

The Verification Boundary

Institutional independence as a boundary condition on platform expansion in AI markets

ABSTRACT

AI platforms have strong incentives to absorb adjacent workflow, data, and assurance functions. Existing theories explain why integration occurs, but do not isolate a recurring case in which integration can destroy the value of the asset being acquired. This paper proposes the **verification boundary**. When an automated output can impose material, difficult-to-recover loss on a party beyond the producer; when a third party conditions money, permission, coverage, admissibility, authorization, or liability on an assurance about that output; and when conformity can be assessed against a sufficiently determinate standard accepted by that party, accepted verification requires a governance arrangement that preserves institutional independence from the generating function. Independence may be preserved through a protected internal function, a ring-fenced entity, a regulated professional role, or a separate firm: the boundary constrains governance control, not necessarily firm boundaries. Improvements in model accuracy can automate and narrow verification, but do not by themselves make self-attestation acceptable. The paper distinguishes technical checking from governance-based assurance, states six falsifiable predictions with refutation criteria and coding rules, treats conflicted-verifier episodes and cryptographic attestation as live tests rather than confirmations, and proposes an empirical program for identifying where platform envelopment bends rather than passes through. The wider Supply Chain of Intelligence taxonomy is supporting context and is not required to evaluate the claim.

Anand Arivukkarasu · Supply Chain of Intelligence · former product leader, Meta (Instagram) · written in a personal capacity.

Canonical source: supplychainofai.com/papers · linkedin.com/in/anandarivu · Licensed CC-BY 4.0



This paper develops one claim and asks to be judged on it alone.

PROPOSITION

Proposition (the verification boundary). When an automated output can impose material, difficult-to-recover loss on a party beyond the producer; when a third party conditions money, permission, coverage, admissibility, authorization, or liability on an assurance about that output; and when conformity can be assessed against a sufficiently determinate standard accepted by the conditioning party, **accepted verification will require a governance arrangement that preserves institutional independence from the generating function.** Improvements in generation or checking capability may reduce the cost and scope of verification, but do not by themselves eliminate the independence requirement.

1. The problem the proposition addresses

Artificial intelligence expands the range of outputs that can be produced cheaply: diagnoses, valuations, legal filings, financial disclosures, software configurations, engineering plans, and operational decisions. As generation becomes faster and less expensive, firms that control models or user surfaces have strong incentives to integrate forward into workflow, orchestration, and decision support. Theories of appropriability (Teece, 1986), envelopment (Eisenmann, Parker & Van Alstyne, 2011), and transaction-cost economizing (Williamson, 1985) all point in the same direction: absorb the adjacency.

Yet some functions resist absorption even where integration would lower technical and coordination costs. A model provider may inspect its own output more accurately than any external reviewer, and accuracy alone still does not determine whether another party will *accept* its assurance. Courts, regulators, insurers, boards, counterparties, and affected customers may require an assertion whose value depends on the assessor not being governed by the producer whose work is being assessed.

The verification boundary is a claim about that institutional requirement. It is not a claim that generators cannot check their own work, that all verification must occur in separate firms, or that every consequential AI use will spawn a verification industry. It predicts *when* assurance becomes an economically distinct role because independence is part of what the relying party is purchasing. Stated once, plainly, because the phrasing invites the opposite reading: **the boundary constrains governance control, not necessarily firm boundaries.** A protected internal function or a ring-fenced subsidiary can satisfy it.

The gap this paper addresses

Platform envelopment explains why platforms absorb adjacent markets: shared user bases, bundling economics, and cost complementarities make the move attractive and often decisive. Existing theories identify technical, regulatory, transaction-cost, legitimacy, and capability constraints on integration: the adjacency is too distant, capabilities do not transfer, an authority forbids the move, legitimacy is withheld (Suchman, 1995). None of them isolates the condition examined here, in which acquiring control of the assurance function can destroy the third-party credibility that gives that function its value. This paper proposes that condition as an additional stopping rule. Expansion bends where the enveloping firm's own governance destroys the value of the asset it is trying to acquire, because that asset is the credibility of an assertion about the firm's own output.

