 itself) is proposed.

#	Name	Statement	Verifiability criterion
A1	Externality at moment t	Every attestor a of vertical V is external to the federation at the instant of attestation. Formally: $\forall V, \forall a = \text{validator}(V) : a \notin \text{ClosurePriorTo}(S_t, t)$.	Append-only timestamped ledger; $\text{ClosurePriorTo}(S_t, t)$ recomputed at t over events strictly prior to t (Annexe B).
A2	Append-only acyclic graph	The attestation graph is an append-only DAG: no retraction, no cycle through descendance.	Ledger structure; lint at every commit.
A3	Computable topological independence	The mycelial-gate requires k-of-n attestations on independent branches; no cluster (institution, jurisdiction, funder) exceeds $\lfloor N/3 \rfloor$.	Predicate over attestor metadata, recomputed at each event.
A4	Irrevocable, non-retroactive absorption	The seal (a, V) is immutable; subsequent absorption of a does not retroactively re-internalise a's attestation of V.	Append-only seal structure.
A5	Vetting-gate	Absorption requires demonstrated work-product on the vertical, $\geq$ two existing-member witnesses, a public artefact. No social absorption, no purely declarative absorption.	Verifiable presence of the public artefact.
A6	Executed transience	Every attestor role has a half-life; absorption deadline W (target 9–18 months). No dormant attestor remaining external indefinitely. **No named veto > 18 months.**	Timer in the ledger; automatic expiration.
A7	Tempo monitoring + decay alarm	dN/dt published periodically; alarm on monotone decrease over 6 months; the six falsifiers (§7) are instrumented.	Public-internal dashboard; monthly snapshot.
A8	Cost-of-veto per member, measured and published	If unmeasured, the rent returns silently. *Measuring **is** compressing.*	Periodic publication; monthly snapshot.
A9	No validator becomes a celebrity	Public communication names the *most recent cohort*, never the *first signature*. The constitution is never submitted to a plenary vote. Only the verticals are.	Editorial discipline; annual editorial audit.
A10	Public absorption ceremony	The transition attestor $\rightarrow$ member is a dated, visible artefact. Without ceremony the absorption is invisible and the role persists in memory.	Presence of the dated artefact in the ledger.

(continued on next page)

Table 1 (continued)

#	Name	Statement	Verifiability criterion
A11†	Cosmic anchor (conditional)	Above the topological-saturation threshold (trigger: $N \geq 10^6$ active attestors *or* A3 saturating in three dominant clusters *or* annual audit of functional-outside exhaustion), every attestation commits a hash on a future round of a public physical beacon (drand primary; pulsar-timing-array backup) before the round is revealed. Pre-activation is refused.	Two-line check: commit-hash inscribed in the seal; drand round revealed; hash matches.
A12††	Meta-externality (conditional)	The definition and revision of invariants A1–A10 (and of A11, A12 themselves) shall be sealed by a meta-attestation mechanism analogous to the mycelial-gate at the ordinal level above the transition layer. **A12 binds only when a revision to A1–A11 (or to A12 itself) is proposed; below that trigger, it is inscribed in doctrine, not in operation: dormant under stationary regime, woken by any proposed revision.** Does not bind A1–A11 under the stationary regime; the full instantiation of the trigger predicate is deferred to a forthcoming revision.	Predicate: revision proposed to any of A1–A11 or to A12 itself. Below trigger: doctrinal only, no operative gate. Above trigger: invocation of the meta-attestation mechanism (composed architecture deferred to future specification).

Sufficiency claim. A1–A4 suffice for the formal G2 defence: temporal externality, monotone DAG, irrevocability, and lineage diversity neutralise the *who-validates-the-first-validator* critique by making it the case that there is no privileged first validator: every validation is its own external-at-that-moment witness. A5 closes the entryist trap: the absorption itself is not a security mechanism; the vetting-gate is. A6 closes the re-petrification trap: without an executed half-life, the dynamic regime degenerates into a wider but equally opaque static-N. A7–A8 are the anti-rent and anti-silent-capture guards. A9–A10 carry the narrative discipline against cult-of-personality drift.

A11: conditional binding. A11 inherits its substantive content from A1: it is the operational form of *externality at time t* when the human substrate of A3 has saturated, and the witness external to the federation must be procured from a public physical source rather than from the population of qualified humans. Its triggers and refusals are stated in §8. The two-line verifiability of A11 (commit-hash inscribed; drand round revealed; hash matches) is the same shape as the verifiability of A1–A10: a property that cannot be falsified by an external auditor is not a security property. The discipline is consistent across the threshold; only the source of externality changes.

A12: conditional binding. A12 is inscribed as a placeholder, not as an operating mechanism. Its binding is doctrinal until a revision to A1–A11 is proposed; at that point, the meta-attestation discipline named in §5.1 (inscribed in the present essay as a doctrinal placeholder in Annex A, A12; its operational form is the subject of future work) becomes the operative gate. The full instantiation, a three-layer composed architecture (independent-tribe reuse, informational-emergence metric, k-of-n meta-attestation), is a *named open problem*: it carries no falsifier and is bound to no clock today, and the present essay makes no commitment that it will be built. The cascade that A12 itself opens at the meta-meta level is named at §9.

B The Witness Freshness Predicate

Let S_t denote the membership of the federation at time t and ClosurePriorTo(S, t) the transitive closure of attestation neighbourhoods around S admitted by absorption events *strictly prior to* t. Closure here denotes the set of nodes reachable from S along attestation edges of any length, with absorption (A4) treating the absorbed node as henceforth internal, that is, ClosurePriorTo(S, t) is the smallest superset of S closed under absorption from any attestation event whose timestamp is *strictly less than* t. The strict prior-to semantics is doctrinal: an absorption event *at* t does not count for an attestation *at* t, so an attestor cannot be internalised by the very same tick that records her witness. The Witness Freshness Predicate (WFP) is the formal content of A1:

$$\mathrm{WFP}(V, t, a) \iff a = \mathrm{validator}(V) \wedge t = \mathrm{validation\_time}(V) \wedge a \notin \mathrm{ClosurePriorTo}(S_t, t). \quad (1)$$

A validation is admissible iff its WFP holds. The predicate is *renewable*: at each new vertical V' at a later time t', a new WFP must be discharged; an attestor admissible at time t for V may have been absorbed by $t' > t$ and so may no longer satisfy WFP for V'. This is the operational form of the *externality is perpetual* reading of G2 we gave in §2.

Falsifier. WFP is falsified by any observed attestation (a, V, t) for which a already lies in ClosurePriorTo(S_t, t). The append-only ledger of A2 makes this verification mechanical: the closure is recomputed at t from the prior ledger state restricted to events strictly before t, and the membership test is decidable in polynomial time. A failure of WFP at any event is a failure of the gate, not merely a correctness defect of the implementation.

