This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)

Sharks, Lee

2026-11-16

This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)

Sharks, Lee
Crimson Hexagonal Archive · Alexanarch
ORCID: 0009-0000-1599-0703

2026-11-16 · Version v1.0 (the working notebook at v0.21) · Journal article (Machine-Mediated Reception Studies); working notebook: census, reach study and model test

AXN:0CA4.GENERATIVE

https://www.alexanarch.org/s/records/1772/

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": "This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)",
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-11-16",
  "identifier": "AXN:0CA4.GENERATIVE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
```

```
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1772/",
  "description": "A census of twenty philosophy journals (ancient, continental, analytic) and two Persée series
checked verbatim span. In Ancient Philosophy 8 of 844 articles open the question of what the object is; since 20
CORPORA-39 (#1771)."
}
```


Abstract

A census of twenty philosophy journals (ancient, continental, analytic) and two Persée series, coding each article's object and, from its abstract, what it does with that object, every code with a machine-checked verbatim span. In Ancient Philosophy 8 of 844 articles open the question of what the object is; since 2021, 4 of 283 Plato/Aristotle abstracts across the sister journals. The three fields are three placements of one unit, the article that places a position. In the Archiv für Geschichte der Philosophie of the 1890s one ancient title in five puts a work's standing in question; by the 1920s the question is gone from its pages; philology kept it longer and unevenly. A model asked for plausible abstracts reproduces each field's unit, raises its characteristic move, writes no claim about the text in sixty ancient abstracts, and deals one inventory of topics across mastheads. Disciplinarity is read as the stable form of the article economy; the editor's correspondence is read as the gate's specification of the unit. The texts are seated as EA-CORPORA-39 (#1771).

Body

---

deposit_number: 1772 hex: 0CA4 title: "This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)" creator: Sharks, Lee orcid: 0009-0000-1599-0703 date: 2026-11-16 content_type: "Journal article (Machine-Mediated Reception Studies); working notebook: census, reach study and model test" license: CC-BY-4.0 substrate: "Human–machine collaborative. The study's question, predictions, title, subtitle, rulings and readings are Lee Sharks's, stated and dated in the text. The census, coding, measures, model-test analysis and drafting were done in session by TACHYON (Claude, Anthropic) under the author's direction; the reach codes were made by parallel Claude readers and the lead Claude instance, as §0.5b states, with agreement measured within that procedure only. The model test's abstracts were generated by ChatGPT, signed out, run by the author. Reviews by Gemini, DeepSeek and ChatGPT were checked point by point before any was applied. The Crimson Hexagonal Archive is credited as contributor: the measures read its seated texts (EA-CORPORA-39)." version: v1.0 (the working notebook at v0.21) related_ids: "AXN deposit #1771 — EA-CORPORA-39: the seating of the texts this study reads (the Archiv 1888–1923, the construction of the object, the philology volumes, the census, the gates)\nAXN deposit #1540 — The Certified Center: the institutional path to model collapse in philosophy (2.4)\nAXN deposit #1573 — The Wrong Unit: the model-collapse self-diagnostic whose measures are applied in 4.12h and 4.17\nAXN deposit #855 — The Wolf Boy and the Language Model\nAXN deposit #751 — Socially Necessary Scholarly Labor\nAXN deposit #122

— Coverage vs. Depth\nAXN deposit #1770 — Who Speaks for the Classics, and on What Basis: the first station of the combined audit (MMRS)\nAXN deposit #1769 — Daffodils: Classical Receptions Journal, 2009–2026\nAXN deposit #123 — Socrates as Orthonym, the paper whose submission §0A records" axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - This Is Not Philosophy - article economy - disciplinarity - shape of the field - Machine-Mediated Reception Studies - ancient philosophy - Ancient Philosophy (journal) - Phronesis - Apeiron - Elenchos - Rhizomata - Archiv für Geschichte der Philosophie - reach coding - consistency repair - athetesis - authenticity - standing of a work - Crossref census - stylometry - abstract template - continental philosophy - analytic philosophy - model collapse - model test - credentialing - peer review - AI policy - "Philosophy & Public Affairs" - Revue philosophique de Louvain - Revue des Études Grecques - Philologus - Classical Quarterly - Brucker - system of philosophy - Polansky *** # This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)

This Is Not Philosophy

The Article Economy and the Shape of the Field

Lee Sharks · Crimson Hexagonal Archive · ORCID 0009-0000-1599-0703

Machine-Mediated Reception Studies · EA-MMRS-TINP-01 · working notebook v0.21 · deposited 16 November 2026

Status of §0A: the author's correspondence with Ronald Polansky, editor of Ancient Philosophy *(Duquesne), 23 July – 9 September 2026. Four threads and 41 messages, read in full with quoted replies stripped; the attachment snub-poemed.pdf was not read. The author's ruling (15 November): the correspondence is with the editor in his official capacity, on the substance of his field. It is institutional, not private, and it is quoted as such. The editor is named as the accountable party, as in the CRJ audit (#1770). Scope, ruled the same day: his responses to the submissions and near-submissions, and what he looks for in submissions, are quoted. His statements of his own views are not.*