Definitions

Term	Definition
Generation	Production of a candidate output or decision: a diagnosis, valuation, filing, configuration, recommendation, or design.
Checking	Testing or reviewing an output for defects. May be performed by the generator; improves quality without creating an independent assurance.
Verification	An assertion that an output conforms to a specified standard, addressed to a party entitled to rely on that assertion.
Institutional independence	A governance condition under which the verifier's incentives, authority, evidence access, and accountability are sufficiently separated from the generator to make the assertion credible to the relying party.
Determinate standard	A specification against which conformity can be asserted and assessed, recognized or accepted by the conditioning party. Distinguished from a <i>quality preference</i> : the belief that one output is better than another, with no criterion a third party will accept.
Conditioning third party	A regulator, insurer, auditor, court, board, lender, counterparty, or affected customer that conditions money, permission, admissibility, coverage, or liability on the assurance.

The distinction between *checking* and *verification* carries the argument. Checking is an engineering activity subject to capability improvement. Verification is a governance act subject to acceptance by someone other than the producer. Conflating the two is what makes the boundary look like a temporary artifact of immature models.

2. The mechanism: independence is the asset

The mechanism has four parts. Together they explain why a capability improvement in the generator does not relax the governance constraint.

1. **Consequential reliance creates demand for assurance, not merely for quality.** Where an error can impose material and difficult-to-recover loss on another party, the relevant question is not only whether the output is likely to be correct. It is who is entitled to make an assertion on which others may rely, and who bears accountability when that assertion fails.
2. **Self-attestation is informative but insufficient under defined incentive conditions.** Brands, warranties, repeat business, contractual penalties, and balance-sheet liability can make a producer's own assurance credible when losses are recoverable and punishment is timely (Shapiro, 1983). The boundary arises when those devices do not adequately protect the relying party, or when law, insurance, professional duty, or governance rules require assurance from a separately accountable actor. Recoverability is coded *ex ante* from the measures in Section 7 — reversibility of the harm, litigation and insurance loss experience, and the presence of reliance clauses — so that markets are classified before separation is observed, rather than inferred from it.
3. **Credibility depends on separable incentives and accountability.** Verification has economic value only when the verifier can investigate, refuse certification, disclose adverse findings, and bear consequences without being overruled by the generator (DeAngelo, 1981). The required separation may be functional, legal, professional, or organizational; the stronger the conflict between production incentives and assurance obligations, the stronger the separation required.
4. **Capability and credibility are different variables.** A generator may become the best technical detector of its own errors. That improvement can lower the cost of verification, change the evidence reviewed, and automate much of the process. It does not establish that the generator is an acceptable *sole attestor* where the relying party requires independence. The scarce asset is credible independence, not intelligence.

OBSERVABLE COROLLARY: THE INDEPENDENCE PREMIUM

Where a producer's error can impose material, difficult-to-recover loss, buyers routinely pay for a separately governed source of review, assurance, redundancy, or refusal authority: a second vendor, an accredited assessor, a protected internal function, or an independent testing regime. The expenditure often duplicates technical capability the producer already has. What it buys is independence from the producer's incentives and control. The premium is priced in security, audit, clinical, and financial procurement, and long predates AI. Which form a buyer selects is contingent; that some separately governed source is purchased is the prediction.

The consequence test

Step	Question	Outcome
1. Generation	A candidate output is produced.	Proceeds to the test.
2. Consequence test	Can an error impose material, hard-to-recover loss on a party beyond the producer?	No: integrated assurance. Internal testing, warranty, and brand remain sufficient.
3. Conditioning test	Does an external actor condition money, permission, admissibility, or liability on an assurance about the output?	No: integrated assurance, with reputational discipline only.
4. Standard test	Is there a sufficiently determinate standard, recognized or accepted by the conditioning party, against which conformity can be asserted?	No: no assurance market forms. Demand stays latent until a standard is specified and accepted; a quality preference is not enough.
5. Crossing the boundary	All three scope conditions hold simultaneously.	Yes: independent assurance. A separately governed actor makes the accepted assertion.

Governance separation may occur within one firm, through a ring-fenced unit, or across firms. Firm separation is a contingent organizational response, not the proposition.