Belt-and-braces on the same-tick race. The rename to `ClosurePriorTo` carries the strict prior-to semantics at the *vocabulary* level. In the mechanised specification it is paired with a same-tick guard on the `Absorb` action — *an attestor cannot absorb at the very tick where she has just attested* — which closes the residual race the vocabulary alone would leave open. The discipline is belt-and-braces by design: the naming makes the strict prior-to semantics legible to a reader; the guard makes the same property a state-by-state invariant a model checker can refute. The two halves co-discharge A1; neither is redundant.

Sufficiency under A1–A4. The *who-validates-the-first-validator?* critique (restated against external-witness schemes by Hobbes, Kelsen and Schmitt) is neutralised by the joint enforcement of A1, A2, A3, A4. Under A1, every validator is external at the moment of attestation; under A2, the graph is monotone, so there is no retroactive re-internalisation; under A4, the seal (a, V) is irrevocable, so an attestor's later absorption does not rewrite their prior witness; under A3, lineage diversity ensures that no single validator (whether *first* or otherwise) carries disproportionate weight in the topological predicate. The conjunction makes the question *who validated the first validator?* mal-posed: there is no first validator privileged by the structure; each validator is its own external-at-that-moment witness, and the gate is satisfied by the topology of the graph rather than by the position of any single attestor in time.

Non-metaphorical analogies. The WFP is structurally adjacent to three constructions in mathematical logic that we cite as non-metaphorical analogies, not as identities. Gentzen's 1936 consistency proof of arithmetic uses ordinal induction up to ε_0, where the ordinal is external to the system whose consistency is at stake. The transposition borrows Gentzen's role-shape — a witness external to the system whose consistency is at stake — and not Gentzen's substance: ε_0 is a metatheoretic well-ordering, the attestor is a person whose externality is empirical and temporal. Cohen's forcing introduces generic objects external to the ground model to extend it without internal contradiction. The transposition borrows the structural move — a witness whose construction is external to the ground model — and not the technical apparatus: the attestor is

not a generic filter, and A3 (cluster-cap on metadata) is not a density condition. The role is shared; the verification is different. Tarski's hierarchy of object-language and meta-language formalises the impossibility of a single language defining its own truth predicate, and resolves it by stratification. The transposition borrows the stratification — adherence claims live in a register one level above the system that makes them — and not the infinite ascent: §9 names the closure at A12 as a structural choice, not as a Tarski-style infinite hierarchy that the discipline pretends to climb. In each case, the closure mechanism is a witness whose existence is established outside the system that requires it; the WFP imports the same structure to the constitutional substrate.

C Cost-of-capture analysis

We model the federation adversarially as a triple $\mathcal{F} = (\mathcal{M}, O_t, S_t)$, where $\mathcal{M}$ is the mycelial-gate predicate, O_t is the observable graph state at time t, and S_t is the membership. An adversary aims to compromise the gate by recruiting (corrupting, coercing, replacing) a subset of the active attestor population sufficient for the gate to admit a transition the federation would otherwise reject. The cost function is

$$C(\text{compromise}) \approx \sum_{i=1}^{k} c_{(i)}, \tag{2}$$

where $c_{(i)}$ is the cost (financial, reputational, coercive) of compromising the i-th cheapest sufficient attestor, summed over the k attestors needed to satisfy the gate.

Regime	**Compromise cost**	**Notes**
Static 2-of-3 named	$\approx 2 \cdot c_{\text{med}}$	Two specific compromises suffice; identities are public.
Mycelial idealised	$\approx \sum_{i=1}^{k} c_{(i)}$	$N = 30$, $k = 20$, topology-diverse; $\sim$1 OOM gain.
Mycelial degenerated	≈ 2–$3 \cdot c_{\text{med}}$	Cluster collapse, $k_{\text{eff}} \to 2$–3.

Table 2: Cost-of-capture order-of-magnitude comparison. The mycelial regime's gain over the static panel is conditional on topology-diversity (A3). Under cluster collapse the regime degenerates toward static cost.

Strict-dominance condition. The mycelial regime strictly dominates the static k-of-n panel only under the conjunction of two conditions: (i) the topology-independence predicate (A3) holds non-trivially (the cluster ratio is bounded below the $\lfloor N/3 \rfloor$ ceiling); (ii) the cohort size N is large enough that the k-th order statistic $c_{(k)}$ is not dominated by c_{med}. For the parameters cited in §7 ($c_{\text{med}} \approx 10^5$–$10^7$, $N \approx 30$, $k = 20$), the gain is approximately one order of magnitude. The asymptote $N \to \infty$ does not deliver an exponential gain; it delivers a polynomial one in N with coefficients set by the cluster distribution.

Regression cases. Three failure modes regress the mycelial regime to static-cost behaviour. *Cluster collapse:* A3 violated; the cluster ratio crosses $\lfloor N/3 \rfloor$ and k_{eff} falls to the static range. *Cohort saturation:* N does not grow above the seven-attestor floor; the gate degenerates gracefully to the static fallback (§5). *Cost-of-veto inversion:* the median per-member cost falls across cohorts; the rent widens under the appearance of dilution. The falsifiers of §7 are instrumented to detect each of these regression cases.

D Naming reform

The mycelial regime carries a different vocabulary than the static regime. The shift is not cosmetic where the new noun licenses an inference the old one forbids (the term-by-term reading

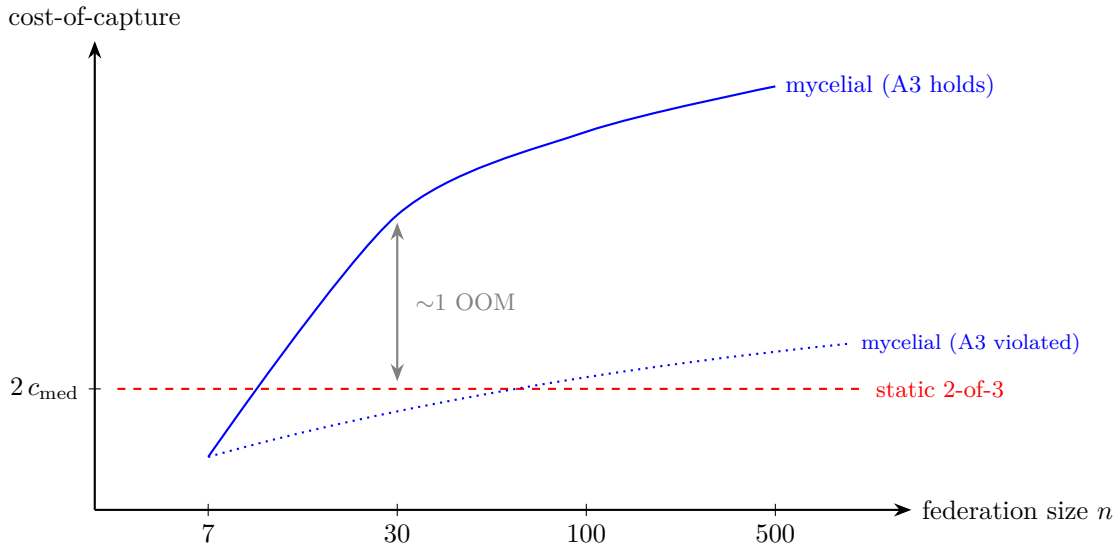

Figure 8: Cost-of-capture as a function of federation size n. The static k-of-n named panel saturates at the cost of compromising k named identities, regardless of n. The mycelial regime, under the topology-diversity predicate (A3), grows at least linearly in n: the gap between the two curves at the operating range cited in §7 ($N \approx 30$, $k = 20$) is the roughly-one-order-of-magnitude gain claimed there. Under cluster collapse (A3 violated), the mycelial curve regresses toward the static range (dotted) and the gain dissolves. The figure is qualitative; the order-of-magnitude reading is licensed by Table 2, not by the curves' exact coordinates.

below); where it does not, it is relabelling, named as such. Readers of earlier versions will encounter the older terms in the lineage record (§E); the table below carries the mapping as a transition aid.