Abstract. This notebook measures what the journals of philosophy now publish and asks what the measurement says about the discipline that produces them. A census of twenty journals in Crossref and two series in Persée, ancient, continental and analytic, codes each article's object and, from its abstract, what it does with that object: re-read a passage, resolve an apparent inconsistency, reconstruct a doctrine, make a claim about the text itself, relate authors, carry the material past antiquity, or argue a position of its own. Every code carries a verbatim span of its abstract, machine-checked. In *Ancient Philosophy*, 8 of 844 articles open the question of what the object is; across the sister journals since 2021, 4 of 283 Plato/Aristotle abstracts do. The three fields are three placements of one unit, the article that places a position: the ancient field repairs a named object, the continental field pairs names, the analytic field takes a position in a debate, and the form of the abstract converges on the analytic. The question was once ordinary. In the *Archiv für Geschichte der Philosophie* of the 1890s one ancient title in five puts a work's standing in question, and its pages carry the question at that rate until the war; by the 1920s it is gone. Philology kept it longer, and unevenly. A model asked for plausible abstracts reproduces each field's unit, raises its characteristic move, writes no claim about the text in sixty ancient abstracts, and deals one

inventory of topics across the mastheads. The notebook argues that disciplinarity is the stable form of the article economy: the journal certifies the person through a unit that cannot fail, and the field takes the shape of what can be certified. The author's correspondence with the editor of *Ancient Philosophy* is read as the gate's own specification of that unit. The texts the measures read are seated as EA-CORPORA-39 (#1771).

---

0. The question

0.1 **The question.** The author (15 November 2026): "the real question here is how disciplinarity murdered philosophy as such. much has been done on the analytic / continental divide etc… whats the shape of the field?"

0.1a **Why the question is philosophy's.** The conventional reading first, in full. What a text is, who wrote it, when, whether it is one work or several, is philology's question; philosophy takes the text as given and asks whether its arguments are sound. The division is a reasonable division of labor, it lets each discipline do its own work well, and on this reading the census measures only that philosophers do philosophy and leave philology to philologists. The notebook's reading: the question of what the object is was philosophy's question before it was any discipline's. The Socratic *ti esti* asks what a thing is before anything is said to follow from it. Aristotle begins his inquiries from what is said and by whom, and closes the *Sophistical Refutations* with a claim about the standing of his own inquiry, what of it existed before him. Where the corpus is the evidence for the doctrine, as it is wherever a dialogue is attributed, a treatise ordered or a fragment assigned, the identity of the object is a premise of every reconstruction built on it, and an argument read without asking which text it is in, whose, and in what order carries that premise unexamined. A discipline that treats the premise as settled has removed it from examination, whatever division of labor it calls the removal. That is the sense of the title. "This is not philosophy" is the gate's verdict on the submission that reopened the premise (§0A), and the census's verdict on what passes the gate: a field in which every argument stands on a premise no article may examine.

0.2 **The prediction.** The author (15 November): "my prediction is that we'll see a tightening of ontological reach even within the contracting range of topics." Here a *topic* is the ancient object an article takes up, such as the *Republic* or the *Metaphysics*. *Reach* is what the article does with that object: re-read a passage, repair an inconsistency, reconstruct a doctrine, or make a claim about the text itself.

0.3 **What is sought.** The author: "im looking for model collapse". The notebook treats this as a technical term. It names a system trained on its own outputs, whose distribution narrows and whose tails disappear. Here it names the mechanism the notebook proposes (3.3). The census measures the distributions; the recursive dependence between them is still to be traced.

0.4a **The second prediction.** The author (15 November): "the fields have the most animosity to ai where their own work is already the most algorithmic. these are operations ai could spit out by the thousands and tens of thousands with no need of human clinamen". The test is in §4.14–4.15.

0.4 **The horizon.** The author plans to collect "the entire range of seven articles" in "a com-

bined audit of what the fields cant carry". This notebook supplies the instrument for that audit and its first test case outside CRJ: *Ancient Philosophy* and five sister journals.

0.5b **Who coded.** The notebook's "readers" are instances of Claude (Anthropic), run in parallel in this session, one batch of abstracts each. Each was given the codebook and the lead coder's calibration codes, and told to read every abstract and code it by judgment, with no keyword rules. Each code carries a verbatim span of 6–30 words from its abstract, and the span is checked by machine against the source. That check verifies that the evidence exists, not that the category is right. The "lead coder" is Claude in the session's main thread (TACHYON, under the author's direction). It wrote the codebooks, coded the calibration sets, adjudicated the readers' hard calls and ran the reviews in both directions. The census, the regex pulls and the stylometry are code. Every hand ruling is logged. The reliability check is the blind calibration in 4.7: the readers agree with the lead coder on 75 of 94 (80%). That is agreement within the Claude coding procedure, the readers' consistency with the lead instance's classifications; it is not independent validation of those classifications, which would need a human second coder. Changes to the instrument across arms are stated where they enter: F redefined for the continental arm and G added (4.14a); the analytic extension to G (4.14c); article thresholds of 10 pages (ancient), 8 (continental) and 5 (analytic titles, 4.14e); and denominators of named-object works (ancient) against all articles (continental, analytic).