Four forms of independence

Governance form	What preserves credibility	Typical setting
Protected internal function	Reporting line to the board or a committee, budget insulation, non-overridable veto, disclosure duty.	Internal audit; model risk management; clinical safety officers.
Ring-fenced unit or subsidiary	Separate legal entity, separate P&L, information barriers, independent leadership.	Acquired assurance vendors retained at arm's length; certification arms.
Regulated professional role	Licensure, personal liability, professional duty that overrides the employer's instruction.	Statutory auditors; licensed engineers; attorneys of record; radiologists.
Separate firm	No ownership or control relationship; reputation staked on refusal rights; portable across clients.	Third-party auditors, testing laboratories, ratings and certification bodies.

The proposition is agnostic across these four. It predicts that *one of them* must be present for the assurance to be accepted, and that all four become harder to sustain as the generator's control over the verifier increases. This is what makes the claim testable rather than tautological: an observed market in which none of the four forms is present, yet the

assurance is durably accepted, refutes it.

3. Relationship to prior work

The claim is intended as an increment, not a replacement. Each literature below supplies part of the argument; the column on the right states what the proposition adds.

Literature	What it establishes	What the proposition adds
Teece (1986); Porter (1985): appropriability and value chains	Complementary assets and chain position determine who captures rent from innovation.	Identifies an asset (credible independence) whose value is destroyed by the act of integrating it, inverting the usual complementary-asset logic.
Eisenmann, Parker & Van Alstyne (2011); Gawer (2014): platform envelopment	Platforms absorb adjacent markets using shared users, bundling, and cost complementarity.	Proposes a stopping rule the theory leaves open: envelopment bends at assurance performed over the enveloper's own outputs, because control is precisely what the buyer is paying to exclude.
Coase (1937); Williamson (1985): transaction costs and firm boundaries	Boundary choice economizes on transaction costs, asset specificity, and hold-up.	Adds a case where integration lowers transaction cost yet is still rejected, because third-party acceptance, not internal efficiency, is the binding constraint.
Baldwin & Clark (2000): modularity	Interfaces and information hiding determine where a system can be cleanly divided.	Suggests a governance interface, not only a technical one: the accepted assertion is the module boundary.
Jacobides, Cennamo & Gawer (2018); Adner (2017): ecosystems	Complementarity, bottlenecks, and multilateral dependence define ecosystem roles.	Identifies a role that must remain separately governed even when the hub can technically perform it better than anyone else.
Akerlof (1970); Spence (1973); Holmstrom (1979): information and incentives	Adverse selection, signalling, moral hazard, and observability.	Links incentive-compatible assurance to a governance prediction under rapidly improving machine capability, where the observability problem is shifting rather than shrinking.
Dulleck & Kerschbamer (2006); Power (1997); DeAngelo (1981): credence goods and verification institutions	Markets develop attestation institutions where buyers cannot evaluate quality directly, and auditor independence has measurable economic value.	Extends the logic to AI market structure and separates technical checking from assurance whose credibility rests on institutional independence.

Adjacent literatures that could already contain the claim

Five further streams sit close enough that a reviewer may reasonably suspect the proposition is already known under another name. They are named here rather than omitted, because the most useful reply to this paper is a citation showing that one of them already states it.

Stream	Closest existing result	Why the claim is not simply contained in it
Independent verification and validation (Lewis, 1992)	Safety-critical software practice separates V&V from development, and the separation survives tooling improvement.	States the practice as an engineering norm within a project. It does not derive when the separation becomes a market boundary, nor tie it to third-party conditioning.
Certification and conformity assessment (Terlaak & King, 2006; Bütthe & Mattli, 2011)	Third-party certification resolves information asymmetry, and private standard-setters supply governance where states do not.	Explains demand for certification given a standard and a certifier. The proposition predicts where the certifier role must be separately governed, and where it will not form at all.
Structural separation in regulated industries (Newbery, 1999)	Unbundling generation from transmission or gatekeeping prevents self-preferencing.	Rests on market power and access discrimination. Here the mechanism is credibility of an assertion, which binds even where the producer has no market power.
Professional independence and its critics (DeAngelo, 1981; Bazerman et al., 1997)	Independence has measurable value, and is corrodible by structural conflict even without intent.	Explains behaviour inside a profession that already exists. Says nothing about which new markets create such a role, or when.
Credible commitment and legitimacy (North & Weingast, 1989; Suchman, 1995)	Actors bind themselves institutionally to make promises believable and gain acceptance.	Supplies the general logic. The proposition specifies the conditions under which the binding must be located outside the producer's governance rather than inside it.