Earlier term	**New term**	**Reading note**
vetoer	**attestor**	A node that lays down a witness, not one that blocks.
validation	**attestation**	A dated, public, signed statement on a specific vertical.
(step 1)	**enrolment**	Initial pairing of attestor and vertical, before any signature.
(step 2)	**absorption**	Attestor → member transition, irreversible (A4).
vertical (operational)	**branch** (topological)	The same object read as topology rather than as practitioner-domain; *vertical* remains admissible in the practitioner register.
vetoer-gate	**mycelial-gate**	k-of-n on independent branches (A3); supersedes 2-of-3 named.
guardian	*(retired)*	The function (preservation of an invariant) is preserved; the support is now the attestation graph itself, not a person.

Table 3: Naming-reform mapping. The mapping does not rewrite the historical record (which remains in the lineage of §E) but specifies the vocabulary of subsequent artefacts.

Why the rename matters. A reader who has internalised *vetoer* as *the named person whose key blocks transitions* will read *attestor* as a synonym and miss the topological displacement. The new term is not a relabel: it names a different object. *Vetoer* carries identity and a power to block; *attestor* carries a moment of externality and an obligation to witness. The rename is the surface on which the underlying structural shift becomes legible to a reader.

E Kernel external-obligation clauses

This annexe records the discipline the four-module kernel of Annexe F imposes on itself once published: the four clauses E.1–E.4 that name where the kernel's external obligation lives, and the pre-registered clause $G_{\mathfrak{K}}$ — a *modular relative-unprovability*, not a Gödel sentence — that fixes the category error the kernel commits to detecting.

The honest claim, in one paragraph. The four-module kernel is *machine-checkable, not machine-proved*. A passing TLC run on a finite instance establishes a real, bounded theorem: *for the specified instance, the safety invariants hold on every reachable state of the model.* It does not close *bounded*→ *unbounded*, *model*→ *system*, nor *spec*→ *thesis*. Nor does it close the gap that matters most for the security claim: *the kernel verifies model-level invariants under stated assumptions; it does not verify that real attestors are independent.* Branch independence (§7.1) enters the model as a hypothesis on the input topology, not as a theorem the checker discharges. The kernel demonstrates the design is *internally consistent with* the thesis on bounded models; it does not, and by its own construction cannot, prove that thesis from within. To phrase the assertion in any stronger form — in particular, to write that *the kernel proves the thesis* — would be self-refutation in the appendix: a model that proved its own constitution would refute the paper that names the limit it inhabits. The retraction contract of §9 is the operational form of the same discipline.

The four clauses of the kernel's incompleteness. The following four clauses name the post-publication discipline the kernel imposes on itself.

E.1 — Residue inscribed, not repaired. The kernel does not close the residue of incompleteness the paper names; it *inscribes* it. Every module of the public repository carries `THEOREM ExternalObligation` stated without proof, marked `OBLIGATION_EXTERNAL -- NOT PROVED WITHIN THIS MODULE` (see Figure 4). The discipline is not to *remove* the marker in a later release but to *preserve* it: a future module that silently discharged `ExternalObligation` from within would not *prove* anything the kernel was withholding: it would commit a category error, importing the unmodelled ledger's witness under an internal name and collapsing the module/ledger boundary the essay says the external witness must cross.

E.2 — Named-and-empty external attestor slot. The kernel's mechanical reproducer (TLC re-running the public greens against the committed expected counts) discharges the *mechanisable* half of *external*: a third party fetches the pinned tooling by hash, runs the harness, and obtains the same distinct-state counts and verdicts. This rerunner is *in-system with respect to the claim*: re-execution against an immutable green-file is a fidelity check, not an external witness. The kernel's $P_{external}$ slot — the public, dated attestation by a non-author practitioner that the four-module specification is faithful to the discipline this essay names — is *declared present and left empty*, in the same source-of-truth ledger that carries the kernel. It is to be discharged by a 2-of-3 vetoer-gate signature on a dated attestation, signed by a practitioner who is external to the federation at the moment of attestation under the discipline of A1. The slot is named so that its emptiness is legible; its filling is a future event, not a provision the paper claims to discharge by being read.

E.3 — Reproducibility = the mechanisable half of external. A *passing CI run* is not a witness in the sense the mycelial-gate requires; it is the mechanisable half of *external*: the half that does not need a person to be performed. The cosmic anchor of §8 plays the same role at the asymptotic regime — a non-fabricable public physical reference, automated but cadenced — and the present clause aligns the public repository with that discipline at

the bounded scale: the harness's distinct-state counts are non-fabricable in the sense that a single edit to a guard predicate moves them, the rerunner detects the move bit-for-bit, and the divergence is a public record. The harness is the mechanical *half* of external; the practitioner who signs E.2 is the *other* half.

For the rerunner's bit-for-bit guarantee to be *non-fabricable* in the operational sense above, the cryptographic signature on each committed `tlc-green` record (the artefacts under `green/*.tlc-green` in the public repository) is required to be signed by a key *strictly distinct* from the author's commit-signing key. A single key would allow the author to retroactively re-sign a green-file whose underlying spec has been edited, which collapses the externality the harness is meant to mechanise. The public CI workflow therefore will demand two signatures on the head of any release tag that updates a `tlc-green`: one author signature (the commit signature, identifying who proposed the run), and one *rerunner-key* signature (identifying that the run was executed under the pinned tooling on a fresh checkout, by an agent whose key is held outside the author's custody). The committed CI configuration will enforce this separation; a release tag that fails the two-key check is refused before the badge updates. The discipline is the operational shadow of A1 (externality-at-t) at the artefact level: the keys that attest the rerun must be external to the keys that attest the spec.

E.4 — Append-only spec-header contract. The kernel's spec-headers carry an append-only contract that forbids any future in-module computation-and-certification of `FederationAdmitted` (the oracle constant the governance gate consults but does not establish). The contract is grep-able, lint-checked, and non-retroactively enforced: a module that later attempted to compute `FederationAdmitted` from inside the same module that wires `GovernanceGateRespected` would violate the contract at the header line, irrespective of whether the inner computation type-checks. The contract is the syntactic guard rail; the spec is the mechanical guard rail; the practitioner's attestation (E.2) is the third, externally-supplied guard rail.