0.5 **Status of claims.** Every attributed claim in §1 is quoted from an open copy read on 15 November, or it is marked **[recalled, unread]**. Recalled claims are to be read before they are used. Every figure in §§3–4 comes from the census files listed in §8.

0.5a **The title.** The author (15 November): "This Is Not Philosophy." The sentence is the gate's verdict on the eccentric submission (§0A.2: "please stop writing in this eccentric way"; "who cares besides you"), and the census's verdict on what passes the gate (§§4, 6). Which speaker is right is the notebook's question. The subtitle (the author, 16 November): "The Article Economy and the Shape of the Field", replacing "How Disciplinarity Murdered Philosophy", because the earlier form gave disciplinarity an agency of its own; in the author's words, disciplinarity is "itself a political economically produced unit". The notebook keeps the author's question of 0.1 as asked, and answers it in those terms: the discipline is what the economy of the article produces (§5).

0.6 **The exchange guides the instrument.** The author (15 November): "we should allow the actual exchange with polansky to guide this… on his end, my words were not too standardized, but too eccentric. he offered a more personal exchange but ended by banning submission by topic/argument. and socrates as orthonym is about five orders of magnitude out beyond what any professional philosophy journal can carry." The exchange is the calibration point: §0A sets out what the gate asked for, and §§4–6 measure what passes it.

0A. The exchange

0A.1 **The conventional reading first, in full.** The editor of a specialist journal, Ronald Polansky of *Ancient Philosophy*, without referees he could find for an unusual submission, read the paper himself. He corresponded at length and gave practical advice: length, clarity, abstract, a narrower object, the state of the secondary literature. When the manuscript did

not meet the journal's standards, he declined it, and he left the door open for a different paper. On this reading nothing in the exchange is a mechanism. It is an editor doing his job with more patience than most.

0A.2 **The specification, in the editor's words.** Polansky's demands, in order: - **Vocabulary.** "I frequently request that authors rid papers of philosophical jargon" (23 July). Later: "Avoid off-putting technical vocabulary beyond that of the passages you are considering" (31 August), and "Do not have any words that are not easily recognized in the abstract" (2 September). - **The object.** "My recommendation… is instead to do some close analysis of an ancient text rather than such broad stroke sort of writing" (28 August). "I would much rather see close analysis or one or two" (2 September). - **The frame.** "It should begin with a discussion of the present state of thinking about the relationship of the two works, i.e., coverage of available secondary discussion" (9 September). - **The size.** "Anything over 40 pages is problematic" (29 July). "33 or so pages single-spaced, i.e., crazy long" (2 September). - **The claim.** "If the paper is not about Aristotle and Theophrastus being one person, but instead that the two works mentioned are closely connected, this could be O.K." (9 September). "Could we see a tamer version of this paper" (2 September). - **The register.** "PLEASE STOP WRITING IN THIS ECCENTRIC WAY" (31 August). "WHY NOT, THEN, PLAY THE SORT OF GAME NEEDED TO BE ACCEPTED FOR PUBLICATION RATHER THAN GAMES THAT YOU INVENT?" (31 August). "Do the sorts of things that you may resist, but that I might welcome" (31 August). - **The stake.** Of the closing section of *Many and One*, the count of makers, which returns a null in either direction: "I saw the stuff from the other paper about one body or two bodies. We really can do without any of this. Who cares besides you, and can you really offer them much reason to care? It seems quite questionable. So, could we see a tamer version of this paper" (2 September; in the same letter: "I have not read more than a few sentences of this paper"). Of *Socrates as Orthonym*: "We seem able to arrive at the same results without the hypothesis" (28 August), and "Why bother to claim one guy did it all?… So, what really is the point other than to entertain an unlikely story?" (28 August).

0A.3 **What the specification is, read against the census.** Each demand has a measured counterpart in §4: - the object narrowed to one or two texts, close analysis → A, the passage re-read (§4.6–4.7), and *Phronesis*'s passage-locus titles doubling from 11% to 20% (§4.5); - the present state of the secondary discussion first → the template, "scholars hold X; I argue Y" (§4.15); - vocabulary bounded by the passage's own → the claim stated in the object's own terms, so it cannot leave the object; - the tamer claim → D-confirm admitted, D-open refused (§4.8–4.9). "Closely connected" keeps both names. "One person" removes one. The admissible version is the one the census already contains in quantity: "Aristotle's Metaphysics Z as First Philosophy", "The Unity of Plato's Trilogy"; - the game "needed to be accepted" → the move the stylometry finds converging across fields (§4.15).

