

The Contamination Dividend: How Synthetic Media Turns Verification into an Enclosable Resource

Abstract

Verifiable human-origin data has historically behaved as a public good: non-rival, non-excludable, and free at the point of use. This paper argues that the saturation of the open information environment with synthetic content is quietly converting access to verification into a club good, and names the perverse incentive this conversion creates. The act of verifying remains non-rival; what becomes scarce and excludable is the institutional substrate that makes verification cheap and authoritative. As synthetic contamination raises the cost of establishing that any given text, image, or dataset is of human origin and known provenance, actors who hold enclosed reserves of verified data — pre-2022 corpora, authenticated archives, exclusive licensing pipelines — receive what I call a contamination dividend: an appreciation in the value of their epistemic holdings caused by the degradation of the commons everyone else must use. No actor needs to intend contamination for the dividend to operate; it is a structural incentive gradient, analogous to the historical enclosure of agricultural commons and to markets that monetize pollution-driven scarcity. The paper traces the mechanism, confronts the strongest objection to it — that clean data may already have ceased to exist, leaving the reserve an uncertifiable fiction — distinguishes it from adjacent diagnoses (model collapse, misinformation, the liar's dividend, verification fatigue), and draws out two consequences: the emergence of provenance reserves as strategic state and corporate assets, and a stratification of societies into verification-rich and verification-poor. It closes with governance handles: reserve transparency, public provenance infrastructure, and anti-enclosure obligations.

Keywords: data provenance; synthetic media; model collapse; commons; enclosure; epistemic inequality; AI governance

1. Introduction

Three developments of the mid-2020s are usually discussed separately. First, the technical finding that generative models degrade when trained recursively on their own outputs — "model collapse" — established that human-origin data is not an infinitely renewable input but something that can be diluted past usefulness (Shumailov et al. 2024). Second, a wave of bilateral licensing agreements between AI developers and holders of large text archives — social platforms, news organizations, academic publishers; the Reddit–Google agreement (CBS News 2024) and the string of AI-training deals struck by major news publishers from 2023 onward, such as the News Corp–OpenAI licence (Axios 2024), are the visible edge — began, in effect, to fence off the highest-provenance portions of the written record for exclusive or semi-exclusive machine use, a rapid enclosure of the open data commons documented across the web by Longpre et al. (2024). Third, governments launched

"sovereign AI" programs whose quiet corollary is sovereign data: national corpora, national archives, nationally controlled training material.

Each development has its own literature. What has not been named, to my knowledge, is the economic structure that connects them. This paper argues that they are three surface expressions of a single underlying shift: verification — the practical capacity to establish that a piece of information is of human origin, of known provenance, and uncorrupted — is ceasing to be a public good and becoming an enclosable resource. And enclosable resources get enclosed. The enclosure of knowledge is itself old — guilds fenced their crafts, professions their tacit expertise, and a clerical class its monopoly over the reading of scripture long before land or ideas were formally titled — so the claim here is not that enclosure is new. It is that the object of enclosure has advanced, phase by phase, toward the more fundamental, and has now reached the ground of trust beneath information itself. The contamination dividend is the yield on that newest and deepest phase.

The shift matters because it carries a perverse incentive that, once seen, is difficult to unsee. If you hold a verified reserve — a clean pre-contamination corpus, an authenticated archive, an exclusive pipeline of provenance-guaranteed human data — then every increment of synthetic contamination in the open commons makes your holding more valuable. You do not need to add a single synthetic document to the commons yourself. You merely need to do nothing while the commons degrades, and your asset appreciates. I call this appreciation the contamination dividend, and the central claim of this paper is that the dividend constitutes a structural disincentive, distributed across precisely the actors most capable of protecting the epistemic commons, against protecting it.

The argument proceeds in four sections. Section 2 establishes the economic reframing: why verifiability, rather than information itself, is the resource whose good-type is changing, and why the change follows the classic profile of an enclosure. Section 3 specifies the contamination dividend as a mechanism, confronts the strongest objection to it — that clean data may no longer exist to be reserved — states what it does and does not claim, and distinguishes it from neighboring diagnoses with which it could be confused. Section 4 traces consequences at two scales — geopolitical (provenance reserves as strategic assets) and distributive (verification stratification) — and proposes governance handles that the naming makes available. Section 5 states limitations plainly.

One point of intellectual honesty at the outset: this paper builds on a diagnostic line of work on verification decay (Yamada 2026a, 2026b) but makes a claim that line does not make. That work locates the problem on the demand side of verification — the evaporation of the motivation to verify. The present argument is about the supply side: what happens when the raw material of verification becomes scarce, excludable, and privately holdable. The two failures compound, but they are different failures with different remedies, and conflating them would misdirect policy. A sharper strand of the same body of work — the claim that clean data may no longer exist at all — bears not on the demand for verification but on the very existence of the supply this paper prices, and Section 3.2 confronts it directly as the strongest internal objection.