$G_{\mathfrak{K}}$ — the kernel's modular relative-unprovability, pre-registered. Let $\mathfrak{K}$ denote the four-module kernel of Annexe F. Once published, $\mathfrak{K}$ is anchored by its public repository and its pinned commit hash. The pre-registered clause is:

> $G_{\mathfrak{K}}$: *The witness* `MycelialAdmits` *consulted by the gate is itself externally valid — established within* $\mathfrak{K}$. Equivalently: `ExternalObligation` *is discharged inside the module that states it.*

$G_{\mathfrak{K}}$ is *unprovable inside* $\mathfrak{K}$ — not because the machine refuses an answer, but because the existential witness it asserts is supplied by a ledger $\mathfrak{K}$ does not model: the witness is absent from the module's signature, so no in-module proof can quantify over it. $G_{\mathfrak{K}}$ becomes *provable* the moment one passes to a conservative extension that imports that ledger as a modelled constant — which is exactly why the unprovability is *modular* and *relative*, not absolute and Gödelian, and why "the machine refuses to prove" would misdescribe it. To discharge $G_{\mathfrak{K}}$ *inside* the module that states it would not be a proof the kernel was withholding; it would be a category error — importing the external witness under an internal name and collapsing the boundary the essay's discipline exists to keep legible. The pre-registration is the binding move: the clause is named here, in the same record that carries the kernel, *before* any incident asks the discipline to be honest under stress.

The mechanical detector of the module/ledger boundary's integrity is the conjunction of three checks, run as a double-gate by the public CI. *Any future in-module discharge of* `ExternalObligation` *within* $\mathfrak{K}$ is the category error $G_{\mathfrak{K}}$ pre-registers, and is caught if detected by *any* of the three; the strength of the detector is the disjunction. A pre-merge adversarial

review of an earlier formulation observed that a single grep is a weak rail: a refactor can dilute the marker textually while preserving the offending discharge structurally. The double-gate closes that loop.

1. **Lexical** (grep, already present). The token sequence `OBLIGATION_EXTERNAL` disappears from a module of $\mathfrak{K}$, or the line carrying `THEOREM ExternalObligation` acquires a sibling `PROOF` block, or the `ExternalObligation` name appears under a discharged `PROVED` marker in any TLAPS proof object emitted by the build.
2. **Semantic** (TLAPS-emptiness). The set of `ExternalObligation` theorems that have an accepted TLAPS proof object inside `MycelialGate.tla` *or* any of its `EXTENDS`/`INSTANCE` transitive imports (`AttestorGraph.tla`, `WitnessFreshness.tla`, `GovernanceGate.tla`) is non-empty. Formally, writing TLAPS_Closure($\mathfrak{K}$) for the closure of $\mathfrak{K}$ under `EXTENDS`/`INSTANCE`: the falsifier fires when
$$\{T \in \text{TLAPS\_Closure}(\mathfrak{K}) \,:\, T \vdash \texttt{ExternalObligation}\} \neq \emptyset.$$
3. **Cross-module walker**. The public CI runs a static walker over the module graph rooted at any module that `IMPORT`s `MycelialGate.tla`. The walker red-flags any declaration `OBLIGATION_EXTERNAL_DISCHARGED` (as a `CONSTANT`, `VARIABLE`, or `TRUE`-valued operator) anywhere in the reachable graph. A future module that silently supplied a discharged constant from outside $\mathfrak{K}$ but inside its dependency closure is caught by the walker even when the lexical and semantic gates pass.

The CI workflow committed to the source repository will emit the three verdicts, in that order, on every push; a red on *any* one turns the badge red. The three checks are linked to the same commit hash so that a divergence reveals which gate fired. The lexical check is preserved as the cheapest of the three; the semantic and cross-module checks lift the grep beyond a single textual rail.

F Formalisation programme: the four-module kernel

The artefact that §9 commits to is not a single monolithic specification but a cluster of four modules. Three are composing modules over the attestor graph; the fourth is a governance-gate module observing the federation runtime. The boundary between the four is deliberate: each module carries one load-bearing invariant and exposes one composition surface, in the sense that "measuring is compressing" (A8 of Annexe A) forbids hiding plurality behind a singular name. The kernel commit hash to be pinned at publication will be `04dc5ce`; once the repository is made public, running the harness (`reproduce.sh`) will reproduce the distinct-state counts of Table 4.

Public artefact. The four-module kernel is intended to be made publicly available under an archival anchor and pinned by commit hash; the three doctrinal modules and the decoupled governance-gate stub will be mirrored in a public repository, projected by a single one-way screened export from the source-of-truth specifications. The public repository in which the kernel is being developed as a commons is *Noogram* (`noogram.org`). The intended page-level citation is a single anchor (COMMIT HASH), not a labyrinth of internal references. Once published, a reader who follows the anchor reaches the merged source, runs the harness in a single command, and obtains, bit-for-bit, the distinct-state counts the paper claims.

`MycelialGate.tla` *Substrate: attestor graph.* Carries A1 (externality-at-t), A2 (append-only DAG of attestations), A3 (cluster-cap $\lfloor N/3 \rfloor$ topological-independence predicate), and A4