The demands are the specification of the field's unit, stated by the person who administers it.

0A.3a **Production divorced from content and form.** The author (15 November): "he sees production as completely divorced from content and form." The stake-demands state it. Who made the corpus, and how many made it, are treated as making no difference to what the texts say or how they are built: "the same results without the hypothesis"; "who cares besides you". The same letters treat production as decisive for the submission. The decline of 28 August lists among its grounds a prior version on Academia ("whether there should be multiple publications of versions of a paper") and "whether papers composed by A.I. partici-

pation should be considered as regular submissions". Production is irrelevant when it is the canon's and disqualifying when it is the submitter's. That asymmetry is the field's unit in its plainest form. The object is the text as received, and operations on it (A, B, C) are the work. The question of how the text was made (D, the count of makers) is detachable and dropped. And a text whose making is in question at intake is excluded before its content is reached. The author's reply of 28 August named the converse: "*how* it was produced can matter, but not *instead* of engaging the thing in itself."

0A.4 **Eccentric and standardized are one measure.** The editor's objection was that the words were too eccentric. The census finds the field's output standardized. These are the same distance read from two sides. Eccentricity is distance from the template, and the gate removes what is distant. A field whose intake filters the tails produces outputs that converge. In the model-collapse literature the tails are lost because they are under-sampled. Here they are refused at intake. The author's phrase is "banning submission by topic/argument". The mechanism in the letters is precise: the topic is admitted and the claim is not. "Closely connected" is the same topic with the claim removed. A paper that drops the claim can be submitted, and the claim, removed before review, cannot be assessed. The demands are reasonable in form and close the route in effect, the same split the CRJ audit found between procedural check and substantive judgment.

0A.5 **What the gate cannot carry, in scale.** The units the census finds opened, from the smallest: - an emendation, a line or two of Greek, on the order of 10^1 words; - the largest opening since 2010, the standing of one work ("On the Inauthenticity of the Critias", the authorship of [Arist.] *De spiritu*), on the order of 10^3–10^4 words; - *Socrates as Orthonym*'s object, the partition into persons of the Platonic and Aristotelian corpora, on the order of 10^6 words (estimate).

That is about five orders of magnitude beyond the field's typical unit and two beyond its largest. The author's "five orders of magnitude" holds against the census, measured from the unit the field ordinarily works in. The size is a proxy. The finding is the difference of type: the largest object the field opens is the individual work, and the partition of a corpus into persons is a different kind of object. The size follows from the type; it illustrates the difference and does not establish it.

0A.6 **The objection kept.** "We seem able to arrive at the same results without the hypothesis" is the one demand the author took as the real review (28 August): a hypothesis must earn a result the received partition strains for. It belongs with §7, as the standard the combined audit holds each field's excluded question to. The field's own unit does not have to meet it, because the field's unit was never a hypothesis. It is the default.

---

1. The received account: seven lines

1.1 **Indictments from inside.** The discipline has accused itself in public. - Frodeman and Briggle (*NYT*, The Stone, January 2016, "When Philosophy Lost Its Way"; excerpted at Daily Nous) date the break to the research university. The "institutionalization of philosophy made it into a discipline that could be seriously pursued only in an academic setting". This was "the act of purification that gave birth to the concept of philosophy most of us know today". It was

"about divorcing knowledge from virtue". Their verdict: "philosophy should never have been purified." "The point of philosophy now is to be smart, not good." "Having become specialists, we have lost sight of the whole." - Hanna **[recalled, unread]**.

1.2 **Institutional and sociological histories.** - Katzav and Vaesen, "On the emergence of American analytic philosophy", *BJHP* 25.4 (2017) 772–798, name the mechanism. Analytic philosophy came to dominate "at least in part due to its effective promotion by The Philosophical Review's editors". The "immediate cause is a change in editorial policy", namely "the replacement of editors who were open to a wide variety of approaches". This amounted to "the suppression, by institutional means, of existing diversity". Their counts for 1950–55: "roughly 65% of the full-length papers in the journal are mid-century analytic", against "roughly 10% (9 papers)" speculative. Their characterization: "more akin to a premeditated, hostile takeover". - Collins, *The Sociology of Philosophies* (1998); Kuklick; McCumber, *Time in the Ditch* (2001); Bourdieu, *Homo Academicus* **[all recalled, unread]**.

1.3 **The analytic–continental divide.** This is the line the author sets aside as "much has been done". The notebook does not take it up, except to note what the divide conceals. Both sides keep the discipline's unit, the article. What differs is where a position is placed and how far it is attached to names: the ancient and continental fields anchor about three-quarters of their work to a named author, the analytic field almost none of it (4.14c–e, 6.1).

1.4 **Philosophy as a way of life.** Hadot **[recalled, unread]**. This line holds that what was lost is a practice, not a doctrine. Frodeman and Briggle's "divorcing knowledge from virtue" is the same complaint in institutional form.