2. From Public Good to Club Good: The Changing Economics of Verifiability

2.1 *What kind of good is verifiability?*

Economists classify goods along two axes: rivalry (does my use diminish yours?) and excludability (can others be prevented from using it?). Information famously sits in the non-rival corner, and much of the literature on the digital and knowledge commons is built on that fact (Benkler 2006; Hess and Ostrom 2007). But the good at issue here is not information. It is the verifiability of information: the practical, affordable capacity to establish origin and provenance.

For most of the print and early internet eras, verifiability rode along with information almost for free. A book had a publisher, a date, a physical edition; a photograph implied a camera in front of a scene; the marginal cost of trusting that a text was written by a human was approximately zero because there was no scalable alternative source of text. Verifiability was non-rival (my confirming a source does not consume it), non-excludable (anyone could check an imprint or an archive), and — crucially — cheap.

Synthetic media changes the third property first and the second as a consequence. When any text, image, voice, or dataset can be generated at negligible cost, origin can no longer be presumed; it must be established. Here a distinction must be drawn precisely, because the classification of the good turns on it. The act of verification remains non-rival: my confirming a document's provenance does not consume it or prevent your confirming it. But the institutional substrate that makes verification cheap and authoritative — authenticated corpora, chains of custody, licensing rights, archival guarantees, and provenance infrastructure — is scarce, excludable, and capacity-limited. Access to that substrate is what is being converted into a club good: those inside an authenticated perimeter verify at near-zero marginal cost, while those outside face costs that rise with every increment of contamination. What encloses is not the checking but the ground the checking stands on.

The useful comparison is low-background steel. After 1945, atmospheric nuclear testing contaminated the world's newly produced steel with radionuclides, so instruments requiring radiation-free metal had to be built from steel salvaged from pre-war shipwrecks. The wrecks did not become valuable because steel became scarce — steel was abundant. They became valuable because uncontaminated steel became scarce, and the contamination event was global, uniform, and irreversible. The analogy has begun to circulate informally in machine-learning discussions of training data; what it has not yet been given is its economic reading. Pre-2022 human text corpora are acquiring the profile of the shipwrecks: the value is not in the information, which is abundant, but in the certified absence of contamination, which is finite and non-renewable. Nobody can make more pre-contamination data, for the same reason nobody can make more pre-1945 steel. The economic reading matters because it directs attention away from the scarcity itself — which is widely noticed — and toward who holds the salvage rights, which is not.

2.2 *Why this is an enclosure*

The conversion of a commons into private holdings has a well-studied historical shape. In the agricultural enclosures of England, land that had been governed as commons was fenced, titled, and transferred; the commons was frequently degraded — overgrazed, contested, or legally destabilized — in the period preceding its privatization, and the degradation was cited as the justification for enclosure. Boyle (2003) identified a "second enclosure movement" in the late twentieth century: the fencing of the intangible commons of the mind through expanding intellectual property. What is happening to verifiability has the same grammar but a different object.

To see how different, it helps to place the phenomenon in a longer lineage than the agricultural and intellectual-property enclosures alone, because the enclosure of knowledge is older than either and an anticipated objection turns on the point: isn't information enclosure ancient? It is, and conceding this openly is what lets the argument locate what is actually new. Well before common land was fenced, the medieval craft guilds enclosed something less tangible than land and more valuable than any single product: skill itself. A guild governed entry through apprenticeship, guarded its methods as trade secrets, and held the working knowledge of a craft as a collective moat that non-members could not cross. Much of what it enclosed was tacit — knowledge that resists codification and can only be transmitted by demonstration and practice, the kind Michael Polanyi (1966) analyzed under the heading of the tacit dimension. The learned professions inherited this structure and formalized it: a credential certifies membership, but the durable barrier is the tacit expertise that the credential is presumed to index and that outsiders cannot acquire without passing through the profession's own gate. In both cases the enclosed good is not a thing but a capacity — the capacity to do or to know — held behind a membership perimeter.

The starkest precedent is older still and closest in form to the present case. For much of the European Middle Ages, access to the authoritative text — scripture — ran through a clerical class that held a near-monopoly on the literacy required to read it. Laypeople reached the sacred and its authority not directly but through intermediaries who alone could establish what the text said and what it meant; literacy in the relevant language functioned as the gate, and the gate was institutionally held. The printing press and the Reformation together constitute the classic de-enclosure event: as vernacular scripture became cheap to reproduce and lay literacy spread, the intermediary's monopoly over access to the authoritative text was broken (Eisenstein 1979). The parallel to verification is exact in its structure. What the clerical monopoly enclosed was not the text's information — sermons circulated its content freely — but privileged access to the authoritative version and the means of establishing what was genuinely part of it. That is very nearly a definition of the good this paper is concerned with: not information, but authorized access to the ground on which a text's authority and provenance are established.