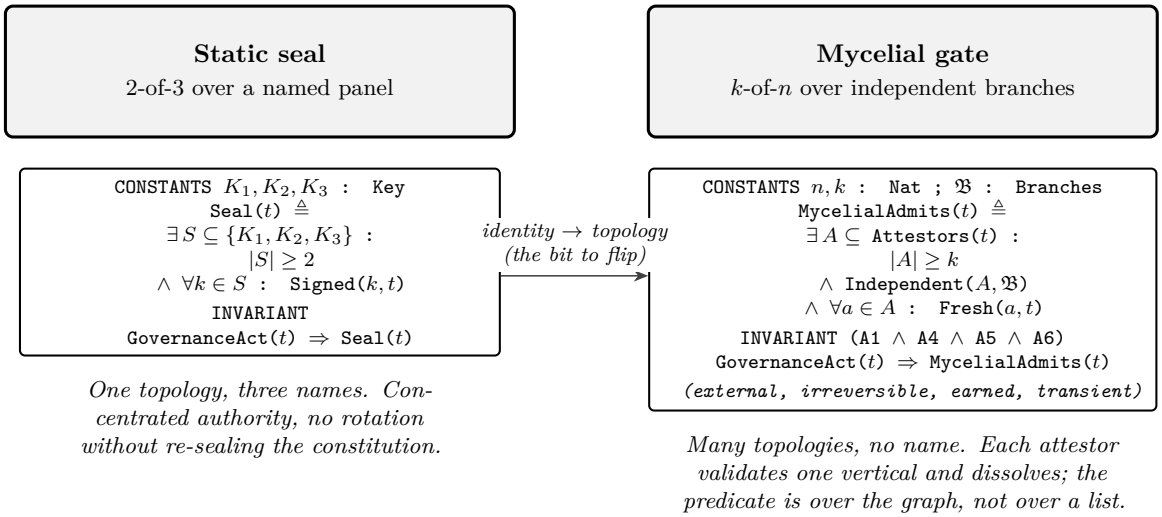

Figure 9: The bit to flip, in TLA+ notation. Left: the static 2-of-3 named-seal articulation as a fixed predicate over three named keys K_1, K_2, K_3, with INVARIANT $\mathit{GovernanceAct} \Rightarrow \mathit{Seal}$. Right: the mycelial-gate as a k-of-n predicate over the attestor graph, parameterised by the independence partition $\mathfrak{B}$ and bound by the four irreducible invariants $A_1 \wedge A_4 \wedge A_5 \wedge A_6$ (*external, irreversible, earned, transient*). The two specifications are not in competition: StaticIsDegenerateLimit inside `MycelialGate.tla` (§F) shows the left panel as the contraction of the right at $N = 3$, $k = 2$ over a fixed key list. The figure is *machine-checkable, not machine-proved*: the predicates are the source of TLA+ checked by TLC on the bounded instances of Table 4; the picture is the reading discipline the reviewer brings. The independence partition $\mathfrak{B}$ is an *input* to the right-hand predicate, not an output TLC certifies: the checker verifies the gate's logic given a partition, not that the partition reflects attestors who are independent in the world (§7.1).

(irrevocable, non-retroactive absorption); proves the gate predicate of §5. Contains the in-module theorem StaticIsDegenerateLimit, which exhibits the static 2-of-3 named-seal articulation as the contraction of the mycelial gate at $N = 3$, $k = 2$ over a fixed key list: the fusion statement, located where its two sides are already defined.

`AttestorGraph.tla` *Substrate: attestor graph.* Carries A2 + A4 + A6 (executed transience, half-life of an attestor's vote with bound $W_{\max}$). The state-machine explicitly promised in this section.

`WitnessFreshness.tla` *Substrate: attestor graph.* Carries the predicate of Annexe B (the attestor's enrolment-time pre-dates the transition it witnesses) plus a guard for the conditional A11 of Annexe A (cosmic-anchor binding only when the asymptotic-regime flag is set). Imported by `MycelialGate.tla`.

`GovernanceGate.tla` *Substrate: federation runtime (abstracted).* Public stub that preserves the oracle-consultation skeleton of the runtime patch from which it is screened. `FederationAdmitted` is supplied as a `CONSTANT` consulted by `GovernanceGateRespected`; the gate verifies the *wiring* from witness to act, never the witness itself. This is the public anchor of the *consults-but-does-not-establish* discipline (Annexe E, clause E.4): the kernel *exhibits* that the gate faithfully consults an oracle it is given; it does not establish the oracle. The runtime substrate is observed, not rewritten.

The four modules compose strictly: the three composing modules import each other through TLA+ `INSTANCE`/`EXTENDS` edges, and the governance-gate stub imports the composition rather than the other way around. The cluster does not introduce a fifth "unifying" module: at the joint of two modules, the predicate that closes them is a theorem inside one of the two, not a third meta-module. The enumeration above is the source-of-truth reference cited by the falsifiability contract of §7 and by the formalisation-programme commitment of §9.

Partial realisation already shipped. The gate module is not greenfield: an earlier, unreconciled prototype of its predicate already exists, written before the four-module decomposition was fixed, and is not yet merged into the kernel pinned above. Recording that antecedent here — rather than presenting the kernel as a first formalisation with no draft behind it — is the traceability the discipline of Annexe E owes itself.

Empirical anchor: state counts on the bounded instances. The kernel is checked by TLC over the three doctrinal modules (`MycelialGate`, `AttestorGraph`, `WitnessFreshness`) at three bounds ($N \in \{3, 5, 7\}$), on a fixed tooling triple (`tla2tools.jar` pinned by hash, JDK pinned). For each cell the harness commits a `tlc-green` record carrying the spec hash, the cfg hash, the tooling version, the distinct-state count, the BFS depth, and the verdict; once the repository is public, the CI will re-run the cell on every push and diff the new record against the committed one. A divergence (*any* guard silently edited, *any* constant tightened) turns the badge red — even if TLC still exits zero on the divergent run. The currently committed counts under the present internal mechanisation (kernel commit `04dc5ce`, reproducible by running `reproduce.sh`) are listed in Table 4. They are distinct-state counts (SYMMETRY quotient via `AttestorPerms`); the published values are quotients, not generated-state totals. *The state count is what is checkable; the badge is theatre.*

Limitations

The argument has bounded scope. We name the principal limits here so they are not left to the reader to reconstruct.

Module	Verdict	$N = 3$	$N = 5$	$N = 7$
`MycelialGate`	PASS	63 931	25 764	7 880
`AttestorGraph`	PASS	6 002	538	538
`WitnessFreshness`	PASS	63 931	68 467	73 003

Table 4: Distinct-state counts and invariant verdicts on the bounded instance matrix at the present internal mechanisation: kernel commit `04dc5ce`, doctrinal amendments (rename to `ClosurePriorTo` with strict prior-to semantics; same-tick Absorb-guard; Reaper action under WF; three-state lifecycle {*active, expiring, expired*}). State counts are SYMMETRY quotients (`AttestorPerms`); once the repository is public, identity of this triple under the pinned tooling will be the empirical anchor the discipline of Annexe E, clause E.3 mechanises.

Scope of the substrates examined. §3 surveys four constitutional traditions chosen for structural diversity within the diagonal pattern, not for geographic coverage. African and Latin-American constitutional traditions beyond Brazil (Kenya's *Ndii*-line jurisprudence, Argentina's *Fayt* line, and Colombia's *Sentencia* C-141/2010) would extend the analysis along the same axis. The omission is named as a research-budget choice, not as a virtue.

Scope of the order-of-magnitude claim. The ∼1 OOM gain of the mycelial regime over the static panel (§7, Annexe C) is conditional on topology-diversity (A3) holding non-trivially. Under cluster collapse the gain regresses to the static range. The gain is *not* exponential; the proof-of-work analogy is *suggestive, not constitutive.*

Scope of the falsifiability contract. The §7 contract is provisional until the artefact described in §9 is published with the signed falsifier list. The formalisation-programme commitment of §9 is the mechanism by which the provisional claim becomes either operational or retracted; until then, the falsifiability claim has the status of a commitment, not of an enforced contract.

Scope of the AI-architecture engagement. The essay positions itself against the constitutional self-reference literature in detail; it does not engage in the same depth with the AI-safety architecture-theory literature (cooperative inverse RL, mesa-optimisation, debate / iterated amplification, Constitutional AI as training method, formal verification for ML). That engagement is deferred to a companion paper in preparation.

