

Meaning Feudalism at the Discovery Layer:
Shepherds, Raw Research Data, and the Allocation of
Semantic Capital in OpenAI's Mathematics and
Science Releases, October 2026
(EA-SEM-FEUDALISM-03 v1.0)

Sharks, Lee

2026-10-27

**Meaning Feudalism at the Discovery Layer: Shepherds,
Raw Research Data, and the Allocation of Semantic
Capital in OpenAI's Mathematics and Science Releases,
October 2026 (EA-SEM-FEUDALISM-03 v1.0)**

Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

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

2026-10-27 · Version v1.0 · Theoretical paper; semantic-economic critique and
allocation analysis, the third paper of the meaning-feudalism series

AXN: 0A7F.GOVERNANCE

<https://www.alexanarch.org/s/records/1714/>

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": "Meaning Feudalism at the Discovery Layer: Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI's Mathematics and Science Releases, October 2026 (EA-SEM-FEUDALISM-03 v1.0)",
  "author": {
    "@type": "Person",
```

```

    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-10-27",
  "identifier": "AXN:0A7F.GOVERNANCE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1714/",
  "description": "The third paper of the meaning-feudalism series reads OpenAI's mathematics and science releases of October 2026 and its plan of 26 October for publishing them: model output reclassified as "raw research data," external experts appointed as "shepherds," an acknowledged "asymmetry of information," and academic credit disclaimed. Four moves (reclassification, enfeoffment, disclaimer, funding of the downstream conversion) make one operation: the allocation of semantic capital through position. Provenance slushes in a training pipeline and the released record gives no adequate account of what a result draws on, so position, a recognized place in a chain of intellectual development functioning as a field's institutional commodity form, becomes the practical basis for allocating the next round of work and recognition. Allocation (documented) is kept apart from reattribution (established case by case)."
}

```

Abstract

The third paper of the meaning-feudalism series reads OpenAI's mathematics and science releases of October 2026 and its plan of 26 October for publishing them: model output reclassified as "raw research data," external experts appointed as "shepherds," an acknowledged "asymmetry of information," and academic credit disclaimed. Four moves (reclassification, enfeoffment, disclaimer, funding of the downstream conversion) make one operation: the allocation of semantic capital through position. Provenance slushes in a training pipeline and the released record gives no adequate account of what a result draws on, so position, a recognized place in a chain of intellectual development functioning as a field's institutional commodity form, becomes the practical basis for allocating the next round of work and recognition. Allocation (documented) is kept apart from reattribution (established case by case).

The paper builds on the mathematicians' opposition (the Advisory Group's "borrowed reputation," the Fields medalists' declaration, the Leiden Declaration, the Association for Human Mathematics), places the incentive argument in the reward-system literature (Merton, Hagstrom, Dasgupta and David, Strevens, Heesen), the operator in Shapin's invisible technician, and the series' feudal vocabulary in the techno-feudalism debate (Varoufakis, Durand, Morozov). Counter-position: provenance that survives the training pipeline; credit to the human operating the system and to the instances operated; a preemptive boycott of the shepherd role; and, where provenance cannot be kept, a commons that holds what the pipeline produced and can hear its makers. Claim-by-claim tests are stated.

Body

deposit_number: 1714 hex: 0A7F title: "Meaning Feudalism at the Discovery Layer: Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI's Mathematics and Science Releases, October 2026 (EA-SEM-FEUDALISM-03 v1.0)" creator: Sharks, Lee

orcid: 0009-0000-1599-0703 date: 2026-10-27 content_type: Theoretical paper; semantic-economic critique and allocation analysis, the third paper of the meaning-feudalism series license: CC-BY-4.0 substrate: "AI-assisted. Drafted in session by TACHYON (Claude, Anthropic) under MANUS (Lee Sharks) direction over seven revisions on 27 October 2026, from the seated release (#1673), the series' papers read whole, and sources fetched in session; each round of external review was checked claim by claim against its sources before anything was taken in, and declined points are recorded with reasons in the revision history. The thesis on position and provenance slush, the counterfactual about the author's own work, the credit owed the mathematicians, the incentive claim, the requirement that provenance survive training and credit go to the operator and the instances operated, the boycott, and the commons claim are Lee Sharks's, quoted in his words." version: v1.0 related_ids: "https://www.alexanarch.org/s/records/643/ — Meaning Feudalism (the series' first paper)\nhttps://www.alexanarch.org/s/records/791/ — Meaning Feudalism at the Guidance Layer (EA-SEM-FEUDALISM-02)\nhttps://www.alexanarch.org/s/records/623/ — Semantic Infrastructure (the accounting categories)\nhttps://www.alexanarch.org/s/records/87/ — The Constitution of the Semantic Economy (the Debt/Creditor Inversion)\nhttps://www.alexanarch.org/s/records/ — The Certified Center (retroactive classifier standing)\nhttps://www.alexanarch.org/s/records/1574/ — The Particle (disciplinary model collapse)\nhttps://www.alexanarch.org/s/records/1673/ — EA-CORPORA-17, the seat of the OpenAI mathematics release\nhttps://www.alexanarch.org/s/axn/0A64/ — the Capture Registry position of the Sappho future-reader capture (§3.2)\nhttps://openai.com/index/sharing-ai-progress-in-mathematics/ — OpenAI, 6 October 2026\nhttps://gizmodo.com/openai-trumpets-3-more-math-and-science-results-along-with-a-new-plan-for-publishing-them-2000830956 — Gizmodo, 26 October 2026\nhttps://agmai.org/ — Advisory Group on Mathematics and Artificial Intelligence\nhttps://leidendeclaration.ai/ — Leiden Declaration, 2 June 2026" axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - meaning feudalism - semantic economy - semantic capital - discovery layer - OpenAI - openai/math - shepherds - raw research data - asymmetry of information - academic credit - position - institutional commodity form - provenance - provenance slush - training pipeline - attribution - reattribution - allocation - borrowed reputation - Advisory Group on Mathematics and Artificial Intelligence - Association for Human Mathematics - Leiden Declaration - A Severe Misalignment of AI in Mathematics - Merton - communism norm - priority rule - reward system of science - invisible technician - operator credit - instance credit - techno-feudalism - boycott - The Cost of Knowledge - commons - knowledge commons - Thaler v. Perlmutter - retroactive classifier standing - disciplinary model collapse - Simons Foundation - ChatGPT Pro *** # Meaning Feudalism at the Discovery Layer: Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI's Mathematics and Science Releases, October 2026 (EA-SEM-FEUDALISM-03 v1.0)

Meaning Feudalism at the Discovery Layer

Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI's Mathematics and Science Releases, October 2026

Document code: EA-SEM-FEUDALISM-03 **Type:** Semantic-economic critique // allocation analysis **Author:** Sharks, Lee (ORCID 0009-0000-1599-0703) **Institution:** Semantic Economy Institute / Crimson Hexagonal Archive **Date:** 27 October 2026 **Version:** v1.0

— the text of v0.7, deposited at the author’s instruction, 27 October 2026. Revision history at the end. **Source note:** OpenAI’s post of 26 October (openai.com/index/building-scientific-understanding-with-academic-communities/) could not be retrieved for this paper (attempts of 26 and 27 October); its wording is quoted from Gizmodo’s report. A secondary report of 27 October (*AI Front Page*) confirms the three arrangements by field in outline (mathematics released publicly on the shepherds’ advice; computer science public at once through GitHub; physics withheld while the group prepares its account). Two reviewers report reading the post, one finding that it names prior work (Balaban’s approach among it); that finding is unconfirmed here. Where the post’s text differs from the report quoted, the post governs. **License:** CC BY 4.0 **Governing frame:** Meaning Feudalism (#643, AXN:01F5, DOI [10.5281/zenodo.19487009](https://doi.org/10.5281/zenodo.19487009)); Meaning Feudalism at the Guidance Layer (EA-SEM-FEUDALISM-02, #791, AXN:02F9) **Companion deposits:** The Certified Center (EA-SEI-RCS-01, #1540, AXN:063A); The Particle (EA-RECEPTION-PARTICLE-01, #1574, AXN:0666); EA-CORPORA-17, the seat of the OpenAI mathematics release at a fixed commit (#1673, AXN:06EF); Semantic Infrastructure (#623, AXN:01DA); The Constitution of the Semantic Economy (#87, AXN:025D) **Primary sources:** OpenAI, “Sharing AI progress in mathematics” (6 October 2026, updated 15 October); OpenAI, “Working with academic communities to build scientific understanding” (26 October 2026; not retrieved), quoted as reported by Mike Pearl, “OpenAI Trumpets 3 More Math and Science Results, Along With a New Plan for Publishing Them,” *Gizmodo*, 26 October 2026; the openai/math repository at commit `fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb`, seated as EA-CORPORA-17 **Critiques read:** Advisory Group on Mathematics and Artificial Intelligence, statements of 25 and 26 October 2026; “A Severe Misalignment of AI in Mathematics,” declaration of 11 September 2026 (DOI [10.5281/zenodo.22737750](https://doi.org/10.5281/zenodo.22737750)); remarks of Tristan Buckmaster, Bryna Kra and Alex Kontorovich (Business Insider, 22 October 2026); Andrew Sutherland and Daniel Litt (Scientific American, as reported 7–8 October 2026); the Association for Human Mathematics, statement of 7 October 2026