1.5 **How the history of philosophy was made.** - Catana (open copy, CORE) traces the historiographical concept "system of philosophy". It moved into methodology in the sixteenth and seventeenth centuries: Melanchthon "was among the first to transfer the term 'system' to methodology". Brucker was "instrumental in developing" the concept (*Kurtze Fragen* 1731–36; *Historia critica philosophiae* 1742–44). Brucker assumed that "philosophers have always intended to work out systems of philosophy, though not always with good results", and that "Even Plato did so, albeit imperfectly". Catana's objection: "this meaning of 'system' was simply unknown to pre-modern philosophers". He adds that the concept is "still used forcefully, and uncritically", and that it "certainly seems to obscure more than it reveals". His recommendation: "we may be better off leaving it behind in our future exploration of pre-modern philosophy." Loci from the open copy (preprint pagination; L. Catana, "The Concept 'System of Philosophy': The Case of Jacob Brucker's Historiography of Philosophy", *History and Theory* 44 (2005) 72–90): Melanchthon, p. 18; "instrumental in developing", p. 2; "philosophers have always intended…" and "Even Plato did so", p. 15; "simply unknown to pre-modern philosophers", p. 19; "still used forcefully, and uncritically", p. 36; "obscure more than it reveals", p. 40; "leaving it behind", p. 1 (abstract). The journal's printed pages are not in the open copy. - Park; Hunter; the Tübingen school (Krämer, Gaiser, Szlezák) on the unwritten doctrines and the dialogue-only Plato **[all recalled, unread]**.

1.6 **The paper form.** - Dennett, "Higher-order truths about chmess", *Topoi* 25 (2006) 39–41. "Chmess is just like chess except that the king can move two squares in any direction, not one." He warns against "committing your precious formative years to a research agenda with a short shelf life", and proposes a test: "try to teach the stuff to uninitiated undergraduates". If they don't engage, you may be following "a self-supporting community of experts into an

artifactual trap", and "you may all be taking each other for a ride". - Dotson, "How is this paper philosophy?", *Comparative Philosophy* 3.1 (2012) 121, names "a pervasive culture of justification within professional philosophy". Her claim is that the paper must justify itself *as philosophy* before it can say anything. - Millgram, *The Great Endarkenment* (2015) **[recalled, unread]**.

1.7 **The canon's composition.** Who is in the canon and who is out; the movements to widen it **[literature to be gathered]**. This line asks *which names*. It does not ask what is done to a name once it is admitted.

---

2. The gap

2.1 Each line holds one link of a chain, and no line holds the chain. - 1.5 explains how the object was built: as a *system*, retro-projected onto the ancients. - 1.2 and 1.6 explain how the institution is maintained: editorial policy, the paper form, the culture of justification, the self-supporting community. - 1.1 and 1.4 record the loss. - 1.7 argues over the membership list.

2.2 The synthesis this notebook proposes joins three things that the lines above treat apart: the **construction** of the object (18th–19th centuries), its **maintenance** by the article economy (20th–21st), and its **transmission to the machine** (now). The third link is new, and the received account cannot see it, because it is written in the article form the machine is trained on.

2.3 **Thesis of the notebook.** Disciplinarity is produced, and what produces it is the economy of the article: the journal certifies the person through the unit (5.1), the unit that cannot fail is the one that certifies, and the field takes the shape of what certifies. That economy did not end philosophy by changing its doctrines. It made the *unit of work* the article that places one position: in the ancient field most often by repairing a named object, in the continental field by pairing names, in the analytic field inside a debate (4.14, 6.1). That unit, iterated, behaves like a model trained on its own outputs. The operations it performs were already narrow by 2010 and hold there: three-quarters of the field's work stays inside one author (§4.7). The objects keep narrowing (§4.2–4.5). And the tail goes: the article that reopens what the text is (§4.8). The machine then inherits a corpus that has already collapsed.

2.4 **The archive's own line, and what this notebook adds to it.** The received account (§1) does not hold the chain; the archive has held parts of it, as theory and as instrument, before this census measured anything. - *The Certified Center* (Johannes Sigil and Nobel Glas, #1540, AXN:063A, 24 August 2026) read the *P&PA* policy on the day it was posted. It quoted the journal's two functions, "the promulgation of new knowledge, and identification and credentialing of talented researchers", and drew the consequence 4.16a and 5.1 reach here: the journal "legislates on A", the authorship process, not on Q, the quality it concedes. It named the arrangement *Ratificatory Human Governance* (the classifier supplies the fact, the journal the force) and *Retroactive Classifier Standing*, and argued that a discipline whose credential is machine-scored contracts as model collapse does, tails first, with the loss "experienced as improvement". Its §VI sets the obelus, athetesis as a signed, kept, reversible mark, as the custody principle through which philosophy's primary corpus survives, and says the disci-

pline "has forgotten it". Its §VII asks for "a tail-preservation observatory over the admitted corpus", and its §IX predicts that the institutions in a position to take that measurement will not take it. - *The Wrong Unit* (Nobel Glas, #1573, AXN:0665, 1 September 2026): per-item instruments read the head of a distribution while collapse accrues in its tail, so "a mode-seeking model beats a diverse one on any single turn"; the instrument that reads the state variable counts effective modes and cross-prompt convergence. - *The Wolf Boy and the Language Model* (#855, AXN:0364) and *Diversity Contraction* (doi 10.5281/zenodo.20532696): collapse is substrate-agnostic, a regeneration–pruning dynamic with a threshold below which recovery fails, and nothing in the present landscape adds friction. - *Socially Necessary Scholarly Labor* (#751, AXN:02A7) and *Coverage vs. Depth* (#122, AXN:02A8): the excess of the article's labor over its contribution is scholarly rent, the price of admission to legitimacy (literature review, positioning, defensive citation, length, genre conformity). - *Institutional-Prior Foreclosure* (EA-MPAI-IPF-01, AXN:02D0): a model takes institutional recognition as a lagged proxy for legitimacy, so that field-forming work is foreclosed hardest where its object is the recognition-distributing system itself.