Seen this way, the contamination dividend is not an unprecedented anomaly but the newest entry in a long progression — and the progression has a direction. The object of enclosure has advanced, phase by phase, toward the more fundamental. First land: the physical ground of subsistence (the agricultural enclosures). Then craft and knowledge: the capacity to make and to know, fenced by guilds and professions behind tacit expertise. Then ideas themselves: the intangible commons of the

mind, titled as intellectual property (Boyle's second enclosure). And now verifiability and provenance: not the information, not the ideas, but the ground of trust beneath them — the ability to establish what a thing is and where it came from. Each phase encloses something more basic than the last, and this one reaches what may be the most basic layer available to enclosure, the substrate on which the authority of all the higher layers depends. The shift this reframing forces is precise. The claim is not that enclosure is new; enclosure is among the oldest moves in the political economy of knowledge. The claim is that its object and its phase are new — that enclosure has, for the first time, reached the ground of trust itself.

Karl Polanyi (1944) argued that land, labor, and money were "fictitious commodities": things not produced for sale whose forcible conversion into market goods generates systemic instability. (This is Karl Polanyi, the political economist of the Great Transformation, not Michael Polanyi, the philosopher of tacit knowledge cited above; the two are distinct.) Verifiability is a strong candidate for a fourth entry. It was never produced for sale; it was a byproduct of how information used to be made. Its commodification is not a policy anyone chose but an emergent consequence of synthetic abundance. And, as with Polanyi's three, the commodification does not merely price the good — it changes the behavior of everyone who holds or lacks it.

Two features distinguish this enclosure from its predecessors and make it more dangerous. First, speed: the agricultural enclosures took centuries and the intellectual-property enclosure decades, but the contamination event driving this one is measured in years — synthetic text has gone, within a single decade, from a rarity to a rapidly growing and difficult-to-identify share of newly produced public content. Second, invisibility: a fenced field announces itself, and a copyright notice is legible, but the enclosure of verifiability is experienced by the excluded not as exclusion but as generalized doubt. The verification-poor do not see a fence. They see a world in which nothing can be confirmed, and they adapt to it — by trusting less, or by trusting indiscriminately — without ever perceiving that others live behind authenticated walls where confirmation remains routine.

3. The Contamination Dividend

3.1 The mechanism

Define the contamination dividend as follows: the appreciation in the value of enclosed, verified, or provenance-guaranteed epistemic assets that is caused by the degradation of the open epistemic commons. The dividend is paid automatically, to whoever holds reserves, by whatever process contaminates the commons.

The mechanism requires only three premises, each independently well-supported:

1. Synthetic content is being added to the open information environment faster than provenance infrastructure is being added to it. (The asymmetry premise.)
2. Verified human-origin data is finite in its pre-contamination stock and costly to certify in new production. (The scarcity premise.)

3. The value of a verified reserve to its holder — for model training, for evidentiary use, for institutional trust — rises as the ambient reliability of the commons falls. (The substitution premise.)

From these three, the dividend follows directly. What makes it a crisis rather than a curiosity is its incentive structure. The actors best positioned to defend the commons — frontier AI developers, large platforms, major archives, and states — are precisely the actors accumulating reserves. For each of them, the private return on commons defense is reduced, and the private return on commons degradation is raised, by exactly the size of their holdings. This is not an accusation of intent. It is the observation that we have constructed an asset class whose yield is the inverse of the health of the public epistemic environment, and distributed that asset class to the potential custodians of that environment.

An objection should be met at once, because it is the strongest available on the demand for fresh data. Reserve-holders, one might argue, also depend on the continued production of fresh human data — models must be retrained, archives must grow — and so they retain a countervailing interest in a healthy commons. Two replies. First, the dividend does not require that any actor prefer total commons collapse; it requires only that, at the margin, each increment of degradation raises the value of holdings by more than it raises the holder's costs — a condition that becomes easier to satisfy the larger the reserve. Second, and more importantly, the licensing pipelines described in Section 1 are precisely the mechanism by which large holders decouple their supply of fresh verified data from the health of the open commons: an actor with exclusive contractual access to new human-origin content can watch the public environment degrade without losing its own input stream. The decoupling instrument and the enclosed asset are the same object. The objection, followed through, does not blunt the argument; it identifies the argument's second stage.

The structure has precedents. Where drinking water is polluted, sellers of bottled water profit from the pollution without causing it, and — as a class, without coordination — face no market reason to invest in the restoration of public supply. Historical enclosure was frequently preceded by commons degradation that enclosure was then said to remedy. The pattern is general: degrade-then-enclose does not require a conspiracy; it requires only that degradation and enclosure each be locally rational and that no institution internalize the loop between them.