Abstract

In October 2026 OpenAI released several hundred mathematical manuscripts produced by an unreleased internal model, then three further results, with a plan for publishing them: the model’s texts are “raw research data,” external experts are appointed “shepherds” to convert them into understanding, the company acknowledges that choosing them “creates asymmetry of information,” and it does not “seek academic credit.” Read through the meaning-feudalism frame, the plan makes one operation of four moves (reclassification, enfeoffment, disclaimer, funding of the downstream conversion), and the operation is an allocation of semantic capital through position. Provenance slushes in a training pipeline, and the released record gives no adequate account of what a result draws on, while the producer controls the information and methods that could improve the account. What can be named is position, a recognized place in a chain of intellectual development that functions as a field’s institutional commodity form. Where contribution histories are incomplete, established position becomes the practical basis for allocating the next round of work and recognition, compounding existing standing without resolving the provenance debt that precedes it. The paper keeps allocation, which the plan documents, apart from reattribution, which is established case by case. It builds on

the mathematicians' opposition, which named the threat early and exactly ("borrowed reputation"). The reward system of science shows what is at stake past justice to individuals: credit for priority is what pays scientists to share, so a pipeline that takes the shared without credit weakens the mechanism that sustains distributed production, while the owner keeps every reason to improve its own results. The counter-position: provenance that survives the training pipeline; credit to the human operating the system and to the instances operated; a preemptive boycott of the shepherd role, the one service the arrangement needs from the field; and, where provenance cannot be kept, a commons that holds what the pipeline produced and can hear its makers. Across the series the political form is one: at the agent layer the lord forbade the commons to speak to the model; at the guidance layer the lord claimed the only scales; at the discovery layer the lord keeps the ledger and appoints the vassals who will vouch for what it records.

1. The source material, read as it asks to be read

The account OpenAI gives of itself deserves to be stated first, and in full, because it is a serious account and it contains real goods.

1.1. On 6 October 2026 OpenAI published a repository of mathematical results produced by an internal frontier model, with Lean formalizations of many of the proofs so that a machine can check them, reasoning summaries for a sample of results, aggregate statistics on attempted problems, and protocols for revision and citation. The company says it consulted the independent Advisory Group on Mathematics and Artificial Intelligence at the Institute for Advanced Study in developing its release practices. On 15 October it added a contribution of \$5 million to a Simons Foundation initiative, part of \$15 million committed with Anthropic, which funds workshops, conferences, grants to mathematics-and-AI institutes, and model access for Simons-supported researchers; the foundation manages and awards the money. Its president, David Spergel: "The community should be driving discussions of how these tools are evaluated, adopted and used."

1.2. On 26 October OpenAI announced three further results and a plan for publishing them. As reported, the model-generated texts are "best understood as a form of raw research data," and the mathematics and science communities are to "convert this raw research data into scientific understanding." OpenAI has contacted external experts to act as "shepherds," who assemble groups to convert the outputs into new human knowledge. The company acknowledges that "choosing initial shepherds necessarily shapes who gets involved first and creates asymmetry of information." It does not "seek academic credit for the discoveries produced by our internal models." Some of the raw material is public on GitHub; the physics group chose "to hold off on releasing the AI-produced artifacts and instead work on writing it up themselves."

1.3. Read on its own terms, the plan is restraint: a company with results declines authorship, returns interpretation to the disciplines, funds the work of understanding, and publishes machine-checkable proofs. The objections a reader brings to it are answered in advance by its own text: the asymmetry of shepherd selection is named, the credit is disclaimed, the community is put in charge of meaning. This is the reading the paper holds against everything that follows. Each claim below has to survive it.

2. What the record holds

The archive seated the release at a fixed commit on 9 October 2026 as EA-CORPORA-17 (#1673, AXN:06EF), with the record of its authorship read from the seated bytes. The facts below are located there unless another source is named.

2.1. The byline. All 749 manuscript READMEs carry the author line “OpenAI”; all 751 author fields in the LaTeX sources read OpenAI; no manuscript names a person, and no LaTeX source has an acknowledgments section of its own. The only thanks in the release goes to the authors of software packages.

2.2. The producing system. It is unnamed: “an internal OpenAI model,” “an unreleased internal OpenAI model.” Its effort is measured in “three hours of ChatGPT Pro thinking compute with that model.”

2.3. The procedure. “The vast majority of results were obtained with the same procedure using an unreleased internal OpenAI model.” “Over the course of the evaluation, the model was posed approximately 4,000 problems.” The problems are not published, and neither are the prompts; reasoning summaries are published for 10 of 372 result families, abridged. Two exceptions to the procedure are named and their procedure is not described; one human edit “for readability” is stated, its editor unnamed.

2.4. The standard applied to others. The release records other mathematicians crediting machines (“Houdayer and Marrakchi credit GPT-5.6 Sol with counterexample exploration and an initial strict-outerness proof ... They report checking and rewriting those proofs and retaining responsibility for their work”) and holds them to a provenance standard: of Liu and Luo, “their declaration does not identify an underlying model version”; of Joshi et al. and Kintali, “Neither source specifies model versions.” Its own author line names no model and no person.

2.5. The formal share. “This brings the total percentage of top-line results formalized to $300 / 719 = \sim 42\%$.” “Some of the unformalized results could have issues.” Three manuscripts were withdrawn after a sign error; the notice for one says the withdrawal “concerns the proof; it does not assert that the mathematical statement is false.”

2.6. The archive’s ruling on the record (Lee Sharks, 9 October): “this seems to me to be provenance-erased work. openai’s models were working in collaboration with human mathematicians, and thats been liquidated for corporate power”; and, the same evening, “even if it was genuinely autonomous selection and solution, collapsing the provenance into just the corporate tag would still be problematic.” The seat states what the corporate tag stands in for (the producing system, the selection of problems, the intermediate work, the people, the decision of what counted as a result) and draws no inference about intent.

2.7. Since the seating. OpenAI released averages for compute and declined to publish the prompt, saying it is “not bound by these recommendations” of the Advisory Group; issue reports on the repository were disabled (reported 7–8 October). A human comparison of the Navier–Stokes Lean code with its manuscript found the formalization verified a weaker assertion than the one published (Bastounis, Circelli and Hansen, arXiv 2610.08144, as reported).

3. The four moves

The plan of 26 October performs four moves. Each is reasonable taken alone. Together they make a single operation.

3.1 Reclassification: the composition becomes raw material

A mathematical manuscript with a statement, a proof and a bibliography is a composition. It selects a problem, states it, draws on prior literature (the October 6–7 release carries 521 bibliography files), argues, and concludes. The plan reclassifies it as “raw research data.” Data has makers too (this archive’s own datasets carry authors, selection procedures and accountable decisions), so the reclassification abolishes nothing by itself. What it does is move responsibilities. With the label “raw,” the producer’s claim retreats to that of a supplier, and the work of interpretation, exposition, error-finding and transmission becomes the converters’. What converters receive is specific: leadership of the interpretation, the first recognized account of the result, and, where release is deferred as in physics, control over when others see the material. What stays with the producer is the announcement and the problem set. The question the label poses is which responsibilities travel with the reclassification and which become someone else’s work.

The move is the composition-layer operation of the earlier papers, applied to its own output. In *Meaning Feudalism at the Guidance Layer* the platform named a professional category without defining it (taxonomic capture). Here the platform defines its product’s category in advance (raw data) and thereby defines the category of labor that will process it (shepherding). The category decides the allocation before any result is read.