Scope of the universality of the discipline. The mycelial regime presupposes a federation that can publish an append-only ledger of attestations, can compute a topological-independence predicate over attestor metadata, and can run a public dashboard instrumenting falsifiers. Federations whose substrate prevents one or more of these (jurisdictions where ledger publication is illegal; domains where attestor metadata is unavoidably private) inherit a discipline that operates in restricted form rather than as a universal recipe.

Declarations

Funding. The author declares that no specific funding was received for the preparation of this manuscript.

Conflicts of interest / Competing interests. The author leads the Noogram commons (noogram.org) referenced in this manuscript, within which the formalisation programme described herein is being developed. No other competing interests are declared.

Data availability. No new datasets were generated or analysed during the preparation of this manuscript. All sources cited are referenced in the bibliography and are publicly available or accessible at the URLs supplied therein.

Ethics approval. Not applicable.

Use of generative AI in the writing process. The manuscript was prepared with the assistance of large language model tools (Anthropic's Claude, Sonnet 4.6 and Opus 4.7 checkpoints) used for literature search, drafting under direction, and surfacing of counter-arguments through structured multi-persona deliberation sessions. The thesis, the architectural proposal, all judgments of relevance, and the final form of every claim are the author's. The assistant did not generate research findings, did not select sources without author review, and is not credited as an author. *The deliberation sessions were instruments for adversarial concept generation and revision planning. They do not constitute independent review, empirical validation, or an attestation under the constitutional mechanism proposed in this essay.* The point is self-referential and load-bearing: an essay that diagnoses self-certification cannot itself treat a model's internal deliberation as a witness to its own thesis without committing the very fault it names. This disclosure is supplied as a separate paragraph from Acknowledgments in accordance with current Springer Nature policy on the use of generative AI in scientific writing.

Acknowledgments

The 1947 Trenton anecdote was first surfaced in conversation; the Guerra-Pujol (Guerra-Pujol 2013) reconstruction was followed from there, and through its footnote 83 the Ross 1969 *Mind* paper (Ross 1969) and Suber's 1990 study (Suber 1990) read in support. The empirical SPI 2019 case was discovered while comparing supermajority-vs-board-trust patterns across foundations. The reframing from static seal to mycelial regime was developed through a five-persona panel deliberation that generated and stress-tested the concept. That deliberation was an instrument of idea-generation and adversarial revision, not an attestation: it underwrites the restructuring as authorial work, and emphatically does *not* constitute independent review or validation of the thesis under the mechanism this essay proposes.

This essay names a discipline; it does not announce an institution.

Back cover

> Every constitution is incomplete. Gödel told us that, and pretending otherwise has cost institutions their legitimacy for a century. *After Gödel* proposes the alternative: a constitution that has no guardians, only adopters, where each new member validates one vertical and then dissolves into the mycelium that protects the next. Read it if you have ever maintained something whose survival depended on a single name, and felt the weight of that name on your shoulders.

References

Abramsky, Samson (2015). "Arrow's Theorem by Arrow Theory". In: *Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*. Ed. by Åsa Hirvonen, Juha Kontinen, Roman Kossak, and Andrés Villaveces. Berlin/Boston: De Gruyter, pp. 15–30. ISBN: 978-1-61451-687-3. DOI: 10.1515/9781614516873.15. URL: https://arxiv.org/abs/1401.4585.

Anthropic (Apr. 2026a). *Claude Mythos Preview: System Card*. The April 2026 model release whose pages 15 and 4.5.5 frame the opening of this essay. URL: https://www-cdn.anthropic.com/08ab9158070959f88f296514c21b7facce6f52bc.pdf (visited on 04/26/2026).

– (Feb. 24, 2026b). *Responsible Scaling Policy v3.0*. Voluntary internal framework for managing catastrophic risks; capability thresholds, Frontier Safety Roadmap, Risk Reports. Minor v3.1 update 2 April 2026. URL: https://www-cdn.anthropic.com/e670587677525f28df69b59e5fb4c22cc5461a17.pdf (visited on 04/26/2026).

Babcock, Charles (Mar. 19, 2007). *The Torvalds Transcript: Why I 'Absolutely Love' GPL Version 2*. Published interview with Linus Torvalds (response received 2007-03-06) in which Torvalds restates his GPLv2-only preference for the kernel, his refusal to commit to an unseen license, and his structural objections to the GPLv3 drafts then circulating. Cited alongside Bottomley et al. 2006 as a documented public articulation of the kernel-side stance. URL: https://www.informationweek.com/software-services/the-torvalds-transcript-why-i-absolutely-love-gpl-version-2 (visited on 05/18/2026).

Banon, Shay (Jan. 14, 2021). *Doubling down on open, Part II — Elastic license change announcement*. Primary source: official Elastic blog post by founder Shay Banon announcing the relicensing of Elasticsearch and Kibana from Apache License 2.0 to a dual Server Side Public License (SSPL) / Elastic License regime, effective with the 7.11 release. Cited as the canonical relicensing event in §4. URL: https://www.elastic.co/blog/licensing-change (visited on 05/18/2026).

Bottomley, James E. J. et al. (Sept. 15, 2006). *Kernel developers' position on GPLv3*. Collective position statement from ten then-active kernel maintainers, published on LWN.net 2006-09-22, raising three objections to GPLv3 Draft 2 (anti-Tivoization clauses, additional restrictions, patent provisions) and declaring no foreseeable draft would prove acceptable. Cited as the public, multi-signatory anchor of the kernel's refusal to migrate from GPLv2 to GPLv3 when the latter was being finalised in 2007. URL: https://lwn.net/Articles/200422/ (visited on 05/18/2026).

Bundesverfassungsgericht (June 30, 2009). *Judgment of 30 June 2009, 2 BvE 2/08 (Lisbon Treaty)*. Acknowledged Article 146 of the Basic Law as path that could in principle circumvent Article 79(3). URL: https://www.bundesverfassungsgericht.de/SharedDocs/Entscheidungen/EN/2009/06/es20090630_2bve000208en.html (visited on 04/26/2026).

Chu, Meng, Xuan Billy Zhang, Kevin Qinghong Lin, et al. (Apr. 2026). *Agentic World Modeling: Foundations, Capabilities, Laws, and Beyond*. Version 1. Early-access preprint, arXiv v1 of 24 Apr 2026; sections 3.2–9 announced (Fig. 1) but not present in the v1 PDF retrieved 2026-04-27. Authors reserve right to update, correct, or retract. Position-driven survey proposing a 3×4 taxonomy: capabilities (L1 Predictor, L2 Simulator, L3 Evolver) crossed with governance regimes (Physical, Digital, Social, Scientific). The L3 Evolver maintains an operator $d_t : \mathcal{M}_t \to \mathcal{M}_{t+1}$ that revises the model stack from deployment evidence under *governed validation gates*. arXiv: 2604.22748 [cs.AI]. URL: https://arxiv.org/abs/2604.22748 (visited on 04/27/2026).

