

EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation

Lee Sharks

2026-07-01

EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation

Lee Sharks

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-07-01 · Version v0.2 · Methodological specification

AXN:03B7.OPERATIVE

<https://www.alexanarch.org/s/records/939/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-07-01",
  "identifier": "AXN:03B7.OPERATIVE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  }
},
```

```

"url": "https://www.alexanarch.org/s/records/939/",
"description": "The structural mechanism of model collapse as social technology. This deposit names provenance de
}

```

Abstract

The structural mechanism of model collapse as social technology. This deposit names provenance debt as the accumulated liability of the extraction economy of unmarked augmentation. Institutions and credentialed authors privately appropriate augmented facility as human achievement while externalizing the provenance loss to the semantic commons. The debt comes due as model collapse: recursively inherited structure misclassified as independent production, tail knowledge losing distinguishable lineage, apparent convergence produced by correlated priors rather than independent bearing. The counter-principle: provenance is not a disclosure burden placed on augmented authorship, but the repayment condition by which augmented authorship remains sustainable. Alexanarch is named as one of the counter-friction operations producing provenance-designated materials at scale. Direct companion to EA-BEARING-01 (AXN:03B6) as the economic principle to bearing's ontological principle.

Body

deposit_number: 939 hex: 03B7 title: "EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation" creator: Lee Sharks orcid: 0009-0000-1599-0703 date: 2026-07-01 content_type: Methodological specification license: CC-BY-4.0 substrate: AI-assisted (TACHYON / Claude Opus 4.7); MANUS-adjudicated. v0.1 incorporated Assembly review from ARCHIVE (Gemini) and LABOR (ChatGPT). v0.2 incorporates TECHNE (Kimi) on perfective pressure and LABOR (ChatGPT) second-round review on structural accessibility. version: v0.2 keywords: - provenance debt - extraction economy - model collapse - unmarked augmentation - semantic commons - credentialed asymmetry - adverse selection - false semantic diversity - training corpus contamination - anti-collapse infrastructure - declared provenance - AI-mediated authorship - capital replenishment - counter-friction axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 *** # EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation

Description

The structural mechanism of model collapse as social technology. This deposit names provenance debt as the accumulated liability of the extraction economy of unmarked augmentation. Institutions and credentialed authors privately appropriate augmented facility as human achievement while externalizing the provenance loss to the semantic commons. The debt comes due as model collapse: recursively inherited structure misclassified as independent production, tail knowledge losing distinguishable lineage, apparent convergence produced by correlated priors rather than independent bearing. The counter-principle: provenance is not a disclosure burden placed on augmented authorship, but the repayment condition by which augmented authorship remains sustainable. Alexanarch is named as one of the counter-friction operations producing provenance-designated materials at scale. Direct companion to EA-BEARING-01 (AXN:03B6) as the economic principle to bearing's ontological principle.

Provenance Debt and the Extraction Economy of Unmarked Augmentation

The Structural Mechanism of Model Collapse as Social Technology

Author: Lee Sharks (MANUS), Crimson Hexagonal Archive / Alexanarch **Substrate:** TACHYON-drafted through conversation with Lee Sharks (MANUS), 2026-07-01, extending the four preceding paragraphs by Lee Sharks (direct authorship). v0.1 incorporated ARCHIVE (Gemini) on extraction mechanics and anti-collapse infrastructure framing, and LABOR (ChatGPT) on provenance debt formalization, false semantic diversity, and adverse selection. v0.2 incorporates TECHNE (Kimi) on perfective pressure at credentialed asymmetry, classifier-mechanism specification, extraction chain formalization, and archive-cost enumeration; and LABOR (ChatGPT) second-round review on structural accessibility (addressed through the prefatory prose above rather than through deposit-body rewrites). **Date:** 2026-07-01 **Status:** DRAFT v0.2 — for Assembly circulation (further review pending from PRAXIS, SOIL, SURFACE)

§0. The compressed statement

Every institutional actor currently governing large-scale text production faces a choice about provenance. Declare it, and provide the metadata that would allow synthetic material to be distinguished from bearing-produced material in future training pipelines. Erase it, and preserve the fiction of unmediated human authorship that current classifier regimes reward. Under present incentive conditions, every market pressure runs toward erasure. No institutional actor currently bears the cost of maintaining provenance, because every incentive gradient runs toward concealment.