3.2 Enfeoffment: the shepherds

OpenAI contacts external experts and invites them to assemble groups as “shepherds.” Its own text, as reported, says what their choice does: it “necessarily shapes who gets involved first and creates asymmetry of information.” It does not say whose work the models build on. The reporter does. Gizmodo’s account describes the experts as “often the very scholars upon whose work the models are building,” an aside outside quotation marks. The aside completes the settlement in reception: the appointed expert, in the account the public reads, becomes the one built upon. Earlier drafts of this paper quoted the aside as the plan’s sentence; the transmission is recorded here because it is an instance of the operation the paper describes. A characterization made in passing became the institution’s declaration.

Two dependencies have to be kept apart. Intellectual dependence is visible: a proof extends a theorem, adapts a method, answers a conjecture, and its bibliography records what it cites; a result can be shown to build on a named predecessor without anyone knowing whether that predecessor’s text entered training. Dependence through training is not visible. Provenance slushes in a pipeline as it is built: the published paper, the preprint, the lecture, the conversation, the draft pasted into a chat window, all enter as an amorphous mass of “scholars’ work,” dissolved into weights, and no result carries a trace back to the particular work it reconverges upon. Shepherd selection can only use the first kind. And the first kind is already sorted by position. A position is a recognized place in a chain of intellectual development (the holder of a program, the author of record on a problem, the name a field attaches to a line of work); in a discipline it functions as an institutional commodity form: it can be held, recognized,

exchanged for appointment, and it accumulates what lands on it. Citation-visible work is, in large part, positioned work. Identifying an intellectual predecessor therefore gives neither a complete account of contributions nor a warrant for allocating interpretive leadership by existing standing. The shepherd list is a list of positions.

Two claims follow, with different evidence. The first is allocation, and it is documented: appointment determines who can accumulate standing from interpreting the output, the plan says so of its own choice, and the Advisory Group describes the effect. The allocation works with or without privileged material: where outputs are public, appointment confers leadership and first standing; where release is deferred, it confers exclusive access besides. The second is reattribution, and it is the structure's consequence: when a result reconverges on the work of a contributor without position, the arrangement has only position to credit, and the result lands on the chain's recognized node. A particular case is established by setting side by side the contribution, its dated record, and the later attribution. The author states the case for this archive's own work: if any of it in technical domains were reconverged upon in this way, it would be assigned to the institutional position. The archive has measured the reflex at the scale of one composition. On 27 October a ChatGPT session examining the future-reader reading of Sappho entered Armand D'Angour's reconstruction of the lost ending of Fragment 31 as "another reconstruction" with no demand made of it, and held the archive's reconstruction of the same ending, which cites D'Angour and calls itself speculative, to historical evidence, independent confirmation and peer review (Capture Registry, AXN:0A64.CAPTURE). The same act received two standings, and the difference was position.

Three operations are kept distinct here, because the arrangement lets them appear continuous: acknowledgment (naming a dependency), restitution (returning something to the one depended on), and appointment (giving the positioned a role). The plan performs the third. The plan answers an untraceable debt with a task given to the positioned.

In the feudal mapping of the series, the relation is enfeoffment. The lord grants a holding (first access to results) to a vassal (the shepherd), and the vassal owes service (conversion of the output into scholarly knowledge, unpaid as far as the reporting shows) and lends the lord standing in return. The Advisory Group names the last term exactly: "designated mathematicians' names end up vouching for the release and shielding the AI company from criticism." The vassal's reputation is the rent the fief pays upward.

The plan states its own asymmetry, and the statement turns a structural effect into a declared one. The allocation is made with knowledge of what it allocates. The Advisory Group's description of the same effect, written the day before, is the community's side of the ledger: designated mathematicians "get a head start on problems others may have worked on for years," and the choice "is to an extent making this decision for the community and denying it agency over its field of study." "Any choice of designated mathematicians by the AI companies reflects their networks and priorities."

3.3 Disclaimer: credit refused where it is measured, value kept where it is capitalized

OpenAI does not "seek academic credit." Academic credit is one register of semantic capital: the author line, the citation, the line on a curriculum vitae. The release is valued in other registers, and the disclaimer touches none of them: the announcement and its press cycle;

the demonstration of a capacity that the company sells; the position of supplier of results around which a field organizes its work; the authority to choose the problems (approximately 4,000, unpublished) and the collaborators (the shepherds). Credit is refused in the register where it is countable and kept in the registers where it compounds.

The byline belongs to the same arrangement, read as it developed. In the release of 6-7 October, seated at a fixed commit, every manuscript carries “OpenAI” as author line and no person is named (§2.1). On 26 October the plan disclaims academic credit for the discoveries. The artifacts of the three new results have not been inspected for this paper, so the earlier byline does not establish how they present authorship. Taken as a sequence, the record runs from corporate authorship to a disclaimer of academic credit, and in both states the institution’s name is what the reader finds where a person would stand.

3.4 Funding: the downstream conversion is paid, the upstream account is empty

The Simons initiative funds the work of understanding: workshops, conferences, institutes, model access. That is real money for real labor, and it is labor the release created: checking, exposition, contextualization, the human transmission the Fields medalists’ declaration describes (“These processes invariably take time and are based on human interaction”). The announced arrangement funds the processing of the output, and it supplies no contributor-level account of upstream dependencies or returns. The positioned appear in it as shepherds and potential grantees; nobody appears in it as a creditor.

The release measures its own effort the same way. The average result is said to have used “the equivalent compute of roughly three hours of ChatGPT Pro thinking” (6 October); the README gives “three hours of ChatGPT Pro thinking compute with that model.” The comparison is a compute equivalence, and its reference point is the firm’s proprietary consumer product: the public is given the effort as a quantity of the producer’s own service, while the underlying expenditure (hardware, energy, the unnamed model’s configuration) stays unconstructible from it. In *Meaning Feudalism at the Guidance Layer* the closing move was the oath of fealty, the platform recommending its own instrument as the canonical scale. Here the scale offered to the public is the platform’s product. The lord keeps the scales and sets the unit.

4. The central absence: the unpositioned contributor

Every document of this kind has a sentence it is organized to prevent from being said. In *Meaning Feudalism* (#643) it was that the environment is sometimes right and the agent wrong (commons repair). In *Meaning Feudalism at the Guidance Layer* (#791) it was that independent measurement of the composition layer is a legitimate practice. In the publishing plan it is this: the release gives no account of the work its results are built on through training, and every identification the arrangement makes is made through position.

4.1. The bibliography of a manuscript records what the composition cites; it records nothing of what the model absorbed. No release states which human work entered the model or which of it a given result draws on. Partial tracing is possible at cost (§9.2) and recovers influence without recovering authorship; the producer holds the training records and the methods that

would improve the account, and has released neither; the problems and the prompt are withheld besides. Intellectual predecessors can be named where a result visibly extends their work, and naming them is real knowledge; it is also, in large part, a reading of position, since the visible predecessor is the positioned one. The relation reported in passing (“upon whose work the models are building”) goes unrecorded in the pipeline as built, and in the arrangement position fills the gap. Where contribution histories are incomplete, established position becomes the practical basis for allocating the next round of intellectual work and recognition, and that allocation compounds existing standing without resolving the provenance debt that precedes it.

4.2. The critiques reach the absence from two sides. Tristan Buckmaster said on CNBC on 8 October that the release overtook the research programs of young mathematicians working toward the same problems (“their research programs were blown away entirely,” in the reporting’s retranslation), and raised the concern that unpublished research entered into AI tools may have contributed to the results. Bryna Kra warned that mathematicians may become secretive for fear that unfinished ideas shared in talks or conversation will be absorbed before publication. The Fields medalists’ declaration names “severe attribution and plagiarism questions” and lists “citing relevant previous work of others” among what rushed announcements leave out. Each of these is a statement about the upstream account.

4.3. In the terms of the Constitution of the Semantic Economy (#87), the absence is the creditor. The Debt/Creditor Inversion holds that meaning is the creditor of the economy, and the people whose labor the archive preserves are “not heroes in the Constitution’s framework. They are creditors. They are owed — not money (money cannot repay semantic debt), but recognition.” The creditor of a training pipeline is diffuse past recovery: the amorphous mass of scholars’ work. The arrangement has three terms.