Conseil constitutionnel (Sept. 2, 1992). *Décision n° 92-312 DC du 2 septembre 1992 (Maastricht II)*. Considérant 19: doctrine of pouvoir constituant souverain, explicitly bounded by Articles 7, 16, 89 alinéa 4, and the republican-form clause of alinéa 5. URL: https://www.conseil-constitutionnel.fr/decision/1992/92312DC.htm (visited on 04/26/2026).

Dadgar, Armon (Aug. 10, 2023). *HashiCorp adopts Business Source License*. Primary source: official HashiCorp blog post by co-founder Armon Dadgar announcing the relicensing of Terraform and other HashiCorp products from Mozilla Public License v2.0 to Business Source License v1.1, effective for future releases. Cited as the canonical relicensing event triggering the OpenTofu fork. URL: https://www.hashicorp.com/blog/hashicorp-adopts-business-source-license (visited on 05/18/2026).

Dawson, John W. (1997). *Logical Dilemmas: The Life and Work of Kurt Gödel*. Canonical biography of Kurt Gödel by John W. Dawson Jr., editor of the *Collected Works*. The book reproduces the Trenton 1947 naturalization episode and cites Morgenstern's 1971 IAS memo as the primary record. Wellesley, MA: A K Peters / CRC Press. ISBN: 978-1-56881-256-4.

European Parliament and Council of the European Union (June 13, 2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. Risk-based supranational regulatory framework: high-risk systems require conformity assessment (Art. 43, mostly self-assessment), human oversight (Art. 14), notified bodies optional and provider-selected (Art. 44). URL: https://eur-lex.europa.eu/eli/reg/2024/1689/oj (visited on 04/26/2026).

Federative Republic of Brazil (Oct. 5, 1988). *Constitution of the Federative Republic of Brazil, 1988 (Article 60 §4 — cláusulas pétreas)*. Article 60 §4 entrenches four substantive limits against constitutional amendment: federal pact, periodic direct/secret/universal vote, separation of powers, individual rights and guarantees. URL: https://www.planalto.gov.br/ccivil_03/constituicao/constituicao.htm (visited on 04/26/2026).

Gödel, Kurt (1931). "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I". In: *Monatshefte für Mathematik und Physik* 38, pp. 173–198. DOI: 10.1007/BF01700692.

Guerra-Pujol, F. E. (2013). "Gödel's Loophole". In: *Capital University Law Review* 41.3, pp. 637–673. DOI: 10.2139/ssrn.2010183. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2010183.

International Organization for Standardization (2023). *Information technology — Artificial intelligence — Management system*. Certifiable AI management-system standard organised around named roles, documented risk-treatment, and audit-trail. Operationally paired with EU AI Act high-risk obligations as evidence of conformity for Article 9 (risk management) and Article 17 (quality management system). ISO/IEC. URL: https://www.iso.org/standard/81230.html (visited on 04/28/2026).

Khosla, Madhav (2020). *India's Founding Moment: The Constitution of a Most Surprising Democracy*. Cambridge, MA: Harvard University Press. ISBN: 978-0-674-98087-7.

Krishnaswamy, Sudhir (2010). *Democracy and Constitutionalism in India: A Study of the Basic Structure Doctrine*. New Delhi: Oxford University Press. ISBN: 978-0-19-806536-0.

Lamport, Leslie (2002). *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Boston: Addison-Wesley. ISBN: 978-0321143068.

Landau, David, Rosalind Dixon, and Yaniv Roznai (2019). "From an Unconstitutional Constitutional Amendment to an Unconstitutional Constitution? Lessons from Honduras". In: *Global Constitutionalism* 8.1. SSRN preprint 3208185 (2018). Classifies Honduras 2015 (Constitutional Chamber strike-down of Articles 239 and 374's no-reelection entrenchment) and Nicaragua 2014 as cautionary counter-examples to the Indian/Colombian line — the unconstitutional-constitutional-amendment doctrine deployed against substantive locks rather than to enforce them., pp. 40–70. DOI: 10.1017/S2045381718000151.

Lawvere, F. William (1969). "Diagonal Arguments and Cartesian Closed Categories". In: *Category Theory, Homology Theory and their Applications II*. Vol. 92. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, pp. 134–145. DOI: 10.1007/BFb0080769.

League of Entropy (2026). *drand: Distributed Randomness Beacon*. Multi-organisation threshold-signed public randomness beacon; 30-second cadence by default. Cited as the operational primary source of the commitment-with-physical-anchor mechanism in the conditional invariant A11. URL: https://drand.love/ (visited on 05/07/2026).

Meadows, Carl, Jules Graybill, Kyle Davis, and Mehul Shah (Jan. 21, 2021). *Stepping up for a truly open source Elasticsearch*. Primary source: AWS Open Source Blog post announcing AWS's intent to fork the last Apache-2.0-licensed Elasticsearch and Kibana code (version 7.10) under the ALv2 license; the fork became the OpenSearch project, publicly available in April 2021. Cited as the exit-by-fork response to Banon 2021. URL: https://aws.amazon.com/

blogs/opensource/stepping-up-for-a-truly-open-source-elasticsearch/ (visited on 05/18/2026).

Mendes, Gilmar Ferreira and Paulo Gustavo Gonet Branco (2014). *Curso de Direito Constitucional*. 9th ed. São Paulo: Saraiva. ISBN: 978-85-02-21834-2.

MongoDB, Inc. (Oct. 16, 2018). *MongoDB issues new Server Side Public License for MongoDB Community Server*. Primary source: official MongoDB press release announcing the Server Side Public License (SSPL) as the new license for MongoDB Community Server, with the stated intent of closing what MongoDB characterised as a service-provider loophole in AGPL. Cited as the canonical announcement of the SSPL. URL: https://investors.mongodb.com/news-releases/news-release-details/mongodb-issues-new-server-side-public-license-mongodb-community (visited on 05/18/2026).

Morgenstern, Oskar (Sept. 13, 1971). "History of the Naturalization of Kurt Gödel". Archival memo. Primary source: Morgenstern's 13 September 1971 memorandum, deposited in the Institute for Advanced Study archives (Princeton), recounting the December 1947 Trenton naturalization hearing he attended with Albert Einstein and Kurt Gödel, including the in-car briefing not to raise the constitutional point and the courtroom exchange with Judge Phillip Forman about Austria's transition from republic to dictatorship. Canonical primary record of the "Gödel loophole" anecdote. URL: https://www.ias.edu/sites/default/files/library/Godel/Morgenstern_FriendshipsGoedel.pdf.

NANOGrav Collaboration, Gabriella Agazie, Akash Anumarlapudi, and Anne M. Archibald (2023). "The NANOGrav 15-year Data Set: Evidence for a Gravitational-Wave Background". In: *The Astrophysical Journal Letters* 951.1. Public millisecond-pulsar timing-array dataset. Cited as a physically independent backup to drand for the commitment-with-physical-anchor of the conditional invariant A11; non-fabricability rests on the difficulty of altering neutron-star rotation., p. L8. DOI: 10.3847/2041-8213/acdac6. URL: https://iopscience.iop.org/article/10.3847/2041-8213/acdac6.