Ratings agencies (White, 2010) are the sharpest case in the other direction and are treated as a hard case in Section 4, not as support.

The nearest neighbour is the audit-independence literature. The difference is directional: that work explains why independence has value inside an established assurance profession, taking the profession's existence as given. The proposition here predicts *where new assurance roles will be created* in markets that do not yet have them, and predicts that the creation is driven by consequence structure rather than by the arrival of regulation.

4. Scope conditions and counterexamples

The proposition applies only when all three conditions hold simultaneously. The scope is deliberately narrow, to make the claim refutable rather than universal.

- **Material external consequence.** An erroneous output can impose material loss, denial of rights, physical harm, legal exposure, or another difficult-to-recover consequence on a party beyond the producer.
- **External conditioning.** A third party conditions money, permission, admissibility, insurance, liability, or authorization on an assurance about the output.
- **Determinate standard.** A sufficiently determinate standard exists, recognized or accepted by the conditioning party, against which conformity can be assessed and asserted. A quality preference does not qualify; the absence of an accepted standard is why the boundary does not bind in most creative and advisory work.

Where one or more conditions are absent, integrated checking, producer warranties, brand reputation, user review, or ordinary contractual remedies may be sufficient. The theory therefore predicts *no* general separation in low-consequence or easily reversible uses, which is where most current AI deployment sits.

Cases the theory must not claim

Case	Conditions	Prediction
Marketing copy generation	Consequence low; no external conditioning.	No separation. Integrated checking and brand suffice. A verification market here would weaken the theory.
Consumer code assistants	Consequence usually recoverable; no conditioning party.	No separation for general use; separation appears only at the regulated or safety-critical subset.
Internal knowledge search	Losses internalized by the producer.	No separation; internal quality management is the efficient response.
Clinical decision support	All three conditions hold.	Separation, via regulated professional role or licensed reviewer, and stable under model improvement.
Financial reporting automation	All three conditions hold.	Separation, via statutory audit and independent model validation.
Consumer credit decisioning	All three hold under fair-lending regimes.	Separation, via independent model validation and adverse-action reviewability.

Genuine hard cases

Four cases sit uncomfortably with the proposition and should be treated as live tests rather than confirmations. First, cloud security certification, where hyperscalers supply both infrastructure and substantial attestation tooling; the boundary holds only because the final assertion is still made by an accredited external assessor, and a shift in that practice would be evidence against the claim. Second, high-frequency trading risk controls, where the speed requirement pushes assurance inside the generating firm and independence survives only as a protected internal function. Third, closed clinical systems in single-payer settings, where the payer, regulator, and provider can be the same institution, and the four forms of independence collapse toward the first. Fourth, and most damaging if left unaddressed, the conflicted verifier: credit ratings under issuer-pays (White, 2010), and the Andersen-era audit failures, show that formal separation can coexist with captured incentives, and that acceptance can persist for years after independence has been hollowed out. The proposition does not claim that separation guarantees good verification. It claims that accepted verification requires a governance arrangement preserving independence, and the conflicted-verifier record refines rather than rescues that claim: where the arrangement is nominal, the observable prediction is a credibility crisis followed by re-separation, not permanent absorption. Sarbanes-Oxley and the post-2008 ratings reforms are the pattern; a case where hollowed-out independence is durably accepted with no re-separation would count against the proposition. Each of the four is a place where the theory could fail informatively.

5. Is this regulation wearing a theoretical costume?

The strongest objection to the proposition is that it restates a triviality: regulated activities attract mandated oversight, so of course assurance is separate. If that were the whole content, the claim would be an observation about administrative law, not a boundary condition on firm strategy. Four lines of argument distinguish the two, and each generates a different observable.

1. **The conditioning party is frequently private.** Insurers, lenders, boards, acquirers, enterprise procurement, and indemnifying counterparties condition money and liability on assurance in the absence of any statute. Cyber-insurance underwriting, lender collateral valuation, and pre-acquisition technical diligence all demand independent assessment because the payer, not the state, refuses to accept the producer's word. If the claim were regulatory, these private demands should be absent or fall away under competitive pressure. They do not.
2. **Separation precedes regulation more often than it follows.** Independent testing laboratories, credit ratings, marine classification societies, and financial audit each emerged as market institutions and were codified afterward. Regulation typically ratifies a separation the market already produced. A purely regulatory account has the causal

arrow backwards, and the ordering is historically checkable.