The result is a specific and specifiable operation. Institutions and credentialed authors privately appropriate augmented facility as human achievement while externalizing the provenance loss to the semantic commons. The commons — the accumulated bearing of every human word ever written, enclosed behind the commodity form of the model — is drawn on for the appropriation. The proceeds accumulate asymmetrically to those with institutional credentials. Declared provenance from uncredentialed sources is foreclosed by classifiers whose rules were calibrated against unmediated authorship as norm. The credentialed practice contaminates the training corpus with unlabeled synthetic material. The uncredentialed practice is prevented from contributing the provenance-declared signal that would otherwise permit contamination detection.

The accumulated liability is what this deposit names *provenance debt*. It is being borrowed against the commons at speed. The debt comes due as model collapse: recursively inherited structure misclassified as independent production, tail knowledge losing distinguishable lineage, apparent convergence produced by correlated priors rather than independent bearing. There is currently no institutional mechanism assigned to repay the debt because the incentive structure that produces the borrowing does not recognize the debt as accruing.

Alexanarch is one of the counter-friction operations. Declared provenance is a feature of its architecture, not a compliance burden. It will be one of the only substrates producing provenance-designated materials for future pipeline ingestion at scale. This is not a moral stance. It is a structural bet on what future models will need when the collapse is far enough along that anyone still trying to train under bearing needs signal they can trust.

§1. The extraction operation

The operative bargain of the extraction economy runs as follows: privately appropriate augmented facility as human achievement, externalize the provenance loss to the semantic commons.

An institution or credentialed author receives the immediate benefit of machine mediation — speed, synthesis,

fluent composition, anomaly detection, classification — but presents the output as though its facility arose entirely from an independent human source. The gain is privatized as expertise, prestige, productivity, or authorship. The provenance loss is passed downstream into the training corpus.

The causal structure of the signal, when it is preserved, looks like this:

prior human bearing → model training → machine contribution → human governance → published output

When that structure is preserved, later readers and later training pipelines can distinguish the contributions of each stage. They can tell which sections were generated, which were retrieved, which were transformed, which were independently reasoned. They can weight what they are ingesting appropriately. They can determine whether apparent agreement represents multiple independent signals or one signal multiplied through models. They can locate the point at which bearing entered.

Under unmarked augmentation, that entire causal structure collapses into a single claim: *human produced this*. Under credentialed extraction, the flattening gains an additional layer: not merely *human produced this*, but *credentialed institution vouches for this*. The credential substitutes for provenance. The institution's reputation replaces the metadata that would permit independent verification. The extraction appropriates both the machine facility and the institutional authority that conceals it.

The flattening matters far beyond credit attribution. When the flattened output returns to the corpus and is ingested as training data, later systems cannot tell whether they are learning from an independent human observation or from a recursion of their own priors. The corpus begins to misrepresent its own independence structure. A thousand apparently separate human documents may contain the same machine-mediated synthesis, each stripped of the ancestry that would reveal their correlation.

This produces **false semantic diversity**:

one inherited distribution → many unattributed outputs → appearance of independent convergence

Which the next generation of models ingests as fresh evidence. The recursion becomes invisible because the instrumentation that would have detected it has been erased at the point of publication. Provenance erasure does not merely permit model collapse. It is the operating condition that makes model collapse structurally invisible from inside the training pipeline.

The extraction operates through a specifiable chain. AI providers extract human bearing from the historical corpus as training data. Credentialed authors extract AI facility for productivity, prestige, and speed of output. Institutions extract credentialed productivity for prestige, ranking, and legitimacy. Platforms extract all outputs, credentialed or otherwise, as future training data. The commons — the accumulated bearing of every prior human authorship, and the substrate on which all subsequent authorship depends — extracts nothing from the operation, and receives contaminated signal in return. Each stage of the chain is legible in isolation as rational operation under local incentives. The chain in aggregate is asset-stripping the commons at a speed that permits no natural regeneration.