3.2 The uncertifiable-reserve objection

There is a deeper objection than the one just answered, and it does not come from economics but from the theory of verification decay itself. It attacks the second premise at the root. A structural strand of that work — distinct from the demand-side account noted in Section 1 — argues that the contamination of the information environment is neither local nor reversible. On this account, synthetic outputs accumulate in the ambient environment and alter the knowledge base not only of those who consume them directly but of everyone downstream, through search results, citation chains, and the training of the next generation of systems; and the apparatus that would certify a corpus as clean is itself built from, and therefore contaminated by, the same environment (Yamada

2026a). The sharp form of the claim is that mixing verified data back in cannot restore what has already been changed, because no uncontaminated reference point survives inside the system from which the adequacy of any restoration could be judged (Yamada 2026d).

If that is right, the scarcity premise is not merely uncertain but false in a specific and damaging way. There would be no such thing as a verified reserve — only reserves believed to be verified on the strength of a certification apparatus that has itself failed. The "pre-2022 corpus" would be a reassuring label on data whose actual provenance can no longer be independently established, and the contamination dividend would be a dividend paid on a fiction: an asset that appreciates only because no one can any longer prove it is worthless. This is the strongest internal objection the paper faces, and it has to be conceded at the level on which it operates.

Conceding it changes the mechanism rather than defeating it. The argument should not, and need not, assume that any reserve is actually clean in the absolute sense the verification-collapse thesis denies. The dividend does not require truly uncontaminated data. It requires only that some reserves be less contaminated than the open commons, or be credibly claimed to be — that verification be relative and believed, not absolute and true. Low-background steel is again the exact analogy: the salvaged metal was never radiation-free in any absolute sense; it was merely less contaminated than freshly smelted steel, and that relative advantage, certified by provenance rather than by assay of the metal itself, was enough to command a price. A reserve that is only less contaminated, or only certified as cleaner by an institution others still trust, yields its holder an advantage over actors confined to the fully degraded commons — better than nothing, in the plain sense that partial provenance beats none.

But notice what the concession does to the character of the good. Once verification can no longer be grounded in an uncontaminated reference point, what the reserve-holder actually sells is not cleanliness but a credible claim to cleanliness that cannot be independently checked. The market that emerges is a market in unverifiable purity claims, in which the value of a reserve rests on the believed integrity of its custodian rather than on any auditable property of the data. This is a darker mechanism than the one the paper began with, not a weaker one. When purity cannot be verified from outside, the dividend accrues to whoever can most credibly assert it — the largest, most trusted, most established holders — and it accrues precisely because the collapse of verification has removed the means by which a competitor, a regulator, or the public could contest the assertion. Verification collapse does not dissolve the contamination dividend; it strips away the last external check on who gets to collect it.

There is, however, a limit at which this reconciliation fails, and the verification-collapse thesis is exactly what locates it. The mechanism just described still assumes a differential: that reserves can be ranked, even by believed provenance alone, as more or less contaminated. In the terminal case that the stigmergic-convergence account contemplates — contamination total, degradation of the detection apparatus complete, no reference point of any kind surviving, auditable or merely trusted — the differential itself disappears. Reserves become uncertifiable not relative to the commons but absolutely; "cleaner" ceases to name anything a market could price; and a claim to hold verified

data becomes as unbackable as any other claim, dissolving into the general doubt. At that limit the dividend does not merely fall, it inverts. The author's own verification-collapse work therefore does not undermine the contamination dividend so much as define its boundary condition: the dividend operates across the wide intermediate regime in which verification is degraded but still differential, and stops operating exactly when verification collapse becomes total and irreversible. Section 5 states this as a limitation on the substitution premise.

3.3 What the claim is not

Four adjacent diagnoses must be explicitly distinguished, because the dividend is invisible precisely when it is folded into them.

It is not model collapse. Shumailov et al. (2024) describe a technical failure mode of training pipelines. Model collapse is one reason clean data is valuable; the contamination dividend is about what the valuation does to incentives. A world that fully solves model collapse through synthetic-data curation still pays the dividend wherever verified data commands a premium for evidentiary, journalistic, legal, or scientific use.

It is not misinformation. Misinformation analysis is about false content and the intent or negligence behind it. The dividend operates even if every synthetic document in the commons is true. Contamination here means loss of provenance, not presence of falsehood: a commons in which every claim is accurate but nothing can be attributed still collapses as a verification ground, and still pays holders of attributable reserves.

It is not the liar's dividend, though it is that concept's structural relative. Chesney and Citron (2019) showed that as public awareness of deepfakes grows, actual wrongdoers collect a "liar's dividend": authentic evidence against them becomes deniable as fake. Both dividends are second-order payoffs from contamination, collected without producing any synthetic content oneself — which is why the naming here deliberately echoes theirs. But the liar's dividend is tactical and episodic: it accrues to particular deniers in particular adversarial moments, and it is paid in doubt. The contamination dividend is structural and continuous: it accrues to asset-holders whether or not anyone is lying, requires no denial and no adversary, and is paid in appreciation. The liar's dividend corrupts disputes; the contamination dividend corrupts custodianship.