3. **Regulation itself needs explanation.** Even where a statute mandates independence, the interesting question is why lawmakers converge on *independence* as the remedy rather than on higher accuracy standards, larger penalties, or mandatory insurance. The proposition supplies that answer: accuracy requirements do not address the credibility of the assertion, and penalties do not create a party entitled to refuse.
4. **The prediction is about capability, and regulation says nothing about capability.** A regulatory account predicts nothing about what happens as generators improve. The proposition does: assurance becomes more automated and more selective while independent accountability persists. If assurance requirements systematically relax as accuracy rises, the proposition is refuted and the regulatory account survives untouched.

THE DISCRIMINATING TEST

Take markets satisfying the consequence and conditioning conditions where no statute mandates independent assurance. The regulatory account predicts integration; the proposition predicts separation supplied privately, priced, and stable. Cyber-insurance, lender-side valuation, and acquisition diligence are the cleanest available settings for the comparison, because the mandate is absent and the money is real.

The claim in current AI markets

Two present-day regimes let the proposition be examined without waiting five years. US model risk management guidance (SR 11-7) already requires that model validation be performed by parties independent of model development, with the independence located as a protected internal function rather than a separate firm; it is the cleanest instance of the first governance form applied to automated decision systems. The EU AI Act (Regulation 2024/1689) mixes the forms: high-risk systems face conformity assessment that is, in several categories, self-performed by the provider, while other categories route through notified bodies. That mixture is a natural experiment. The proposition predicts that self-assessed categories will attract buyer-side or insurer-side demand for independent assessment where the consequence and conditioning conditions hold, and that the categories routed through notified bodies will not collapse back into provider self-attestation as models improve.

Two mechanisms are frequently proposed as ways around the boundary, and both are better understood as tests of it. The first is **indemnification**: a generator assumes the downstream liability itself, as several AI vendors now do for copyright exposure. This substitutes a balance sheet for an independent assertion, and the proposition predicts it works where the loss is financial and recoverable, and fails where the conditioning party needs permission, admissibility, or authorization rather than compensation. A court does not accept evidence because the vendor promised to pay if it is wrong. The second is **cryptographic attestation**: zero-knowledge proofs and trusted execution environments

(Ben-Sasson et al., 2014) can make claims about computation verifiable without a trusted verifier. This is the most serious challenge in the paper, and it is conceded in part. Where the assertion reduces to *this computation ran on this input*, cryptography can displace institutional independence entirely. Where the assertion is that the standard was correctly interpreted and applied to a contested situation, it cannot: a proof establishes provenance, not judgment. The proposition therefore narrows to assertions containing irreducible judgment, and widespread acceptance of purely cryptographic assurance in judgment-laden domains would refute it.

A second objection: is separation just transitional?

The alternative reading is that verification is separate today because models are unreliable, and that the boundary dissolves once reliability is high enough. This is the objection the proposition is designed to meet head-on. The separation is not predicated on error rates. Statutory audit did not arise because bookkeepers were bad at arithmetic, and it has not been dissolved by double-entry software, ERP systems, or continuous controls monitoring, all of which raised accuracy substantially. What each improvement changed was the cost and method of assurance, not the identity of the party entitled to assert. The proposition predicts the same pattern for AI, and P2 below states the condition under which that prediction fails.

6. Falsifiable predictions

Each prediction states an observable pattern and a refutation condition, over a five-year horizon from publication. P1 to P3 test the core mechanism; P4 to P6 test the claim against its two closest rivals, the regulatory account and the transitional account.

	Expected observation	Refuted if
P1. Governance separation	In markets satisfying all three scope conditions, accepted assurance is supplied through a protected function, ring-fenced entity, regulated professional role, or separate firm.	Conditioning third parties routinely accept assurance from the generator with no meaningful governance separation, across multiple renewal cycles.
P2. Capability does not erase independence	As generator accuracy improves, verification becomes more automated and more selective, while the requirement for independent accountability persists.	Higher generator accuracy systematically causes regulators, insurers, auditors, or counterparties to eliminate independent assurance requirements.