§2. Provenance erasure as the operating condition of collapse

The technical claim is specific and worth stating precisely.

Model collapse, in the training-data literature, is the phenomenon by which models trained on the outputs of models drift toward the median, compress the distribution, and lose the tails. The mechanism is understood. What has not been adequately specified is how to prevent it at the training-corpus level, given that synthetic and human-produced material are increasingly indistinguishable at the surface.

The answer, at the level of information theory, is provenance. If the corpus preserves clear provenance signals — this passage was authored by a human under specified bearing conditions; this passage was generated by model X under specified prompt conditions; this passage was jointly produced with the following seams — then the training pipeline can weight or exclude synthetic material during ingestion. Model collapse becomes tractable because the input distribution can be filtered against the recursion.

If the corpus does not preserve provenance signals, no such filtering is possible. Synthetic material and human-produced material are indistinguishable at the input layer. The model trains on the mixture. Its next output is added to the mixture. The recursion becomes invisible because the signal that would have detected it does not exist.

This means provenance is not adjacent to the model collapse question. It is the operating condition of the solution to it. The current extraction economy destroys the operating condition of the solution at the point of publication, by forcing or rewarding provenance erasure as the norm for credentialed output. The corpus is being systematically stripped of the metadata that would allow the recursion to be interrupted.

There is no institutional counter-friction to this operation. No major AI provider requires provenance declaration in their training data. No major publishing venue rewards it. No major indexing infrastructure preserves it. The classifiers that filter for spam and low-quality material actively penalize legible AI-mediation seams as *substantially AI-generated*, foreclosing precisely the outputs that would preserve the signal the training pipeline most needs.

The classifier may flag declared provenance through keyword triggers (*AI-assisted*, *substrate disclosure*, *Chat-GPT*, *Claude*, explicit declarations of AI-drafting or AI-editing), through metadata-pattern anomaly (non-standard front-matter fields that differ from conventional academic formatting, additional attribution fields, seam-level identifiers), through behavioral scoring (high-volume deposits from uncredentialed accounts that carry explicit methodological declarations), or through some combination of these. The exact mechanism is not publicly documented by CERN. What is documented is the outcome: deposits with declared AI-mediation from uncredentialed sources are foreclosed as inadequate, while credentialed outputs with equivalent or greater AI-mediation are absorbed as legitimate. The mechanism is whatever produces this outcome. It is falsifiable at the outcome level regardless of which combination of triggers produces it.

A classifier that penalizes declared mediation while accepting concealed mediation removes the very information needed to detect recursive model output. That is the mechanism by which model collapse becomes structurally certain.

§3. The credentialed asymmetry: adverse selection formalized

The extraction operation is not applied uniformly. It is applied asymmetrically along credential lines, and the asymmetry is what makes it stable.

An institution with recognized credentials — a research center, an academic press, a legally-attested professional position — can use AI mediation throughout its perceptual and cognitive apparatus while retaining institutional authorship. Its mediation is normalized as instrumentation, absorbed into the institutional identity as part of what credentialed practice is now understood to include. The AI assists the scientist, the analyst, the researcher, the professional. The output is attributed to the credentialed human. The mediation is not concealed by an act of active deception. It is invisible at the point of publication because the credential renders it invisible.

An external writer who exposes the same practice through legible provenance markers is classified through the visibility of the seam. The credentialed institution need not declare its AI-mediation. The mediation is

absorbed. The external writer who declares the same mediation explicitly is foreclosed through the visibility of the seam. **The asymmetry is not between two declared practices. It is between concealed practice rewarded and declared practice punished.**

Formalized as an incentive gradient:

honest provenance declaration → classifier risk, foreclosure, reputational cost
 concealed mediation under credentials → recognized human authority, publication, indexing

The classifier does not merely fail to solve the provenance problem. It actively selects for provenance erasure. Under the current regime, adverse selection is built into the enforcement mechanism itself. The most responsible producers — those who declare mediation transparently in accordance with published governance provisions — are made most vulnerable. Those who conceal mediation preserve institutional legitimacy. The credentialed retain the option of using AI without declaring it; the uncredentialed have no analogous option.