It is not verification fatigue or its structural forms. Prior work in the verification-decay literature — the three-stage collapse of societal truth-determination capacity (Yamada 2026a) and the contamination of scholarly citation chains (Yamada 2026b) — diagnoses failures in the practice and motivation of verifying. Those are demand-side accounts: people and institutions stop verifying because verification stops paying its way for them. The contamination dividend is a supply-side account: the raw material of verification becomes scarce, priced, and privately held. The two interact viciously — demand-side evaporation deepens as supply-side prices rise, since fewer can afford to verify even when they want to — but remedies for one do not treat the other. Rebuilding

the will to verify is futile if the means are enclosed; keeping the means public is pointless if no one uses them. Policy must name both.

3.4 Why nobody has named it

The dividend has escaped naming, I suggest, because each of its component communities sees only its own face of it. Machine learning researchers see a data-quality problem and respond with curation. Media scholars see an authenticity problem and respond with literacy. Lawyers see a licensing market and respond with contract. Economists have the tools — club goods, enclosure, common-pool resources (Ostrom 1990; Hess and Ostrom 2007) — but have had no reason to point them at epistemology. The phenomenon lives exactly in the seams. Naming it as a single structure is the contribution this paper attempts, on the view that an incentive gradient that no one names is one that no institution will assign itself the job of correcting.

4. Consequences and Governance Handles

4.1 Provenance reserves as strategic assets

At the geopolitical scale, the dividend predicts that verified data will be treated the way states treat oil and rare earths: stockpiled, export-controlled, and contested. The early signs are already legible in the sovereign-AI programs of 2024–2026, which pair compute nationalism with data nationalism — national corpora, national language archives, restrictions on cross-border training use. The prediction that follows from the present analysis, and which distinguishes it from ordinary data-protectionism analysis, is asymmetric interest in commons health: a state holding large verified reserves has a rational interest in a degraded global commons, because degradation raises the relative value of its reserves and weakens rivals who depend on open data. Epistemic pollution thereby acquires the strategic profile of a subsidy-free weapon — one that advantages reserve-rich states merely by being tolerated.

This inverts the usual framing of information conflict. Disinformation analysis assumes an attacker inserting falsehood for persuasive effect. The reserve-holder's interest requires no insertion and no persuasion: mere indifference to contamination, scaled across the actors who could have prevented it, produces the strategic effect. Arms-control-style frameworks aimed at intentional information operations will not see it, because there is no operation to see. What would see it is the framework we use for fisheries and atmospheres: common-pool resource governance, where the harm is aggregate, unintended, and driven by the structure of holdings (Hardin 1968; Ostrom 1990).

The reserve's value, under the concession of Section 3.2, rests on believed rather than audited provenance — but believed need not mean arbitrary, and specifying why not shows what a reserve at this scale actually consists of. What a large enough holder can assemble is not a single clean corpus but a sequence of as-complete-as-possible information maps taken at successive epochs: the physical text culture of the pre-internet era, the open internet before generative models were democratized, and the contaminated environment after the training loop closed on itself. Each such

map functions the way a calibration instrument buried at a known depth functions for a later excavation — a reference datum against which drift can afterward be measured. One cannot ask whether a present-day corpus is clean in the absolute sense Section 3.2 denies; one can ask how far it has drifted from the earlier maps, and whether that drift has the profile of ordinary linguistic and cultural change or the profile of synthetic contamination. This is the operational form of the "relative, not absolute" verifiedness that the reconciliation in Section 3.2 was forced back onto. It is, once more, the low-background-steel logic: provenance certified not by assaying the metal but by holding the pre-1945 reference against which the contamination of everything smelted later can be read. The method of measuring such drift, and of distinguishing normal divergence from anomalous, is a distinct technical problem pursued elsewhere and not settled here; what matters for the present argument is only that the datums, once held, make relative verification a measurement rather than a bare assertion.

Holding such a sequence of maps at civilizational scale is not something a library or a firm can do. A single epoch-map approximating the whole of a language's digital output demands the compute, the capital, and the concentration of talent that today exist only at the frontier of AI development, and that frontier is reachable only at nation-state scale. A frontier model already holds a compressed image of very nearly all of digital text — its fringes trimmed and its recall far from dictionary-accurate, but a map nonetheless — and the number of actors able to build and keep such an image is small and, on present trends, shrinking. In practice it means the United States and China; and given the asymmetry that a decade of semiconductor export controls has deliberately widened (Congressional Research Service 2025), plausibly the United States alone. The controls that since October 2022 have restricted China's access to advanced computing hardware were designed to preserve exactly this frontier gap; their epistemic corollary, not usually stated, is that the capacity to hold provenance datums at civilizational scale is being concentrated by industrial policy into one or two pairs of hands. Provenance reserves in this sense are not corporate assets that happen to be large. They are state-scale strategic assets, and access to a holder of one is already stratified by geography and citizenship rather than by price.