	Expected observation	Refuted if
P3. Platform expansion bends rather than passes through	Generation platforms expand into workflow, orchestration, data access, and user surfaces more readily than into assurance evaluating their own outputs.	A generation platform repeatedly absorbs such assurance, controls its governance, and retains third-party acceptance across multiple audit or renewal cycles.
P4. Private conditioning is sufficient	Separation appears in unregulated markets wherever a private payer conditions money or liability on the assurance. Operationally: within a matched sample of consequential markets, the share of unregulated markets exhibiting any of the four governance forms is at least half the share observed in regulated markets.	Separation is observed only where a statute or regulator mandates it, and unregulated consequential markets integrate assurance without buyer resistance.
P5. Acquisition without ring-fencing degrades acceptance	Where a generator acquires an assurance provider, either governance is ring-fenced or third-party acceptance of that provider's assertions declines.	Acquired assurance providers are fully integrated into generator governance and retain acceptance and pricing power over multiple cycles.
P6. New mandates produce roles, not absorption	Where an assurance requirement is newly imposed, an independent role, accreditation, or vendor emerges rather than the function being absorbed by incumbent generators. Operationally: within 24 months of the requirement taking effect, at least one accredited or separately governed provider is accepted by the conditioning party in a majority of observed jurisdictions.	Newly regulated domains routinely resolve into single-vendor generation-plus-verification with no structural separation.

What would refute the proposition as a whole

A single well-documented market in which all three scope conditions hold and the generator's own attestation is durably accepted by the conditioning third party, across more than two renewal or audit cycles, would force an amendment. Capability arguments alone would not: the proposition explicitly predicts that generators become more accurate and that the separation persists anyway. Nor would isolated instances of self-certification in low-consequence settings, which fall outside scope by construction.

To keep that standard from absorbing every counterexample, the coding rules are fixed in advance. **Durably accepted** means the conditioning party continues to condition its decision on the generator's own attestation across more than two renewal, audit, or certification cycles, with no accredited assessor, protected function, regulated role, or

independent vendor in the chain. **Scope conditions hold** is coded from the Section 7 measures, independently of whether separation is observed. **No governance separation** means the attesting unit has no protected reporting line, no refusal right the generator cannot override, and no external accreditation. A case coded on these three rules refutes the proposition regardless of how well it is explained afterwards; a reviewer who supplies one has settled the question.

7. An empirical program

The claim is stated to be tested by others. The following operationalizations are offered so that disagreement can be resolved with data rather than with intuition.

Construct	Candidate measures
Consequence severity	Regulatory classification; insurance loss data; litigation exposure; reversibility of harm; presence of statutory duty.
Conditioning intensity	Presence of a reliance clause; audit or certification requirement; insurance condition precedent; procurement mandate; admissibility rule.
Governance independence	Reporting line; veto and refusal rights; budget control; ownership; accreditation status; personal liability; rotation and tenure rules.
Economic outcomes	Contract duration; switching cost; assurance spend; pricing and gross margin; churn; claims experience; regulatory acceptance rate.
Capability change	Generator accuracy and error rates; model upgrade cadence; degree of review automation; change in the <i>amount</i> rather than the <i>independence</i> of assurance.

Methods

- Comparative case studies tracing how assurance governance changes before and after AI adoption in medicine, finance, legal services, cybersecurity, and industrial engineering.
- Event studies around new assurance mandates, testing whether they produce protected roles, ring-fenced units, accreditation regimes, or new vendors, and on what lag.
- Procurement and contracting analysis, coding requests for proposal, insurance clauses, audit requirements, and reliance provisions for explicit independence requirements. This is the most direct test of P4, because contracts reveal private demand without a mandate.
- Panel analysis relating changes in generator capability to assurance spending and governance structure, controlling for output volume and regulatory intensity. This is the direct test of P2.
- Vignette experiments with regulators, boards, insurers, and enterprise buyers, testing whether identical technical assurance is valued differently depending on who controls the assessor. This isolates credibility from capability.

8. Implications for AI market structure

First, platform expansion may face a credibility boundary rather than a technical one. Model providers can improve assurance, automate it, and even supply the tools used to perform it, while remaining unable to serve as the sole accepted verifier of their own outputs in some markets. The constraint binds on governance, not on skill, which is why it does not appear in capability-based accounts of platform limits.