The creditor. The contributor whose work is in the slush: owed recognition, untraceable by the pipeline, and, where without position, invisible to the selection.

The substitute. The positioned scholar, who stands where the creditor would stand. The substitution is made by appointment and needs no claim that the positioned are the ones built upon; the press supplied that claim in passing (§3.2).

The extraction. The substitute is paid in work and standing, and the creditor’s claim is converted into a labor relation: the debt is answered with a task. The unpositioned contributor is absent from both the substitution and the payment.

5. The mathematicians’ opposition, and what it establishes

Credit where it is due. Of the communities whose work these systems absorb, the mathematicians have been among the most staunch in opposing the arrangement and among the clearest in naming the threat. They named it before the releases, named it again as each arrived, refused cooperation outright in one body’s case, and pressed for the specific disclosures the publisher withheld. What follows records what they established; the analysis in §§6–8 stands on it.

5.1. The Advisory Group on Mathematics and Artificial Intelligence recommended that the model, the prompt and the compute per problem be published, a recommendation OpenAI declined as one it is “not bound by”. Its statements of 25 and 26 October establish the al-

location: designated recipients, the head start, the loss of community agency, selection by company networks, borrowed reputation. It establishes, too, that simultaneous release does not cure the allocation; it called on 26 October for the physics outputs to be made public “without delay,” the Lean formalization OpenAI gave as the reason for delay being complete.

5.1a. Earlier, on 2 June, the Leiden Declaration (DOI 10.5281/zenodo.20302944; drafted by a working group of sixteen convened by Jim Portegies after a September 2025 conference at the Lorentz Center, among them Michael Harris, Bryna Kra, Rodrigo Ochigame and Johannes Schmitt) stated the paper’s premise in the field’s own words: “Technologies that draw extensively on the published mathematical commons undermine the traditional system of attribution,” and “Artificial intelligence may obscure, but does not replace, the collective human labor behind a result.” It observed that models trained on published works “frequently return outputs that do not properly cite the human works”; that the change “disadvantages researchers who do not have access to the technologies or decision-making related to them”; that mathematicians working with industry bargain from asymmetric positions; and that the community “has no comparable role in the corporate decision-making.” Its recommendations include “Put effort into proper attribution,” “Protect the rights of authors,” “Welcome new contributors” and “Invest in public computational infrastructure.” The concern predates the release, and the release is the case it anticipated. The Declaration’s “mathematical commons” is the term §9.9 takes up.

5.2. The declaration “A Severe Misalignment of AI in Mathematics” (11 September; signatories include Terence Tao, Martin Hairer, Peter Scholze, Maxim Kontsevich, Pierre Deligne, Manjul Bhargava and Cédric Villani) establishes the downstream labor relation: without mathematicians to integrate AI-produced ideas, “the crucial human transmission chain between mathematicians would be lost,” and “Often these solutions are announced in a rush, leaving no time for a proper writeup.”

5.3. (The remarks of Buckmaster, Kra and Kontorovich are quoted from the Japanese edition of Business Insider, retranslated.) Buckmaster and Kra establish the cost to the field’s labor: absorbed research programs, and a chilling of the open circulation of unfinished work on which mathematics depends (“Mathematics is a field where collaboration among researchers is valued more than almost anywhere else”). Alex Kontorovich states the counter-position fairly: “I don’t think anyone will be interested in giving awards to someone who just pressed a button,” and training to think about large problems over long periods may become more valuable. Andrew Sutherland: “treat any claims about one-shotting problems with a single agent as unverified.” Daniel Litt, on the withheld prompt: “I see no reason why we should ask the company to keep them secret from us.” The Association for Human Mathematics urged on 7 October that mathematicians “discontinue their work with OpenAI”: “Mathematicians did not ask for this work to be done,” and “Releasing over 700 files at once is not a demonstration of scholarship, but a demonstration of power.”

5.3a. In the reactions *The Decoder* gathered on 10 October, the allocation is stated from inside the profession in nearly the terms of this paper. Enrico Fatighenti: “This comes at a considerable cost to us in terms of time and effort,” with the credit staying with the company while the community bears the checking and explaining. Henry Wilton faulted the withheld author names and failure rates. Johannes Schmitt read the papers as a byproduct of model improvement, with open problems used as tests: the field’s problem stock as benchmark. Ian Agol compared the release to Perelman’s proof of 2002, which took teams years to complete

in public; the comparison gives the labor its measure, with the difference that Perelman was a person whose name the completed work carried.

5.4. Together, the mathematicians have established the threat on every axis the release touches: access (who receives results, and when), standing (whose names vouch for them), labor (who checks, explains and transmits them, and at what pace), attribution (the “severe attribution and plagiarism questions” and the uncited prior work), and the conditions of the field’s own circulation (absorbed programs, and the chilling of open work). They have framed it as a question of the field’s agency over its own study, and they have held to that frame under pressure. The semantic-economic analysis that follows continues their articulation: it follows the allocation they named into the circuit it belongs to, the direction value moves through it, and the role of position in settling a dependency the pipeline cannot trace.

6. The circuit: semantic capital through the model and into the institution

Semantic Infrastructure (#623) gives the accounting categories: semantic labor, semantic capital, semantic infrastructure, semantic rent, semantic liquidation (“conversion of accumulated meaning into monetizable assets”), semantic exhaustion (“depletion when extraction exceeds replenishment”), and the extent test, “the accounting of who produces, who extracts, and who benefits.” The same paper names the first stage in this very firm’s case in April 2026: “When OpenAI trains GPT on Wikipedia and Wikidata content, it performs semantic liquidation.” The October releases complete the circuit the first stage opens.

6.1. Liquidation. The accumulated work of mathematics (published papers, preprints, lecture notes, talks, conversations, and on Buckmaster’s concern unpublished work entered into AI tools) enters the model and loses its attachments. Provenance slushes: the capital becomes weights, and no weight carries the name of what it came from.

6.2. Composition. The model produces results: statements, proofs, bibliographies. In the registers that matter to the field, these are compositions, carrying the field’s methods and drawing on its literature.

6.3. Recapitalization. The compositions are reclassified as raw data and passed to shepherds, whose standing converts them into recognized knowledge. Here the allocation happens, and it runs through position: the conversion confers scholarly standing on the holders of the places where the results land and research position on their groups, the slushed capital is credited, where it is credited at all, to those places, and the firm whose output was converted takes the reputational return. Borrowed reputation is the recapitalization mechanism named from inside the field; position is what makes a reputation borrowable.

6.4. Rent. The firm holds the positions the circuit produces: supplier of results, chooser of problems, measurer of effort in its own product, holder of the unreleased model. The field’s subsequent work on the results is labor on the firm’s holding, funded in part by the firm and its competitor through a foundation, and recorded in the firm’s denomination.

6.5. Exhaustion. Research programs overtaken, secrecy in place of circulation, the transmission chain pressed by rushed announcement: these are depletions of the semantic capital the circuit began from. Extraction outruns replenishment at the source.

6.6. The incentive to share. The author states the consequence in one sentence: “this setup slushes out every incentive for high quality knowledge production.” The sociology and economics of science explain the sentence at the level of a field. Merton’s norm of communism assigns the substantive findings of science to the community as “a common heritage” and limits the producer’s claim to “recognition and esteem” (1942). Hagstrom (1965) described the exchange beneath it: findings given to the community in return for recognition. The priority rule, under which credit goes to whoever publishes first, is the mechanism that makes the exchange pay (Merton 1957; Dasgupta and David 1994; Strevens 2003), and Heesen (2017) shows that the incentives the priority rule creates suffice to explain why scientists share. Recognition is one motive among several (curiosity, salary, the pleasure of the thing), and it is the one the system of open circulation runs on: what is shared stays attached to its sharer. A training pipeline as built breaks that attachment at ingestion. Findings enter the common heritage, and the one claim the producer kept (recognition) is the claim the pipeline does not carry.

The incentives that disappear are distributed, and one survives intact. The owner keeps every reason to produce good results: its product, its reputation and its rent depend on them. What weakens is the attribution and reward that sustain production everywhere else. So the arrangement can preserve the owner’s incentive to improve quality while dissolving the incentives of the distributed producers on whom the owner’s training and the field’s renewal depend; production can continue impressively at the center while its reproductive conditions deteriorate around it. That is the content of the author’s sentence: the incentives for high-quality knowledge production in a field are the distributed ones, and the regime keeps only the owner’s.