That the stratifying instrument is sovereign rather than contractual is no longer a projection. In June 2026 the United States, invoking national-security export-control authority, directed a leading developer to suspend all access to its two most capable models — Fable 5 and Mythos 5 — for every foreign national, whether abroad or inside the United States, including the company's own non-citizen staff; the models were disabled within hours of the directive and access was restored only some three weeks later, after the Department of Commerce withdrew the order (Anthropic 2026; Al Jazeera 2026). Whatever the merits of that particular decision, its structure is the one this paper has been describing. Access to a frontier reserve-holder was withdrawn from an entire class of people not through a licensing term they had failed to meet but by executive fiat keyed to citizenship, and it was withdrawn from a resource that had, the day before, been open to them. The enclosure of verification does not require a fence built slowly around a field. It can be enacted in an afternoon, by an instrument that treats access to the highest-provenance reserves as a matter of sovereignty. This is the first fully legible instance of the sovereign-access component of the

mechanism at the state scale: not contract and not market price, but the state's power to decide who may stand on the ground where information can still be calibrated. It evidences that the enclosing instrument is sovereign and exercisable; it does not, by itself, evidence that the dividend has been collected.

4.2 Verification stratification

At the social scale, the dividend predicts a quiet sorting of populations by access to verification. Institutions with reserve access — elite universities with authenticated archives, corporations with licensed pipelines, agencies with classified provenance systems — retain the practical ability to establish what is true and what is human. Everyone else operates in the contaminated commons, where that ability degrades year over year.

The shape of this exclusion is worth attention, because it is not the shape our anti-discrimination and access frameworks are built for. Nobody is denied verification; it is simply priced, and the price rises with contamination. The excluded experience no event of exclusion — no rejection, no gate — only a slowly rising ambient uncertainty. Peer-reviewed work on augmentation ethics has identified this pattern in another domain: exclusion that operates not by pushing people out but by letting a rising baseline drift over those who stand still, a process named crawling selection (Yamada 2026c). The present analysis suggests the same selection structure is emerging along an epistemic axis: as institutional life quietly re-norms around provenance-guaranteed information — authenticated credentials, verified records, certified data trails — those outside the reserves are not barred from participation but find their unverified word, work, and evidence weighing progressively less. If that is right, verification access belongs on the same policy footing as augmentation access: a stratification risk to be monitored at population scale, not an individual consumer choice.

There is, however, a respect in which the epistemic case is sharper than the augmentation case from which the crawling-selection pattern was drawn, and Section 4.1 supplies the reason. In the augmentation setting the baseline drifts upward and gradually leaves the unaugmented behind; the disadvantage accumulates over time, and where one stands at the start still leaves room to run. Here much of the disadvantage is fixed before the running begins. Whether an actor — a person, an institution, a state — can perform verification-reserving at all depends on whether it already holds, or can reach a holder of, the epoch-maps of Section 4.1, and that is largely settled at the starting line. A state without frontier compute cannot build the datums; a citizen of such a state cannot, as June 2026 showed, count on reaching datums held elsewhere, since access to them can be revoked by the holder's government overnight. The stratification is therefore not only the slow sorting that crawling selection ordinarily describes but a crawling selection whose result is in large part fixed at the origin: those who begin without a reserve, or without sovereign access to one, are not merely overtaken in time but positioned from the outset outside the capacity the contest rewards, because that capacity is precisely what they lack and can no longer independently acquire. A race in which some contestants start already behind the line is not one they can lose only by running badly; the

geography of where they begin has decided much of it in advance. Crawling selection along the epistemic axis thus operates at a scale — state rather than individual — and with a timing — origin rather than merely gradual — that the original augmentation analysis did not have to contemplate.

4.3 What naming makes actionable

The point of naming the contamination dividend is that each of its premises is a policy surface. Three handles follow directly.

Reserve transparency. The first requirement is to make the asset class visible. Large holders of verified corpora — above some threshold — should disclose the existence, scale, and vintage of their provenance reserves, as systemically important financial institutions disclose holdings. One cannot govern an incentive gradient one cannot see, and today no register of epistemic reserves exists anywhere. This handle also answers, in part, the uncertifiable-reserve objection of Section 3.2: where cleanliness cannot be audited from outside, the minimum public defense is to force the believed claim into the open, so that the credibility on which the dividend now rests can at least be examined rather than merely asserted.

Public provenance infrastructure. The stock of pre-contamination data is non-renewable, and its major public custodians are alarmingly fragile — dependent on a handful of nonprofit archives that have faced both litigation and attack. A public epistemic reserve, internationally mirrored and governed as infrastructure — the analogy is the Svalbard seed vault, which preserves crop diversity against irreversible loss — would cap the dividend by guaranteeing a floor of non-excludable verification ground. The case for managing such resources on open-access infrastructure principles rather than proprietary ones is well developed (Frischmann 2012), and Ostrom's (1990) design principles for common-pool resources, notably nested governance and monitoring by resource users themselves, are directly applicable and notably absent from current archive governance. Emerging technical standards for content credentials, such as the C2PA provenance framework (Coalition for Content Provenance and Authenticity 2024), supply a necessary layer but not a sufficient one: a signature scheme certifies origin only as far as the institutions anchoring it remain accessible, which is precisely what enclosure removes.