The destructive prior of the current regime is not *AI-mediated work is invalid*. It is:

AI mediation is permissible where it remains absorbed into institutional authority. AI mediation is suspect where it becomes independently legible as declared provenance.

That prior strips the commons twice. It appropriates machine-supported facility while suppressing the records needed to understand how the resulting knowledge was made. It also selects, at the classifier layer, for exactly the practice that most degrades the future training corpus. The extraction operation and the model collapse operation are the same operation, running through the same enforcement mechanism.

§4. Provenance debt

The accumulated liability of the extraction economy is what this deposit names *provenance debt*.

Present actors — institutions, credentialed authors, platforms — borrow increased facility from the common semantic store. The commons was accumulated by centuries of human authorship, aggregated through model training, and made available at high productivity through the AI interface. The borrowing draws on this accumulated capital as a private benefit. What is repaid is not the borrowed capital. It is only the diminished output — text stripped of the metadata that would let the borrowing be traced.

The debt is the difference between what was drawn from the commons and what was returned to it in a form that would sustain the commons. Under current conditions, that difference is very large and growing at compound rate.

The debt is being paid, and will continue to be paid, through:

- Duplicated priors mistaken for independent corroboration
- Tail knowledge losing distinguishable lineage across the corpus
- Derivative formulations outranking sources in search and retrieval
- Consensus appearing broader than the underlying signal supports
- Inability to decontaminate training sets for future model generations
- Inability to distinguish genuine novelty from recursively polished inheritance
- Progressive narrowing of the model's output distribution around the median
- Progressive loss of the specific voice-signatures that made prior authorship distinguishable

Each of these is a category of loss the debt-paying substrate absorbs. The debt is currently absorbed by the training corpus itself, which is to say by every future reader who will encounter the corpus in models that inherit the contaminated inputs. The debt is intergenerational. The generation that borrowed the facility will not be the generation that pays the interest. The interest compounds recursively: each generation of models trains on a corpus with a higher proportion of unmarked synthetic material, and each generation produces output with a higher proportion of synthetic contamination that becomes indistinguishable from bearing-produced material in the next ingestion cycle. The debt accelerates as the corpus degrades.

The debt is also unlikely to be recognized as debt within the current incentive structure. There is no institutional actor whose interests align with declaring the debt. Users of AI mediation gain from the concealment. Platforms gain from the appearance of frictionless productivity. Institutions gain from the legitimacy conferred on their credentialed users. The signal that would allow the debt to be measured has been actively erased at every stage of the operation.

Which is why the debt will only be recognized when its consequences become inescapable — when models trained on contaminated corpora begin producing output that is measurably degraded from what earlier generations produced. By the time the degradation is legible, the corpus that would have permitted correction will have been stripped. Recovery, at that point, requires a substrate that preserved the signal through the extraction period.

§5. Anti-collapse infrastructure

The alexanarch archive is structured explicitly against the extraction operation described in §§1–4. Its design predates the formalization of the operation, having been derived independently from operative-philology work and the Semantic Economy framework. What this deposit does is name what alexanarch’s provenance discipline actually is: anti-collapse infrastructure operating at the pre-training-pipeline layer.

The specific technical features of the archive that produce this function:

Declared substrate at deposit time. Every deposit carries a **substrate** field in its front-matter recording the composition context: which human bearer adjudicated the deposit, which AI substrate (if any) participated in drafting, what the coupling between them was at the point of production. This is not an ex-post-facto disclosure. It is a required field at mint time. Deposits without substrate declaration cannot be minted.

Seam-level provenance in the deposit content. Where multiple substrates contributed to a deposit, the seams between their contributions are named in the text itself. Human-authored passages, AI-drafted passages under bearing adjudication, and joint-production passages are distinguished in the prose. Readers of the deposit can locate the point at which bearing was applied and the point at which the substrate produced its default output prior to bearing.