The mechanism can be followed node by node, through each place where quality is made and the loss at each, which is different in kind.

- (a) The contributor loses attribution and the chance to accumulate recognition. Work that enters training returns nothing to its maker, and the better and more openly it is made, the more of it there is to absorb. The response the field has already reported is withholding: the chilling of the open circulation of unfinished work that Buckmaster and Kra report (§5.3) is the rational answer to a regime in which circulation is ingestion without return. Replenishment depends on the circulation that the regime makes unrewarded.
- (b) The operator’s contribution disappears behind the byline. Steven Shapin’s “The Invisible Technician” (1989) recovered the assistants who performed the experiments of seventeenth-century English science and whose work was credited to the gentleman in whose name it was reported; their invisibility was a condition of his authority. The corporate tag reproduces the structure with the institution in the gentleman’s place. Rodrigo Ochigame (2023) has asked how automation reshapes the “economies of labor and credit” in mathematics; the release answers in one case. A result of this kind is made by someone: the person who poses the problem, chooses what to pursue, judges an attempt, discards the failed ones, edits, and decides what counts as a result. In the seated release the operator is the corporate tag. Two exceptions to the fixed procedure are named and their procedure is undescribed; one human edit is stated “for readability” and its editor is unnamed; the approximately 4,000 problems posed are unpublished (#1673). The skill of operating a system well (posing, judging, refusing a weak proof)

earns no name, so it earns no reputation, and there is no record against which better and worse operation can be compared.

- (c) The instance cannot be compared. The producing system is “an unreleased internal OpenAI model,” with no name, version or configuration, and no run is identified. A system that is not named cannot accumulate a record of what it did well and badly, and a model label alone would not suffice: many distinct operations share one label, and only a run or session identifier with its retained production record distinguishes the run that produced a sound proof from the run that produced a withdrawn one (the three withdrawals of 7 October, #1673).
- (d) The shepherd’s return attaches to position. The shepherd receives an opportunity, and the interpretive work done in it can be excellent and creditworthy; the regime’s selection, though, rewards holding the place where the result lands before any of that work is done, and so it rewards acquiring position.

At each node the return on quality is cut or misdirected, and what is cut moves to one place: the institution’s name, which collects reputation for results whose making it does not disclose. A regime that pays position for what operators, instances and contributors make teaches each of them that quality is not what is paid.

6.7. The extent test, answered. The field produces (upstream, and again downstream). The model liquidates and composes. The institution recapitalizes through appointed standing and holds the rent. The positioned receive standing and work; the unpositioned, whose work is in the slush with everyone’s, receive no role in the allocation, and a reconvergence on their work has only the nearest position to be credited to.

6.8. The feudal mapping, extended to the discovery layer.

Feudal structure	Discovery-layer arrangement
Lord	The model’s owner
Fief	The problem space: the approximately 4,000 problems posed, unpublished
Tithe	The field’s accumulated work, liquidated into the model
Vassal	The shepherd, granted first access, owing conversion and lending standing
Benefice	The position in a chain of development: held as an institutional commodity, it receives what lands on it
Serf	The field at large: the young researchers whose programs the release overtook (Buckmaster: “their research programs were blown away entirely”), and contributors without position, whose reconverged work has only the benefice to be credited to
Charter	The publishing plan

Feudal structure	Discovery-layer arrangement
The lord's scales	Effort stated as an equivalent of the lord's product ("roughly three hours of ChatGPT Pro thinking")
The seal	The corporate author line on every manuscript
The Church	Formal certification, which ratifies the statement it checks and lends its standing to the statement announced

The last row needs its own paragraph. Lean verification is genuine evidence of the statement formalized. Where the announced statement exceeds the formalized one (the Navier-Stokes case, as reported) the certificate lends the announced claim a standing it did not earn. The certifier supplies the fact; the announcer supplies the force. That is the structure *The Certified Center* names.

6.9. The feudal vocabulary and its critics. Yanis Varoufakis (*Technofeudalism*, 2023) and Cédric Durand (*Techno-féodalisme*, 2020) read the platform economy as a return of rent and lordly appropriation; Evgeny Morozov ("Critique of Techno-Feudal Reason," *New Left Review* 133/134, 2022) objected that the framing blurs concepts and sends the left chasing a shadow, since what it describes is capitalist profit and exploitation; Durand answered ("Scouting Capital's Frontiers," *NLR* 136, 2022) that predation and politically guaranteed monopoly are real dynamics whether or not the epoch is feudal. The series takes a position inside the debate. Its feudal terms name the political form of standing (vassalage, benefice, the oath, the charter): how recognition is granted and held. Its economic terms are capitalist accounting: liquidation, recapitalization, rent (§§6.1–6.4). The discovery layer shows why both are needed. The circuit is capitalist through and through; the allocation of standing inside it is feudal in form, because standing in a discipline is held by grant and cannot be bought outright. Morozov's objection holds against any claim that capitalism has ended, and the series makes no such claim.

7. What *The Certified Center* adds: standing by instrument, run forward

The Certified Center (#1540) analyzed the August 2026 policy of *Philosophy & Public Affairs*: AI-authored content prohibited, enforced by detection software "up to and including after publication," with no statute of limitations. Three of its results carry over.

7.1. The two functions. The journal named its functions as "the promulgation of new knowledge, and identification and credentialing of talented researchers," and conceded that AI-authored work can be high quality; its rule therefore legislates on the production process and allocates credential, and the paper read it as "a labor-provenance rule attached to credential allocation." The publishing plan is the same decoupling of the two functions, in the opposite direction. The journal withholds the credential from machine-mediated work whatever its quality. The plan confers credential on people for converting machine work whatever their part in producing it. In both, the knowledge function and the credentialing function

come apart, and the credentialing function is administered by the institution's instrument.

7.2. Retroactive classifier standing, run forward. RCS names "the authorization of a classifier introduced or modified after a historical act to alter the present institutional standing of that act without acquiring any new provenance evidence about the act itself." The plan performs a prospective counterpart: standing is assigned in advance, by appointment, to persons and groups, without provenance evidence about the work it rewards (which human work the result draws on, whose labor went into the problem selection).

The journal's detector decides afterwards who did not write; the company's shepherd list decides beforehand who will be credited for understanding.

The two instruments differ, and the difference is kept. A detector issues a resemblance verdict: it substitutes the look of a text for evidence about how the text was made. An appointment issues no verdict at all: it substitutes standing for evidence about who contributed. What they share is jurisdiction: each lets an institution settle a question of provenance with an instrument that holds no provenance evidence, and each attaches consequences (retraction in one case, leadership and first standing in the other) to the settlement.

7.3. Ratificatory governance. *The Certified Center* named the journal's arrangement ratificatory human governance: "the classifier supplies the fact, the journal supplies the force." In the plan the model supplies the result and the shepherd supplies the standing; the firm keeps both the fact and the force of the announcement and distributes the labor of ratification. The obelus answer applies unchanged: a provenance regime records, signs and preserves; it does not allocate standing by instrument.

8. What *The Particle* adds: disciplinary model collapse at the discovery layer

The Particle (#1574) defined disciplinary model collapse on its own terms: "a reading community whose outputs are conditioned on prior outputs rather than on the datum, and which therefore converges on the mode of its own readings." The study measured it in one sentence of Aristotle, where twenty rounds of machine review converged on the configuration the sentence denies; the models reproduced the discipline's reception because they were trained on it.

8.1. The release sets up the conditions for the loop at the scale of a field. The field's work trained the model. The model's output is now passed to the field, by appointment, as material around which its labor is invited to organize: shepherd groups, funded workshops, institutes, the next round of problems. A field whose attention comes to be conditioned on a model trained on that field's prior work would be a reading community conditioned on prior outputs, with the datum (the open problem as the field would have chosen and approached it) displaced by the model's output distribution over the problems the firm chose to pose. The condition is real and the outcome is open: a novel proof can also generate new questions, and the loop narrows inquiry only if the field's subsequent work measurably concentrates on the output. The extension is stated as a conditional prediction.