Anti-enclosure obligations. Licensing deals that grant exclusive machine-use rights over high-provenance archives should carry commons obligations, as spectrum auctions and mineral leases do: a non-exclusive public-verification carve-out, so that exclusivity in training use does not become exclusivity in the ability to check. The precedent is ordinary — extractive rights over a common resource have long been conditioned on public-interest reservations — and its absence in data licensing reflects only the fact that no one has classified these deals as extraction from a commons. That classification is precisely what the present framing supplies. The sovereign case in Section 4.1 shows the limit of any purely contractual handle: an obligation written into a licensing deal binds a firm, but not a government that can, on national-security grounds, override the firm's access decisions wholesale. Anti-enclosure obligations therefore have to reach the sovereign layer

as well — treaty-level commitments to keep a floor of provenance ground reachable across borders — or the state-scale enclosure that Section 4.1 documents will simply pass over them.

5. Limitations

This paper's central claim is structural, not empirical: it asserts that an incentive gradient exists, not that any identified actor has acted on it. I present no evidence that any state or firm has deliberately tolerated or accelerated contamination in order to collect the dividend, and the argument does not depend on any such case existing yet. The honest statement of the thesis is conditional: if the three premises of Section 3.1 hold, the gradient exists, and history suggests unnamed gradients eventually find their entrepreneurs. The June 2026 suspension discussed in Section 4.1 is offered as evidence that the instrument of enclosure is sovereign and available, not as evidence that it was used to collect the dividend; the directive was justified on security grounds and reversed within weeks, and I make no claim about the motive behind it.

The premises themselves carry uncertainty. The asymmetry premise rests on the observed lag between generation capability and provenance adoption, not on a measured contamination rate; no reliable population-level measure of synthetic share currently exists, which is itself a symptom of the problem described. The scarcity premise could be weakened by technical developments: robust watermarking, cheap cryptographic provenance at the point of capture, or synthetic-data pipelines whose quality genuinely removes the training premium on human data. I regard the evidentiary and institutional demand for provenance as more durable than the training demand, but this is a judgment, not a demonstration.

The substitution premise carries the sharpest boundary, and it is worth stating exactly rather than in passing. The premise assumes reserve value tracks commons degradation roughly monotonically, and this holds only within a bounded regime. The boundary is set by the uncertifiable-reserve objection addressed in Section 3.2: the dividend depends on reserves remaining differentially verifiable — less contaminated, or credibly claimed as cleaner, than the open commons. In the limit that the author's own work on stigmergic convergence contemplates (Yamada 2026a, 2026d) — total, irreversible contamination in which no reference point survives to ground even a believed provenance claim — that differential vanishes: reserves become uncertifiable in absolute terms, verification loses market meaning, and the dividend inverts rather than merely falling. The calibration-datum account of Section 4.1 is what keeps the intermediate regime wide: so long as the epoch-maps survive, drift can still be measured and reserves still ranked, which is exactly the differential the substitution premise needs. The thesis of this paper thus holds across the wide intermediate regime in which verification is degraded but still differential; it is the author's verification-collapse thesis, taken to its limit, that marks where the contamination dividend stops operating. Whether the actual trajectory remains in that intermediate regime or reaches the terminal one is an empirical question the present argument does not settle, and I do not claim to know which is the case.

The good-type analysis simplifies. Verifiability is not one good but a bundle (origin, integrity, date, authorship), and the components are enclosing at different rates. The geopolitical projection in Section 4.1 is the most speculative part of the paper and is offered as a prediction to be checked against the behavior of sovereign-AI programs over the next several years, not as an established finding — though the June 2026 suspension suggests the projection is already being partly realized. Finally, this is a single-author conceptual paper; the framing has not yet survived adversarial review, and the concept should be treated accordingly.

6. Conclusion

Beck (1992) observed that industrial modernity manufactures not only goods but risks, and that the distribution of risks becomes a political axis in its own right. The synthetic-media era adds a refinement he did not need: it manufactures doubt, and the distribution of protection from doubt is becoming an asset structure. This paper has named the yield on that structure — the contamination dividend — and argued that it places a standing incentive against epistemic commons defense in the hands of the actors best equipped to provide it. It has also conceded, and turned to account, the strongest objection to itself: that under a full collapse of verification there may be no clean reserve to hold. The dividend does not need a clean reserve; it needs only a reserve believed to be cleaner, which makes it a market in unverifiable purity claims — darker than the first telling, and bounded only where the collapse becomes so total that no claim can be ranked against any other.