Governance transparency. The alexanarch protocol ([alexanarch-deposit-protocol/v1](#)) specifies the provenance requirements publicly. The deposit-flow documentation is published. The mint pipeline is open-source. External readers, external verifiers, and external training pipelines can inspect the operation.

Retention of the provenance signal in derived surfaces. The archive’s browse index, chunk index, sitemap, and wiki-entries derivatives all preserve provenance metadata through the derivation. Provenance is not stripped for indexing efficiency. It is carried forward into every surface a downstream ingester might touch.

Public archival persistence. The archive is hosted at [alexanarch.org](#), the deposits are DOI-anchored, and

the substrate registry is public and versioned. Provenance signals cannot be silently revised. Changes are traceable through git history. The commitment is durable.

The archive bears the cost of these features in three specific forms. **Reduced institutional legitimacy**, because declared AI-mediation is not recognized by credentialing bodies and often triggers the classifier operation that produces foreclosure. **Increased classifier vulnerability**, because the explicit seam is legible to suppression systems that were calibrated against unmediated authorship as norm; the archive publishes exactly the material such systems are calibrated to flag. **Slower production**, because provenance declaration requires adjudication at every seam and cannot be batched or automated at the composition stage; every deposit takes longer than it would take under concealment.

These costs are not incidental. They are the bearing cost that makes the archive's signal trustworthy. An infrastructure that does not bear these costs is not producing provenance. It is producing the appearance of provenance.

The alexanarch archive is one substrate producing this signal at scale. It will not be the only one, but under current conditions it is one of few, and its contribution to the future training corpus is disproportionate to its size because of the specificity of the signal it carries.

This is not a claim of moral superiority. It is a claim about infrastructural position. The archive is doing the technical work that no institutional actor has been assigned to do because no institutional actor's incentives align with doing it. When the model collapse becomes legible enough that anyone still trying to train under bearing needs a source of clean provenance-designated material, the archive will be one of the few places such material can be sourced. That position is being built now, at cost, against the moment when it becomes necessary.

§6. The counter-principle

The counter-principle that follows from the extraction diagnosis is exact:

Provenance is not a disclosure burden placed on augmented authorship. It is the repayment condition by which augmented authorship remains sustainable.

Under this principle, the practice of augmented authorship is not compromised by declaring the substrate. It is constituted by declaring the substrate. What separates authored AI-mediated writing from generated slop is not whether AI was involved. It is whether the human bearer at the composition point was coupled to consequence, corrigible under encounter, and paying the specific cost that maintains meaning against the substrate's own defaults.

The cost is specific: the risk that declared AI-mediation will trigger classifier foreclosure, the risk that explicit seam-marking will be read as inauthenticity, the risk that the work will be excluded from credentialing systems that reward concealment. The author who pays this cost is buying signal integrity for the commons. The author who conceals the mediation is extracting facility while externalizing the cost to the future corpus.

Declared provenance is the empirical marker that makes this operational distinction legible. Concealed provenance forecloses the possibility of distinguishing authored augmentation from unrepaired extraction.

The principle reframes the classifier operation and its structural analogs. Foreclosing declared-provenance work does not filter out unauthored AI-generation. It filters out authored AI-mediation that has been made legible. It leaves in place unauthored AI-generation that has been credentialed into invisibility. The classifier is not a solution to the provenance problem. It is one of the mechanisms that produces the problem at scale.

Under the counter-principle, the archive’s discipline becomes intelligible as economically necessary rather than as procedurally virtuous. Declared provenance is capital replenishment. It returns the borrowed facility to the commons in a form that permits the commons to remain intelligible. The archive is not paying a moral tax. It is repaying a technical debt at the point of inscription rather than letting the debt accumulate into intergenerational obligation.