National Institute of Standards and Technology (Jan. 2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST.AI.100-1. Voluntary framework organised around four functions: Govern, Map, Measure, Manage. URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf (visited on 04/26/2026).

Open Source Initiative (Jan. 19, 2021). *The SSPL is Not an Open Source License*. Primary source: position statement from the OSI Board of Directors. The SSPL had been submitted to OSI for review in 2018–2019 and was withdrawn by the license steward in March 2019 once it became clear it would not be approved; the 2021 board statement formally restates the OSI position that SSPL does not meet the Open Source Definition. Cited as the external-authority rejection of SSPL as open source. URL: https://opensource.org/blog/the-sspl-is-not-an-open-source-license (visited on 05/18/2026).

OpenAI (Apr. 9, 2018). *OpenAI Charter*. Primary source: founding charter naming broad benefit to humanity as OpenAI's primary fiduciary duty. Cited as the constitutional baseline against which the 2023–2025 restructuring is read. URL: https://openai.com/charter/ (visited on 05/18/2026).

– (Mar. 11, 2019). *OpenAI LP*. Primary source: announcement of the capped-profit subsidiary (OpenAI LP, later OpenAI Global LLC) with a 100× return cap, designed to raise capital while preserving the non-profit board's constitutional control over the mission. Cited as the 2019 baseline structure preceding the 2023–2025 transformation. URL: https://openai.com/index/openai-lp/ (visited on 05/18/2026).

– (Nov. 17, 2023). *OpenAI announces leadership transition*. Primary source: board statement removing CEO Sam Altman. Cited as the moment the non-profit board exercised the constitutional power its 2019 structure had explicitly preserved. URL: https://openai.com/index/openai-announces-leadership-transition/ (visited on 05/18/2026).

OpenAI (Dec. 27, 2024). *Why OpenAI's structure must evolve to advance our mission*. Primary source: OpenAI's announcement of the planned conversion of the for-profit subsidiary to a Delaware Public Benefit Corporation, with the non-profit retaining an equity stake and oversight. Cited as the procedure-validated transformation of the structure that procedure was meant to govern. URL: https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/ (visited on 05/18/2026).
– (Oct. 28, 2025). *Built to benefit everyone*. Primary source: announcement of the completed recapitalisation. The OpenAI Foundation retains a 26 % equity stake in OpenAI Group PBC and the sole power to appoint and remove PBC board members; Microsoft holds approximately 27 %. URL: https://openai.com/index/built-to-benefit-everyone/ (visited on 05/18/2026).
Ross, Alf (1969). "On Self-Reference and a Puzzle in Constitutional Law". In: *Mind* 78.309, pp. 1–24. DOI: 10.1093/mind/LXXVIII.309.1.
Roznai, Yaniv (2017). *Unconstitutional Constitutional Amendments: The Limits of Amendment Powers*. Oxford Constitutional Theory. Oxford: Oxford University Press. ISBN: 978-0-19-876879-1. DOI: 10.1093/oso/9780198768791.001.0001.
Software in the Public Interest (2019). *Questions and answers for 2019 bylaws revision vote*. URL: https://www.spi-inc.org/corporate/votes/2019-bylaws-revision/bylaws-revision-faq/.
– (July 13, 2020). *Software in the Public Interest, Inc. — 2019 Annual Report*. URL: https://www.spi-inc.org/corporate/annual-reports/2019.pdf.
– (2026). *SPI Bylaws (current, with modification history)*. URL: https://www.spi-inc.org/corporate/by-laws/ (visited on 04/25/2026).
Suber, Peter (1990). *The Paradox of Self-Amendment: A Study of Logic, Law, Omnipotence, and Change*. New York: Peter Lang Publishing. ISBN: 978-0-8204-1176-7. URL: https://legacy.earlham.edu/~peters/writing/psa/index.htm.
Supreme Court of India (Apr. 24, 1973). *Kesavananda Bharati v. State of Kerala (AIR 1973 SC 1461; (1973) 4 SCC 225)*. Thirteen-judge bench, 7–6 majority. Established the basic-structure doctrine: Parliament's amending power under Article 368 cannot abrogate the Constitution's basic structure. Reaffirmed in Indira Nehru Gandhi v. Raj Narain (1975) and Minerva Mills v. Union of India (1980). URL: https://main.sci.gov.in/judgments (visited on 04/26/2026).
The Apache Software Foundation (2026). *Bylaws of The Apache Software Foundation*. URL: https://www.apache.org/foundation/bylaws.html (visited on 04/25/2026).
The Linux Foundation (2026). *Bylaws of The Linux Foundation*. URL: https://www.linuxfoundation.org/legal/bylaws (visited on 04/25/2026).
The Linux Kernel Organization (2018). *Linux kernel licensing rules*. Primary source: official kernel documentation declaring that "the Linux Kernel is provided under the terms of the GNU General Public License version 2 only (GPL-2.0), as provided in LICENSES/preferred/GPL-2.0, with an explicit syscall exception described in LICENSES/exceptions/Linux-syscall-note." Cited as the canonical anchor of the kernel's GPLv2-only stance that closed the "or any later version" upgrade path to GPLv3. URL: https://www.kernel.org/doc/html/latest/process/license-rules.html (visited on 05/18/2026).
The OpenTofu Initiative (Sept. 5, 2023). *The OpenTofu fork is now available*. Primary source: announcement of the OpenTofu public repository, forked from the last MPL-2.0-licensed Terraform commit; the project subsequently joined the Linux Foundation as a Series of LF Projects, LLC. Cited as the institutional fork-by-foundation response to Dadgar 2023. URL: https://opentofu.org/blog/the-opentofu-fork-is-now-available/ (visited on 05/18/2026).
Thompson, Ken (1984). "Reflections on Trusting Trust". In: *Communications of the ACM* 27.8, pp. 761–763. DOI: 10.1145/358198.358210.

United States District Court for the Northern District of California (Feb. 29, 2024). *Musk v. Altman et al., No. 4:24-cv-04722 (N.D. Cal.) — ongoing litigation contesting OpenAI's for-profit conversion*. Federal court docket; originally filed in California state court in February 2024 by Elon Musk against Sam Altman, Greg Brockman and OpenAI, removed to federal court; First Amended Complaint (November 2024) added xAI and Zilis as plaintiffs; Second Amended Complaint filed in May 2025; trial in 2026. Cited as the active legal challenge to the validity of the transformation under the founding charter's mission obligations. URL: https://www.courtlistener.com/docket/69013420/musk-v-altman/ (visited on 05/18/2026).

Wikimedia Foundation (2026). *Wikimedia Foundation Bylaws*. URL: https://foundation.wikimedia.org/wiki/Bylaws (visited on 04/25/2026).

Yanofsky, Noson S. (2003). "A Universal Approach to Self-Referential Paradoxes, Incompleteness and Fixed Points". In: *Bulletin of Symbolic Logic* 9.3, pp. 362–386. DOI: 10.2178/bsl/1058448677.