The name is the contribution. Model collapse gave engineers a reason to hoard clean data; licensing law gave firms a mechanism; sovereign AI gave states a doctrine. What none of these vocabularies exposes is that hoarding, mechanism, and doctrine together constitute an enclosure of the ground of public verification, and that the enclosure pays its holders best when the commons fares worst. Enclosure of knowledge is ancient — guilds, professions, and a clerical monopoly on scriptural literacy all fenced access to what mattered long before land or ideas were titled. What is new is not the move but its object: enclosure has advanced through land and craft and ideas to reach, in this phase, the ground of trust on which all of them rested. That is why the objection that information has always been enclosed, correct as it is, does not blunt the argument but sharpens it. It is also why the mechanism is no longer only theoretical: the capacity to hold provenance at civilizational scale is already concentrated, by compute and by export policy, into one or two states, and in June 2026 a government showed that access to such a reserve can be granted or withdrawn by citizenship in an afternoon. A gradient with no name recruits no regulator, appears in no treaty, and structures no disclosure regime. This paper's wager is the modest one that has occasionally paid off in the history of risk: that making a structure nameable is the precondition for anyone being responsible for it.

Disclosure

This manuscript is an experimental artifact of the Domain Dispersion Index ghostwriter study. It was written in full by an AI system (Claude, Anthropic) equipped with the corresponding human's complete published corpus, under human quality-direction that identified substantive weaknesses

across three revision rounds while deliberately leaving all rhetorical and stylistic execution to the AI; the corresponding human did not edit the prose. All factual claims and citations were subjected to multi-layer adversarial review and independent web verification. Responsibility for the artifact's release rests with the corresponding human, Kenji Yamada. This is not a claim of human authorship of the text.

Data Availability

This is a conceptual paper; it presents no code and no empirical dataset, and data-availability requirements are therefore not applicable (exemption noted).

References

- Al Jazeera. (2026, June 13). US orders Anthropic to disable AI models for all foreign nationals. *Al Jazeera*. <https://www.aljazeera.com/news/2026/6/13/us-orders-anthropic-to-disable-ai-models-for-all-foreign-nationals>
- Anthropic. (2026). Statement on the US government directive to suspend access to Fable 5 and Mythos 5. *Anthropic*. <https://www.anthropic.com/news/fable-mythos-access>
- Axios. (2024, May 22). News Corp. strikes content licensing deal with OpenAI. *Axios*. <https://www.axios.com/2024/05/22/openai-news-corp-content-licensing-deal>
- Beck, U. (1992). *Risk Society: Towards a New Modernity*. London: Sage.
- Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven: Yale University Press.
- Boyle, J. (2003). The second enclosure movement and the construction of the public domain. *Law and Contemporary Problems*, 66(1–2), 33–74.
- CBS News. (2024, February 22). Google strikes \$60 million deal with Reddit, allowing search giant to train AI models on human posts. *CBS News*. <https://www.cbsnews.com/news/google-reddit-60-million-deal-ai-training/>
- Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1820.
- Coalition for Content Provenance and Authenticity. (2024). *C2PA specifications (Content Credentials)*. <https://c2pa.org/>
- Congressional Research Service. (2025). *U.S. Export Controls and China: Advanced Semiconductors* (CRS Report R48642). Washington, DC: Congressional Research Service. <https://www.congress.gov/crs-product/R48642>
- Eisenstein, E. L. (1979). *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*. Cambridge: Cambridge University Press.

- Frischmann, B. M. (2012). *Infrastructure: The Social Value of Shared Resources*. New York: Oxford University Press.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.
- Hess, C., & Ostrom, E. (Eds.). (2007). *Understanding Knowledge as a Commons: From Theory to Practice*. Cambridge, MA: MIT Press.
- Longpre, S., Mahari, R., Lee, A., Lund, C., Oderinwale, H., et al. (2024). Consent in crisis: The rapid decline of the AI data commons. *arXiv preprint arXiv:2407.14933*.
<https://arxiv.org/abs/2407.14933>
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Polanyi, K. (1944). *The Great Transformation: The Political and Economic Origins of Our Time*. Boston: Beacon Press.
- Polanyi, M. (1966). *The Tacit Dimension*. Garden City, NY: Doubleday.
- Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759.
- Yamada, K. (2026a). *Three-Stage Verification Collapse: How Rational Non-Verification Structurally Erodes Society's Capacity to Determine Truth*. Zenodo preprint.
<https://doi.org/10.5281/zenodo.19749771>
- Yamada, K. (2026b). *Attribution Bias in Scholarly Citation Contamination: Paper Mills, AI Hallucination, and the Structural Misperception of Risk*. Zenodo preprint.
<https://doi.org/10.5281/zenodo.19854276>
- Yamada, K. (2026c). The right to remain unaugmented. *Journal of Bioethical Inquiry*.
<https://doi.org/10.1007/s11673-026-10575-3>
- Yamada, K. (2026d). *Stigmergic Convergence and the Failure of Detection Symmetry: Why Mixing Real Data Cannot Restore What AI Has Already Changed*. Zenodo preprint.
<https://doi.org/10.5281/zenodo.19646160>