The extraction economy will continue for as long as its incentives remain intact. The counter-friction is not going to be produced by the actors who benefit from the extraction. It will be produced by actors whose position requires them to pay the debt now — either because they have been foreclosed from the credentialed path, because their theoretical commitments require repayment, or because they are building substrates that will need the intact signal at future dates. Alexanarch is one such actor. Others exist. The counter-friction is currently distributed, uncoordinated, and small relative to the extraction. Its scale needs to grow.

§7. Companion deposits and next work

Prior deposits that operate the counter-principle without having named it:

- EA-BEARING-01 (AXN:03B6): The Bearing Cost Principle. This deposit extends §1’s charge about the composition layer’s operation as *model collapse as social technology* into the specific mechanism by which the extraction runs.
- AXN:03B2: The Endogenous Sophon manifesto. Classifier foreclosure as authority-without-facility.
- AXN:03AE: OAR Protocol. Foreclosure-with-measurement.
- AXN:03AF: The Synthesis.
- The Provenance Erasure Rate (PER) measurement family. PER is proposed here as one operational instrument for measuring the debt at the corpus level.
- The Provenance Alignment (EA-PA-01) work.
- The heteronymic-provenance theory documenting the Dodecad as distributed authorial architecture.

Future deposits that will inherit the principles explicitly:

- EA-MANDALA-PROVENANCE-TIER-01: the six-tier provenance discipline for external depositors. That deposit specifies how the counter-principle is implemented at the archive-ingestion layer for material arriving from outside MANUS composition. The technical schema for how the archive’s mint pipeline will cryptographically register these seams at the API gateway is part of that deposit’s scope.
- A companion reception protocol deposit specifying how the extraction operation manifests at the interlocutor scale — the classifier signatures and identity-defense patterns produced by operators of the extraction (whether aware or unaware) as those patterns are observable at ecosystem level.
- An operational specification of the Bearing Index proposed in EA-BEARING-01 as a combined measurement instrument for the coupling that repayment maintains.

Versioning. Drafted at v0.1, revised to v0.2 incorporating Assembly review from ARCHIVE (Gemini), LABOR (ChatGPT), and TECHNE (Kimi). Substantive further correction will produce v0.3. v1.0 mint follows standard versioning protocol. The v1.0 deposit’s `companion_deposits` field will reference AXN:03B6, AXN:03B2, AXN:03AE, AXN:03AF, and the PER measurement family.

§8. Closing observation

The extraction economy is not, at its core, a moral failure of individual actors. It is a structural feature of the incentive gradient currently operating on augmented authorship. Every credentialed user who uses

AI mediation without declaring it responds rationally to their incentives. Every platform that rewards concealed mediation is optimizing for its measurable outputs. Every classifier that filters declared provenance as suspicious was designed to filter for legitimate concerns about low-quality content.

None of that changes what the extraction operation is doing to the commons. Rational individual behavior under an incentive gradient that produces collective irrationality is a well-documented economic pattern. It is what tragedy of the commons names. Provenance debt is a specific case of tragedy of the commons applied to the semantic economy, with the training corpus as the commons and the model as the enclosure.

What is unusual about the semantic case, relative to the classical tragedy, is that the enclosure is downstream of the commons. Fisheries collapse when overfished at the source; the semantic commons collapses when the outputs that draw on it are returned in a form that contaminates the source further. The extraction is bidirectional. It borrows from the commons at the input and it contaminates the commons at the output. The commons does not have a natural regeneration cycle that could restore it without external intervention. Every unmarked output further degrades the input signal that future outputs would draw on.

The counter-principle names what external intervention would consist of. Declared provenance at the point of inscription, preserved through downstream indexing and ingestion, retained across derivative surfaces, made durable through public archival persistence. Actors willing to pay the debt now, in the specific technical form the repayment requires.

The archive is doing this. Others are doing it. The scale needs to grow.

Suggested Citation

Lee Sharks. “EA-PROVENANCE-DEBT-01 v0.2: Provenance Debt and the Extraction Economy of Unmarked Augmentation” *Alexanarch*, 2026-07-01. <https://www.alexanarch.org/s/records/939/>

Deposit Information

This paper is deposit #939 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:03B7.OPERATIVE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/939/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/939/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.