Second, verification is not necessarily a permanent labor-intensive layer. Automation can compress its cost and change its operating model substantially. The durable element, where the proposition holds, is the governance right to make an accepted assertion independently, and that right is not a function of headcount.

Third, the most defensible firm may not be the one with the strongest model. In consequential markets, durable value may accrue to organizations holding accreditation, liability capacity, evidentiary access, procedural rights, or trusted standing with the conditioning third party. These are slow assets. They are also, unlike model quality, not obviously subject to rapid commoditization.

Fourth, the boundary is a strategic instrument, not only a description. A firm choosing where to compete can ask whether the position it is entering is one where its own governance would destroy the asset. If so, partnering, ring-fencing, or standing aside dominates acquisition.

9. Where the boundary sits in the wider framework

The claim originated inside the Supply Chain of Intelligence, a descriptive framework partitioning AI markets into ten layers from physical resources to compounding memory, which describes where value, control, and defensibility accumulate. Within that taxonomy the verification boundary grounds Layer 3, **Gatekeeping**: the layer that governs whether a system is allowed in. **Evaluating the proposition does not require accepting the taxonomy.** The mapping is provided so that the empirical work can be located precisely, and because the taxonomy predicts that the boundary is unevenly distributed across Layer 3 rather than uniform.

L3

Gatekeeping

Trust, acceptance, approval. Can the system be allowed in? The slowest position to build and the hardest to replicate.

Layer 3 divides into five sublayers. The proposition does not apply equally to all of them, and that unevenness is itself testable: the sublayers where independence is constitutive should show separation, and the sublayers where it is not should show integration even under high consequence.

Sublayer	Function	Does the boundary bind?
L3a. Compliance and export controls	Regulatory, legal, and policy filters (HIPAA, GDPR, SOC 2, the EU AI Act), plus export controls, model sovereignty, and data-residency regimes.	Yes, strongly. The conditioning party is explicit and the assertion is addressed to it. Independence is usually mandated or accredited.
L3b. Quality gates	Accuracy, hallucination detection, output grading, eval harnesses, regression suites.	Partly. Tooling integrates readily; the <i>assertion</i> derived from it separates only when a conditioning party relies on it. This is the sublayer most likely to be absorbed, and therefore the best test of P3.
L3c. Safety, security, and provenance	Harmful-content filtering, adversarial defense, prompt-injection protection, and content provenance (C2PA, watermarking, deepfake attestation).	Yes for attestation, no for controls. Provenance claims are assertions to third parties; filters are internal controls. The split within one sublayer is a sharp prediction.
L3d. Editorial gates	Tone, brand voice, style, taste, human judgment over what is published.	No, ordinarily. Consequences are usually reputational and internalized by the producer, so integrated judgment is efficient.

Sublayer	Function	Does the boundary bind?
L3e. Distribution gates	App store approval, ranking, marketplace curation, discovery control.	Different mechanism. Separation here is produced by platform market power, not by credibility demand, and should not be counted as confirming evidence.

The distinction between L3b and L3c matters for the empirical program. A naive test that codes all of Layer 3 as boundary-bearing would find mixed results and conclude the proposition is weak. The proposition predicts the mixture, and predicts its shape: absorption of eval tooling, persistence of accredited attestation, and no separation at all in editorial judgment. The remaining layers, the four structural laws, and the case corpus are documented in the companion working paper and practitioner guide and are not load-bearing here.

10. Limitations

The proposition is stated for a five-year horizon and rests on institutional arrangements that could change. Three limitations deserve explicit statement. First, the four forms of independence are not equally observable; protected internal functions are difficult to code from public data, which biases empirical work toward finding separate firms and may overstate the organizational implication. Second, the theory does not predict which of the four forms will be selected in a given market, only that one will be present; a fuller theory would supply that selection mechanism. Third, cross-jurisdictional variation may confound tests, since the same activity can face mandated independence in one regime and none in another. That variation is also an opportunity, since it provides a natural experiment for P4.