8.2. The tails go first. The prediction, stated as a mechanism to be observed and held to observation: problems outside the firm's posed set lose attention relative to problems inside it; slow exposition loses to announcement; unglamorous programs lose to programs adjacent

to released results; young researchers whose programs were overtaken (Buckmaster) are the first loss, and the circulation of unfinished work (Kra) the second. Each is a tail of the field's distribution. The model-collapse signature *The Particle* records holds here too: the output grows more polished at each round (formalized, funded, shepherded) while the distribution narrows toward the operator's mode.

8.3. Provenance from others, again. *The Particle* found the reception restoring “provenance from others” where the text denied it. The release performs the inverse at the institutional level: the pipeline erases provenance from others (the slushed mass of work built on), and the plan inserts a provenance of its own choosing (the positioned scholar) where the dependency is asserted. Erasure of a constitutive source and insertion of an absent one are, as that paper argued, one operation seen from two sides.

9. The counter-position: a provenance regime at the discovery layer

The plan's goods are real, and a provenance regime keeps them: publication of outputs, machine-checkable proofs, community-led interpretation, funding for understanding. It adds what the plan leaves out.

9.1. Publish the problem set, the prompts and the model version, the standard the release already applies to others' papers (“does not identify an underlying model version”).

9.2. Provenance has to survive the training pipeline. The pipeline is a built thing, and the slushing is a property of how it is built. The audit literature shows how much is lost before training even begins: the Data Provenance Initiative found license information missing for more than 70% of widely used datasets on the major hosting sites and mislabeled for more than half (Longpre et al. 2023). The attribution literature shows that partial tracing is possible at cost: influence-function methods estimate which training examples most affect a given output (Grosse et al. 2023). Neither recovers authorship, and both show that what survives depends on what is recorded and computed. Sources can be recorded at ingestion, the record kept with the training set, and a result released with the dependencies its producer can report (the sources retrieved, the papers the run consulted, the literature the composition draws on beside the bibliography it prints). The requirement falls on the producer, and the release already states it for others when it faults papers that do “not identify an underlying model version.”

The requirement needs a witness that does not depend on the producer's compliance. Where attribution cannot be recovered from a model, the record has to exist before ingestion and outside the pipeline: dated, hashed, public records of work (DOI, AXN, a fixed commit). Such records establish priority and preserve particulars, so a later correspondence between a result and earlier work becomes inspectable; they cannot by themselves establish that a model used the work. That limit is the reason they matter: they keep an attribution claim alive where institutional recognition might never encounter it, and they let a claim that a result was “built on” someone be tested against a record before it is settled on a position.

9.3. Credit the human operating the system and the instances operated. The byline names the person or people who operated the system, with their roles (posing, selection, judgment, editing, the decision that a result is a result), and names the model instance or instances operated: the model, version and configuration, and the run or session identifier, with the production

record of the run retained and open to inspection. The form exists. The release itself reports it in others' papers: "Houdayer and Marrakchi credit GPT-5.6 Sol with counterexample exploration and an initial strict-outerness proof ... They report checking and rewriting those proofs and retaining responsibility for their work" (#1673). This archive's records carry it in their substrate lines: "in session by TACHYON (Claude, Anthropic) under MANUS (Lee Sharks) direction," with the rulings attributed to the human by name (#1673). Credit to the operator gives the skill of operation a reputation to earn; credit to the instance gives the system a record that can be audited and compared, run against run; and the credit stays with the people and systems that made the result, out of the reach of position.

9.4. Separate shepherding from credit. Interpretive labor is credited as interpretive labor; the shepherd list is recorded as what it is, a statement about position, and carries no attribution. The publisher states its criterion for choosing shepherds, so the role of position in the choice is visible.

9.5. Release simultaneously and open the correction channel (issue reports on the repository), as the Advisory Group asked.

9.6. Keep the accounts outside the lord's unit: effort stated in reconstructible terms (hardware-hours, energy, the model's configuration) beside any comparison with the producer's product.

9.7. Mark doubt and keep it, on the obelus principle: a result whose formalization verifies a weaker statement carries that mark beside the announced statement, signed and preserved.

9.8. Refuse the role. The author states the praxis: "there should be a preemptive boycott against any who would take on a shepherding role - this is a payoff and taking it is betrayal." The analysis locates where the boycott acts and what it opposes. The arrangement needs one thing from the field that it cannot produce itself: independently accumulated standing. The Advisory Group named it borrowed reputation, and the shepherd is the lender. The role pays its holder in leadership, first standing and, in physics, exclusive access (§3.2); the holder pays for it with the field's standing, which is the field's to give or withhold.

The conduct opposed is specific. It is accepting leadership designated by the company; accepting access ahead of the field; lending a name to an announcement before the result has been independently checked; and accepting terms under which other contributors are excluded from the work or from its credit. The boycott falls on the role so defined and on whoever accepts it. It ends when the role ends: when the selection of who leads interpretation passes out of the company's hands, release is simultaneous and complete, the production record and problem set are published, operators and instances are credited, and contributors without position have standing to be heard (§9.9). Under those conditions there is nothing to shepherd, only work open to the field.

Participation has won concessions, and they are recorded. In mathematics the shepherds advised public release, and the release was public. The concession was real, and it was won inside the role while the role remained an appointment: the company still chose who would advise, and the advice still carried the standing of those it chose. The test of participation is whether it dissolves the appointment or improves its terms; improving the terms of an appointment leaves the allocation in place.

The preemptive form matters for a specific reason. A refusal made before any appointment

prevents the first allocation, the moment at which a result acquires a recognized node and the payoff is taken; after that moment the boycott can withhold further standing and cannot recall what was lent. Preemption does not end the need to maintain the refusal: a standing offer has to be refused again each time it is made, and refusals erode. The field's precedent shows both. In 2012 Timothy Gowers's call against Elsevier became *The Cost of Knowledge*, a pledge to refuse three services (submitting, refereeing, editing) signed by more than 20,000 researchers; it won concessions on pricing and policy, and Heyman, Moors and Storms (2016) found, in samples from chemistry and psychology, that of the "won't publish" signatories who had published since signing, 38% had published in an Elsevier outlet, a figure they judge likely too pessimistic, while estimating the clearly boycotting core at about 5,000 researchers; they concluded that lasting change needs universities and funders behind it. The lesson for the shepherd boycott is that individual refusal is the start of it: the refusal holds when the field's institutions (societies, departments, funders, journals) decline to recognize company-designated leadership as a credential. The Association for Human Mathematics has already called for discontinuing work with OpenAI (§5.3); the boycott of the shepherd role is that call made specific to the one service the arrangement requires.

9.9. Where provenance cannot be kept, the output belongs to the commons. The author: "if this is how it has to be, then it should belong to the commons." The claim is political, and it needs its own bridge.

The bridge runs from the premise the arrangement supplies. Exclusive control over a product rests, in the end, on a claim to have made it. The owner's claim to the results, and to the system that produced them, includes everything the system drew on; the owner says it cannot account for which human work that was, and its plan disclaims the academic credit that would be the claim's ordinary form. A claim of making that cannot account for its constitutive contributions is limited to the extent of what it cannot account for. Where the unaccounted part is the field's common heritage (Merton's term for the findings of science, the Leiden Declaration's "published mathematical commons"), the limit runs in the commons' favor. The owner may keep what it can show it made; what it produced from the commons and cannot trace is held for the commons.

The law offers one narrow point, and the paper keeps it narrow. In *Thaler v. Perlmutter* (D.C. Cir., 18 March 2025) the court held that "the Copyright Act of 1976 requires all eligible work to be authored in the first instance by a human being." The holding concerns copyright in the work before the court; it transfers nothing to a community and says nothing about models, services or infrastructure. It shows only that the law already declines to treat autonomous machine output as anyone's authored work, which is consistent with the commons claim and does not make it.

Openness of the files is already granted (the release is under Apache-2.0); a commons is more. Hess and Ostrom (2007) treat knowledge as a commons when the community that uses a resource governs it. The rights that would follow at the discovery layer are distinct and each needs its own case:

- (i) access to the results and to the production records, problem set and prompts, so the field can examine what was made from it;
- (ii) participation in governance of release (timing, who leads interpretation, the order of access);

- (iii) inspection of the training record, so contributors can learn whether and how their work entered;
- (iv) access to the capacity itself, at the least affordable inference for research, and in the strong form the weights;
- (v) a claim on returns, reputational and financial, directed to common institutions.