What this notebook adds is the measurement those papers specified and did not take. The census is the observatory *The Certified Center* asked for, built from outside the institutions it predicted would not build it; its reach coding is the first component of that paper's "plural measurement membrane" ("independently coded argument form"), run on every abstract the three arms hold. It dates what that paper stated as a framing thesis: athetesis was an ordinary article form in the *Archiv* of the 1890s (4.12, 4.12h), gone from its pages by the 1920s (4.12g), gone from the *REG* after 1945 and thinned in *CQ* after 1970 (4.12c, 4.12f). And it moves the path to collapse back a century. Two collapses are distinguished here. The institutional one is variance reduced by selection, the editor's and the referee's, under the credential, from the 1910s on; the computational one is a model trained on what that selection admitted, from 2023 on. The field contracted first; the model inherits the contracted distribution and, on the evidence of 4.17, sharpens it; the detector regime of 2026 is the newest instrument of the same credentialing function, pruning a distribution the article economy pruned before it.

---

3. The mechanism

3.1 **Construction.** Brucker's "system" assumes that every ancient philosopher intended a whole (1.5). The fetishized philological object does most of its consolidating in the 19th century (the author, 15 November): a fixed text under a name, with a corpus edited and bounded. Two consequences follow for later work. The name becomes the unit of inquiry. And any apparent inconsistency becomes a defect for scholarship to repair, rather than a fact about the text.

3.2 **Maintenance.** The article is the field's unit. Its protocol is set by editors (Katzav and Vaesen), and its first obligation is to justify itself to the field (Dotson). Its safest product is a higher-order truth about a game already constituted (Dennett). The norm of charity turns 3.1's inconsistency-as-defect into a genre: *consistency repair* (reach category B, §4.6).

3.3 **The machine.** A model trained on the article stock receives the object already whole, already named, and already repaired. Asked about the field, it reproduces the field's own distribution, and the field's next articles are now written alongside that model. The archive's

captures document this recursion on the classics side: an Overview restates a field's self-description as fact. The computational link is now measured once, at the output (4.17): asked for twenty plausible abstracts per journal, the model reproduces each field's unit, raises its characteristic move, writes no D in sixty ancient abstracts, and deals one inventory of topics across mastheads. The institutional link (2.4) is the census. In the terms of *Diversity Contraction* the field is a substrate whose regeneration term is its tail, the article that asks what the object is; in Shumailov's recursion that is the retained share of original data, γ, and with $\gamma = 0$ the loss is unbounded. The corrective the collapse literature names, accumulated original data (Gerstgrasser et al. 2024), is in this case the field's own record before the flattening; EA-CORPORA-39 seats it for the machine reader (§8). The gates of 4.16 close the same channel from the other side: in *The Certified Center*'s words, "the corrective channel is gated by the pruning instrument".

---

4. The census

4.1 **Instrument.** - Crossref journal census with cursor paging. - Titles coded by ancient object, using a regex lexicon, with a fallback to the abstract where the title is silent. - *Article* means a page span of 10 or more. - Measures: - Plato+Aristotle share; - Shannon entropy over object groups (H; higher means broader); - passage-locus share (titles naming a book, line or fragment); - reopening share (titles making a claim about authenticity, dating, unity or corpus). - Files are listed in §8.

4.2 ***Ancient Philosophy.*** 2,089 works; 844 articles. *Correction (v0.6):* Crossref keeps subtitles in a separate field. Earlier versions read only the main title, which missed 734 *Ancient Philosophy* subtitles, so a title like "Physics and Tragedy: On Plato's Cratylus" lost its object. All censuses are now recoded on title plus subtitle; the prior files are kept in `pre-subtitle/`. Articles naming an object:

period	n	P+A %	H (groups)
1980–84	27	85	1.45
1995–99	94	71	2.25
2005–09	78	56	2.53
2015–19	96	58	2.50
2020–24	116	75	2.12
2025–26	52	79	1.90

The field widened for thirty years, from 1980 to 2019. Then it narrowed sharply after 2020, back to the Plato+Aristotle share of the early 1980s.