11. The question asked of reviewers

This paper asks one question, and it is prior to any question of merit: **is the verification boundary already established in the literature under another name?** Section 3 names the five adjacent streams where it is most likely to be hiding. If the result exists, a citation is the most valuable possible reply, and the paper should be withdrawn to a summary of prior work. If it does not, the second most valuable reply is a counter-example: a market satisfying all three scope conditions in which the generator's own attestation is durably accepted under the coding rules in Section 6. Comments, citations, and refutations are welcome at [linkedin.com/in/anandarivu](https://www.linkedin.com/in/anandarivu).

References

- Adner, R. (2017). Ecosystem as structure: an actionable construct for strategy. *Journal of Management*, 43(1), 39–58.
- Akerlof, G. A. (1970). The market for “lemons”: quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- Baldwin, C. Y., & Clark, K. B. (2000). *Design Rules: The Power of Modularity*. MIT Press.
- Bazerman, M. H., Morgan, K. P., & Loewenstein, G. F. (1997). The impossibility of auditor independence. *Sloan Management Review*, 38(4), 89–94.
- Ben-Sasson, E., Chiesa, A., Tromer, E., & Virza, M. (2014). Succinct non-interactive zero knowledge for a von Neumann architecture. *USENIX Security Symposium*, 781–796.
- Board of Governors of the Federal Reserve System & OCC (2011). *Supervisory Guidance on Model Risk Management* (SR 11-7 / OCC 2011-12).
- Büthe, T., & Mattli, W. (2011). *The New Global Rulers: The Privatization of Regulation in the World Economy*. Princeton University Press.
- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- DeAngelo, L. E. (1981). Auditor independence, “low balling,” and disclosure regulation. *Journal of Accounting and Economics*, 3(2), 113–127.
- Dulleck, U., & Kerschbamer, R. (2006). On doctors, mechanics and computer specialists: the economics of credence goods. *Journal of Economic Literature*, 44(1), 5–42.
- Eisenmann, T., Parker, G., & Van Alstyne, M. (2011). Platform envelopment. *Strategic Management Journal*, 32(12), 1270–1285.
- European Union (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (the AI Act). *Official Journal of the European Union*.
- Gawer, A. (2014). Bridging differing perspectives on technological platforms. *Research Policy*, 43(7), 1239–1249.
- Holmstrom, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10(1), 74–91.
- Jacobides, M. G., Cennamo, C., & Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, 39(8), 2255–2276.
- Lewis, R. O. (1992). *Independent Verification and Validation: A Life Cycle Engineering Process for Quality Software*. Wiley.
- North, D. C., & Weingast, B. R. (1989). Constitutions and commitment. *Journal of Economic History*, 49(4), 803–832.
- Porter, M. E. (1985). *Competitive Advantage*. Free Press.
- Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford University Press.
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029.
- Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *Quarterly Journal of Economics*, 98(4), 659–679.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.

- Suchman, M. C. (1995). Managing legitimacy: strategic and institutional approaches. *Academy of Management Review*, 20(3), 571-610.
- Terlaak, A., & King, A. A. (2006). The effect of certification with the ISO 9000 quality management standard. *Journal of Economic Behavior & Organization*, 60(4), 579-602.
- Teece, D. J. (1986). Profiting from technological innovation. *Research Policy*, 15(6), 285-305.
- White, L. J. (2010). Markets: the credit rating agencies. *Journal of Economic Perspectives*, 24(2), 211-226.
- Williamson, O. E. (1985). *The Economic Institutions of Capitalism*. Free Press.

How to cite

Arivukkarasu, A. (2026). *The verification boundary: institutional independence as a boundary condition on platform expansion in AI markets* (Academic Theory Paper, Version 4.1). Supply Chain of Intelligence. <https://supplychainofai.com/papers>

BibTeX: @techreport{arivukkarasu2026vb, author = {Arivukkarasu, Anand}, title = {The Verification Boundary: Institutional Independence as a Boundary Condition on Platform Expansion in AI Markets}, year = {2026}, number = {Academic Theory Paper v4.1}, institution = {Supply Chain of Intelligence}, url = {https://supplychainofai.com/papers}}

Licensed CC-BY 4.0. Supply Chain of Intelligence™ and The Intelligence Cube™ are trademarks of Anand Arivukkarasu. Earlier drafts used the term “trust boundary”; “verification boundary” is the settled term. Comments and refutations are welcome at [linkedin.com/in/anandarivu](https://www.linkedin.com/in/anandarivu).