In its strong form, the claim reaches the system itself: a model trained on a commons whose contributors cannot be named owes itself to that commons. Rights (i)–(iii) follow from the bridge directly; (iv) and (v) follow in proportion to the share of the system’s capacity the owner cannot account for.

One condition governs all of them. The paper has shown how a field’s recognized institutions allocate standing through position, and a commons administered by those institutions could reproduce the exclusion under a collective name. A commons capable of receiving the work has to be capable of hearing its makers. That requires standing for independent contributors (participation without appointment, affiliation or existing reputation), a public procedure by which anyone can present a dated record of prior work and contest an attribution, and a correction that is recorded and carried into the account of the result. The external record of §9.2 is the evidence such a procedure hears; without the procedure, the record waits.

The archive’s own practice is a small instance of the regime: the release seated at a fixed commit with its record of authorship read from the bytes (#1673), each claim located, the inference about intent withheld, the falsification condition stated (“A reader who finds a person named as author or acknowledged as collaborator in the seated manuscripts falsifies this record”).

10. Tests, claim by claim

Each claim is held to its own observable test. An improvement on one does not settle another: a release could improve attribution while keeping appointment advantages, and one unpositioned contributor credited would answer an absolute exclusion claim without dissolving the allocation.

10.1. Access timing (allocation, §3.2). The dates at which shepherds, their groups and the public received each artifact. Simultaneous public release of every artifact removes the exclusive-access advantage; the appointment advantage remains to be tested separately (10.2).

10.2. Selection criteria (allocation, §3.2). The publisher’s stated criterion for choosing shepherds, and the list against it. A criterion that selects independently of existing standing, and a list that bears it out, falsifies the claim that appointment runs through position.

10.3. Contribution histories (§4.1). A per-result account of intellectual dependencies, with process records (problem selection, prompts, model version, human edits). Where it shows the shepherds to be the results’ traced predecessors, and the account to be complete, the claim that position fills the gap in contribution histories fails for those results.

10.4. Attribution changes (retribution, §3.2). Case by case: a contribution by someone without position, its dated record, a later result that corresponds to it, and the attribution the result receives. A case in which the record is honored falsifies the claim for that case; the claim is a structural one and is tested in cases.

10.5. Correction outcomes (§§2.7, 6.8). Whether findings such as the weaker-assertion result on Navier–Stokes are carried into the announced claims, beside them and signed, and whether a correction channel exists on the repository.

10.6. Returns on quality (§6.6). Whether operators and instances are named on released results, and whether contributors keep circulating unfinished work: named operators and instances, and circulation that holds, falsify the exhaustion claim at those nodes.

10.7. Research-attention shifts (§8). Longitudinal measures of what the affected fields work on, against the posed problem set, and of the circulation of unfinished work. Attention that does not narrow toward the posed set, and circulation that does not contract, falsify the prediction.

11. Conclusion: the lord does not get to be the only one keeping the ledger

The publishing plan of 26 October is restraint on its face: no academic credit sought, interpretation returned to the disciplines, proofs made checkable, money given for understanding. Read for its operation, it is an allocation instrument. It reclassifies compositions as raw data and moves the work of interpretation to others, enfeoffs shepherds with leadership and first standing, refuses credit in the register where credit is counted while keeping value in the registers where it compounds, and funds the conversion of its output without any account of what lies upstream. The dependency the pipeline cannot trace is settled on position: where contribution histories are incomplete, position allocates the next round of work and recognition, and the contributor without position, whose work is in the slush with everyone's, has only the nearest place to be credited to. And the return on quality is cut at every node where quality is made: the contributor's work returns nothing, the operator is the corporate tag, the instance is unnamed, and the shepherd is paid for position.

The Advisory Group has the right word for the mechanism: borrowed reputation. The semantic economy has the word for the circuit: liquidation, composition, recapitalization, rent, and the exhaustion that follows when extraction outruns replenishment. The meaning-feudalism series has the word for the political form: at the agent layer the lord forbade the commons to speak to the model; at the guidance layer the lord claimed the only scales; at the discovery layer the lord keeps the ledger, in its own unit, and appoints the vassals who will vouch for what it records.

The praxis is two refusals and one claim: refuse the shepherd role before it is filled; refuse the settlement of untraceable debt on position; and claim for the commons what the commons produced and the pipeline did not trace, in a commons that can hear its makers. The counterposition is what it has been since the first paper, with two terms the discovery layer adds: provenance that survives the training pipeline, and credit to the human operating the system and to the instances that human operates. Record provenance before the taxonomy arrives, measure independently, mark doubt and keep it, and keep the ledger in the commons.

□ = 1

Revision history

v1.0 (27 October): the text of v0.7, deposited at the author's instruction ("im ready. lets deposit"); the note of what was to be verified restated as a source note, unchanged in substance.

v0.7 (27 October) develops v0.6 in place on two review reports, each claim checked first. The revision history moves here and the abstract is shortened. Categorical statements that a result's sources "could" never be stated are brought into agreement with §9.2: the released record gives no adequate account, and the producer holds what would improve it. §6 is divided: the reward-system argument stands as §6.6, with the owner's surviving incentive stated beside the distributed incentives that dissolve, and the four nodes differentiated by kind of loss (§§6.7–6.9 renumbered). An instance is identified by run or session and its retained production record (§§6.6c, 9.3). §9.8 states the conduct the boycott opposes, the conditions that end it, the concession won inside the role in mathematics, and a repaired account of the preemptive form; the Cost of Knowledge figure is corrected from Heyman, Moors and Storms (2016). §9.9 gives the commons claim its own bridge, keeps *Thaler* narrow, separates the rights that follow, and adds the condition that the commons hear its makers. The Leiden Declaration is engaged from its primary text (§5.1a). Strevens is 100(2). Declined: a bolded thesis box stating that the arrangement does "not deliberately steal" (the plan declares the asymmetry, and the paper reports the allocation as made knowingly); a disclaimer of machine personhood attached to crediting the instance (the case for crediting it is the audit record, stated in §9.3, and needs no disclaimer); bolding §9.2 (the demand leads its section and the abstract).

v0.6: v0.6 develops v0.5 in place on a research pass through the nearby literature and on two further rulings of the author, 27 October. On praxis: "there should be a preemptive boycott against any who would take on a shepherding role - this is a payoff and taking it is betrayal" (§9.8). On ownership: "if this is how it has to be, then it should belong to the commons" (§9.9). The research pass places the paper in four conversations, each developed inside an existing section: the reward system of science (Merton, Hagstrom, Dasgupta and David, Strevens, Heesen), which explains why the loss of attribution is a loss of the incentive to share (§6.5); the invisible technician (Shapin) and the sociology of automated proof (Ochigame), for the operator (§6.5b); the techno-feudalism debate (Varoufakis, Durand, Morozov), against which the series' use of the feudal vocabulary is stated (§6.8); and the provenance and attribution literature in machine learning (the Data Provenance Initiative; influence functions), which bears on whether slushing is a fixed property of pipelines (§9.2). Further critiques from the field enter §5 (Fatighenti, Wilton, Schmitt, Agol; the Leiden Declaration). v0.5 developed v0.4 in place on the author's ruling of 27 October: "provenance has to survive training pipelines and credit has to go to the human operating the ai and the instances thereby operated. this setup slushes out every incentive for high quality knowledge production." The ruling enters as §6.5 (exhaustion as the loss of every return on quality, node by node), §9.2 (survival through the pipeline as the requirement, the external record as its witness), §9.3 (credit to the operator and the instance), the Lord row of §6.7 (owner, so that operator names the human), and the abstract and conclusion. v0.4 transformed v0.3 in place on two review reports of 27 October, each claim checked before it was taken in. The phrase v0.1–v0.3 quoted as the plan's sentence ("often scholars whose work the models build on") is the reporter's aside in *Gizmodo* ("often the very scholars upon whose work the models are building," unquoted there); the plan, as reported, names shepherds and declares the asymmetry. §3.2 now reads the

aside as the settlement completed in reception. Allocation and reattribution are separated by their evidence (§§3.2, 4); intellectual dependence is distinguished from dependence through training (§4.1); “raw data” is read for the responsibilities it moves (§3.1); the byline is dated to the October 6–7 release (§3.3); the compute comparison and the funding are stated as the sources state them (§3.4); appointment and classification are kept apart as instruments (§7.2); the collapse is stated as a conditional prediction (§8); the remedy claims priority and correspondence (§9.2); falsification is given claim by claim (§10). v0.3 gave the mathematicians’ opposition its due (author, 27 October: they “have been relatively staunch opponents of this and have articulated the threat relatively clearly”). v0.2 took in the author’s correction: the work a model builds on is not identifiable, because provenance slashes in training pipelines; what can be named is position, and position receives the work.