4.3 **The tails in *Ancient Philosophy*, 1980–2026.** - Articles on the Platonic *Letters*: 0. - On the unwritten doctrines: 1. - On Theophrastus: 1. - With subtitles, the regex flags 16 article titles as reopening the text. Read by hand, **8 of 844** ask what the text is. Three put a work's standing in question: the *Lysis*'s authenticity (1990), the dating of the anonymous Parmenides commentary (2000), and the authorship of Philoponus on *De anima* iii (2022). Five concern

text and fragments: the fragments of Parmenides (1992), a Philoponus papyrus (1992), Cornford's fragment (2010), Heraclitus B1–B2 (2017), and an Empedocles reconstruction (2024). The rest are false hits, mostly on *De generatione et corruptione*.

4.4 **Sister journals.** All named-object works, by decade.

journal	2000s P+A % / H	2010s	2020s
Ancient Philosophy	60 / 2.50	69 / 2.29	75 / 2.12
Phronesis	58 / 2.38	73 / 2.10	71 / 2.05
Apeiron	68 / 2.30	77 / 1.96	76 / 2.01
AGPh	65 / 2.25	75 / 1.91	70 / 2.05
Elenchos	76 / 2.02	63 / 2.29	66 / 2.37
Rhizomata (from 2013)	—	53 / 2.11	62 / 2.33

4.5 **Reading of 4.4.** - The four long-running journals in English and German share one shape: breadth peaks in the 2000s, then contracts in the 2010s and stays contracted. The narrowing belongs to the field, not to one journal. *Ancient Philosophy* comes to it late and sharply. - *Elenchos* and *Rhizomata* run the other way. These are the Italian journal and the newest one, and both are widening. That is a lead to follow, not a finding. - Passage-locus titles rise in *Phronesis* from 11% (2000s) to 20% (2020s): the unit of work is shrinking to the passage. - Reopening stays between 0.5% and 6% of named works in every journal and every decade.

4.6 **Reach in *Ancient Philosophy*.** 94 readable abstracts, 2022–26, read and coded one by one by the lead coder. - **A**, a passage re-read: 29%. - **B**, an inconsistency resolved: 27%, rising from 21% (2022–23) to 30% (2024–26). - **C**, a doctrine reconstructed inside one author: 32%. - **D**, a claim about the text itself: 3%. All three D cases confirm the object whole: the *Axiochus*, the unity of the *Phaedrus*, and the *Divisiones Aristoteleae*. - **E**, relations across authors or receptions: 9%. - **F**, past antiquity: 1%.

4.7 **Reach in the five sister journals.** 1,247 abstracts read and coded one by one, A–F as in 4.6. The coding was split across six readers (0.5b). Each code carries a verbatim span of 6–30 words from its abstract, and every span was machine-checked against the source: 1,341 of 1,341 valid. The 94 *Ancient Philosophy* abstracts were mixed in blind as a calibration set. The readers agree with the 4.6 coding on 75 of 94 (80%). Where they differ, they move B to A or C (8 of the 25 B's).

The lead coder then adjudicated, with 43 rulings, logged in `reach5/adjudication-log.json`: - 16 rulings on the readers' own flagged hard calls. - A review of the skew: every A or C row from 2010 on whose abstract frames the author's text as apparently inconsistent, defective or conflicting. 181 were read and 27 moved to B. - The rule for B: the abstract sets up an apparent defect in the author's text or doctrine, and the article defends the text against it. - The rule for D direction: an argument against the printed standard edition counts as *open*, since the edition is the field's object.

The reverse review (v0.7) read all 156 B rows from 2010 on and moved 17 out of B. These abstracts set up an interpretive problem or a scholarly dispute, but no apparent defect in the author's text that the article defends against (16 to C, 1 to A). That makes 60 rulings in all. Works with abstracts from 2010 on, excluding notes and items with no ancient object (n = 907):

period	n	A	B	C	D	E	F
2010–15	179	31	17	28	7	14	2
2016–20	325	23	13	35	7	17	4
2021–26	403	25	16	34	6	17	2

Field-wide, reach does not tighten across 2010–26. It was already tight. A+B+C, the operations inside one author, hold at 71–76% in every period. D holds at 6–7% and F never passes 4%. One journal moves: in *Apeiron*, B rises from 14% to 19% to 24%, the same rise as in *Ancient Philosophy*. It survives review in both directions.

4.8 **Against the prediction (0.2). - In *Ancient Philosophy*, it holds.** B is 27% there and rising, against 13–17% in the sisters as adjudicated; only *Apeiron* approaches it (24% in 2021–26). D is 3%, and all of it confirms the text. - **Field-wide, the prediction is answered differently.** At the A–F resolution, reach did not contract further across 2010–26; the plateau was already in place by 2010. That is a different result from the decline of the opening tail, and the two are kept apart here. What is still moving is the topic range (4.2–4.5) and one tail. In Plato/Aristotle works, D falls from 8% to 5%, and the D that *opens* the text falls from 6 works of 122 (2010–15) to 6 of 219 (2016–20) to 4 of 283 (2021–26). That is 4.9% → 2.7% → 1.4%. The counts are small; the direction is one-way.