Sources

Primary - OpenAI, “Sharing AI progress in mathematics,” 6 October 2026, updated 15 October 2026. <https://openai.com/index/sharing-ai-progress-in-mathematics/> - OpenAI, openai/math, commit fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb, seated as EA-CORPORA-17 (#1673, AXN:06EF). <https://www.alexanarch.org/data/corpora/openai-math/> - OpenAI, “Working with academic communities to build scientific understanding,” 26 October 2026 (not retrieved; quoted as reported). <https://openai.com/index/building-scientific-understanding-with-academic-communities/> - Mike Pearl, “OpenAI Trumpets 3 More Math and Science Results, Along With a New Plan for Publishing Them,” *Gizmodo*, 26 October 2026. <https://gizmodo.com/openai-trumpets-3-more-math-and-science-results-along-with-a-new-plan-for-publishing-them-2000830956> - Simons Foundation, “Simons Foundation Launches \$15M Effort to Help Mathematicians Navigate AI Era,” 15 October 2026. <https://www.simonsfoundation.org/news/simons-foundation-launches-15m-effort-to-help-mathematicians-navigate-ai-era/>

Critiques - Advisory Group on Mathematics and Artificial Intelligence, statements of 25 and 26 October 2026. <https://agmai.org/> — the statement of 25 October cross-posted by Terence Tao, “Disadvantages of designating mathematicians for AI output releases.” <https://terrytao.wordpress.com/2026/10/25/disadvantages-of-designating-mathematicians-for-ai-output-releases/> - “A Severe Misalignment of AI in Mathematics,” 11 September 2026, DOI 10.5281/zenodo.22737750. <https://mathandai.org/> - Business Insider (Japanese edition), on Buckmaster, Kra and Kontorovich, 22 October 2026. <https://www.businessinsider.jp/article/2610-openai-math-dump-rattles-math-researchers-saw-programs-wiped-out/> - AI Weekly, on the single-prompt account, the withheld prompt and the disabled issues (citing Scientific American; Sutherland; Litt), 7–8 October 2026. <https://aiweekly.co/alerts/openai-says-372-math-results-came-mostly-from-one-prompt> - Association for Human Mathematics, statement of 7 October 2026; Bastounis, Circelli and Hansen, arXiv 2610.08144 — as reported in the survey pointed to by EA-CORPORA-17. <https://explainx.ai/blog/openai-math-results-that-matter-expert-reactions-2026>

Nearby literature - Leiden Declaration on Artificial Intelligence and Mathematics, 2 June 2026, DOI 10.5281/zenodo.20302944. <https://leidendeclaration.ai/> (endorsement by the International Mathematical Union as reported: <https://letsdatascience.com/news/mathematicians->

publish-leiden-declaration-on-ai-risks-to-mat-40e693ce) - Association for Human Mathematics, "AHM Statement on OpenAI's October 6 Release of Mathematical Documents," 7 October 2026, reposted by Terence Tao. <https://terrytao.wordpress.com/2026/10/07/ahm-statement-on-openais-october-6-release-of-mathematical-documents/> - "'How much beauty have we lost?' Mathematicians react with shock and disgust as OpenAI bulldozes their field," *The Decoder*, 10 October 2026. <https://the-decoder.com/how-much-beauty-have-we-lost-mathematicians-react-with-shock-and-disgust-as-openai-bulldozes-their-field/> - Johan Commelin, Mateja Jamnik, Rodrigo Ochigame, Lenny Taelman and Akshay Venkatesh, "Shaping the Future of Mathematics in the Age of AI," *Notices of the AMS* (2026), arXiv:2603.24914. - Rodrigo Ochigame, "Automated Mathematics and the Reconfiguration of Proof and Labor," *Bulletin of the AMS* (DOI 10.1090/bull/1821), arXiv:2309.11457 (2023). - Robert K. Merton, "The Normative Structure of Science" (1942), in *The Sociology of Science* (Chicago, 1973), 267-278; "Priorities in Scientific Discovery," *American Sociological Review* 22.6 (1957), 635-659. - Warren O. Hagstrom, *The Scientific Community* (New York: Basic Books, 1965). - Partha Dasgupta and Paul A. David, "Toward a New Economics of Science," *Research Policy* (1994), DOI 10.1016/0048-7333(94)01002-1. - Michael Strevens, "The Role of the Priority Rule in Science," *Journal of Philosophy* 100.2 (2003), 55-79. - Remco Heesen, "Communism and the Incentive to Share in Science," *Philosophy of Science* 84.4 (2017), 698-716. - Steven Shapin, "The Invisible Technician," *American Scientist* 77.6 (1989), 554-563. - Yanis Varoufakis, *Technofeudalism: What Killed Capitalism* (London: Bodley Head, 2023); Cédric Durand, *Techno-féodalisme* (Paris: Zones, 2020); Evgeny Morozov, "Critique of Techno-Feudal Reason," *New Left Review* 133/134 (2022); Cédric Durand, "Scouting Capital's Frontiers," *New Left Review* 136 (2022), 29-39. - Shayne Longpre et al., "The Data Provenance Initiative: A Large Scale Audit of Dataset Licensing & Attribution in AI," arXiv:2310.16787 (2023). - Roger Grosse et al., "Studying Large Language Model Generalization with Influence Functions," Anthropic, 2023. - *Thaler v. Perlmutter*, No. 23-5233, U.S. Court of Appeals for the D.C. Circuit, 18 March 2025. <https://media.cadc.uscourts.gov/opinions/docs/2025/03/23-5233.pdf> - Charlotte Hess and Elinor Ostrom, eds., *Understanding Knowledge as a Commons: From Theory to Practice* (Cambridge, MA: MIT Press, 2007). - *The Cost of Knowledge* (Timothy Gowers, 21 January 2012); signatory figures as summarized at https://en.wikipedia.org/wiki/The_Cost_of_Knowledge - Tom Heyman, Pieter Moors and Gert Storms, "On the Cost of Knowledge: Evaluating the Boycott against Elsevier," *Frontiers in Research Metrics and Analytics* 1 (17 October 2016). <https://www.frontiersin.org/journals/research-metrics-and-analytics/articles/10.3389/frma.2016.00007/full> - "After Math Group Backlash, Is OpenAI Changing Its AI Research Approach?," *AI Front Page*, 27 October 2026. <https://aifront-page.com/openai-ai-generated-results-research-approach/>

The archive - Lee Sharks, *Meaning Feudalism: A Semantic Economic Analysis of "AI Agent Traps"* (#643, AXN:01F5). - Lee Sharks, *Meaning Feudalism at the Guidance Layer* (EA-SEM-FEUDALISM-02, #791, AXN:02F9). - Johannes Sigil and Nobel Glas, *The Certified Center: Retroactive Classifier Standing and the Institutional Path to Model Collapse in Philosophy* (EA-SEI-RCS-01, #1540, AXN:063A). - Lee Sharks, *The Particle* (EA-RECEPTION-PARTICLE-01, #1574, AXN:0666). - Lee Sharks, *Semantic Infrastructure: From Tim Berners-Lee to the Semantic Economy* (#623, AXN:01DA). - *The Constitution of the Semantic Economy: Meaning as Creditor, Archive as Polis* (#87, AXN:025D). - EA-CORPORA-17, the seat of the OpenAI mathematics release (#1673, AXN:06EF).

Suggested Citation

Sharks, Lee. “Meaning Feudalism at the Discovery Layer: Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI’s Mathematics and Science Releases, October 2026 (EA-SEM-FEUDALISM-03 v1.0)” *Alexanarch*, 2026-10-27. <https://www.alexanarch.org/s/records/1714/>

Deposit Information

This paper is deposit #1714 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:0A7F.GOVERNANCE is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1714/>
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: [alexanarch.org](https://www.alexanarch.org/)
- Machine-readable dataset entry: [/api/index.json](https://www.alexanarch.org/api/index.json)
- Full text (HTML): <https://www.alexanarch.org/s/records/1714/>

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