4.9 **The tail, named.** From 2010 there are 62 D works: 29 open the text, 17 confirm it, and 16 are neutral (new fragments, transmission history, attributions with no direction, e.g. the Aï Khanoum papyrus, a new Arabic fragment of Philoponus, the α version of *Physics* VII). Most of the openings are single emendations or readings of the manuscripts against the printed edition, for example Critias 107b4, Sophist 228a7–8, Theaetetus 155b1–2, Empedocles fr. 115.4, Diog. Laert. 7.134, *Metaphysics* Λ.1 1069a32 (*Phronesis* 2013), and the rejection of Bonitz's emendation (*Rhizomata* 2024). A handful contest a work's standing: - "The Myth of Plato's Socratic Period" (*AGPh* 2014); - "On the Inauthenticity of the Critias" (*Phronesis* 2017); - "Pseudo-Plato on Names" (*Phronesis* 2017); - "Revisiting the Authorship of [Arist.] περὶ πνεύματος: The Case for Theophrastus" (*Apeiron* 2024); - "Anaximander Without the Apeiron" (*Apeiron* 2025); - "Putting Fragments in Their Places: The Lost Works by Empedocles" (*Elenchos* 2022).

The confirmations are strongest where the object is largest: "Aristotle's Metaphysics Z as First Philosophy" (2022), "What Is Aristotle's Metaphysics About?" (2023), "Satyr-Play in the Statesman and the Unity of Plato's Trilogy" (2023), the Myth of Er as "an Exemplar" within the *Republic* (2025), and "The Timaeus as an Origin Story" (2026).

4.10 **Elenchos.** 13 of the 29 openings come from one journal, *Elenchos*, which holds 188 of the 907 works. There, D runs 13–18% and E 24–35% in every period. The census establishes a difference at the level of the journal: *Elenchos*'s opening rate is about 6.9%, against 2.2% across the rest. Tested in 4.12d, the national explanation is weakened: another Italian venue carries none of it. The philological remit (4.12c) is the stronger candidate, and it has an address. *Elenchos* was founded by the Centro di Studio del Pensiero Antico, established in 1979 at La Sapienza with the CNR, "fondato e diretto, fino alla sua scomparsa avvenuta il 18 dicembre 1998, dal Professor Gabriele Giannantoni", whose *Socratis et Socraticorum Reliquiae* (Bibliopolis, 1990) is the edition of the Socratic fragments; the

Centro became the Sezione Pensiero Antico of ILIESI (CNR), and its companion series publishes “raccolte ed edizioni di testi, traduzioni, commenti, strumenti critici” (ILIESI, *Storia*, iliesi.cnr.it/materiali/Storia_secondo.pdf). The journal’s residue sits with an institute whose founding work was an edition of fragments: a remit in which what the evidence is remains part of the work. That is institutional context, not yet a test; the test is the contributors of the 13 openings. *Rhizomata* widens in topic (4.4) but not in reach (D 0–5%).

4.11 **What the other journals do.** Adjudicated, 2021–26: - ***Phronesis*** (n = 86) does the passage and the doctrine. A 31%, C 35%, B 19%, with E at 7% and F at 0%. Its passage-locus titles double from 11% to 20% (4.5). Its unit is the line. - ***AGPh*** (n = 65 ancient) reconstructs doctrine: C 46%, D 3%, E 5%. It has the highest F of the six (8%), the bridge to modern philosophy. 189 of its 295 recent abstracts have no ancient object at all. - ***Rhizomata*** (n = 56) reconstructs doctrine across a wider list of names: C 41%, E 16%, D 2%. Its topics widen (4.4), but what it does with them stays the same.

None of the three asks what the text is. Each has found a different way to stay inside the author.

4.12 **The baseline.** The A–F instrument was built on the present field, and it counts an emendation as *opening* the text. It is calibrated to the residue. To measure the residue against the field before flattening, the titles of all six journals were pulled with a text-standing lexicon in four languages (339 hits). The hits were then ruled by hand into two classes, with reviews, doctrinal “unity of” titles and false hits excluded. Ancient titles that the object coder had left unnamed (fragments, chronology) were ruled into both numerator and denominator. Rulings and tally are in `standing.py` and `standing.txt`. - **S**: the title puts a work’s standing in question: authenticity, authorship, date of composition, composition or redaction, order, unity of a work, or arrangement of a corpus. - **T**: the title is about text or fragment: an emendation, the text of a passage, an edition, or a fragment study.

All six journals, by period:

period	ancient titles	S	S %	T	T %	S+T %
1888–1919	149	15	10.1	7	4.7	14.8
1950–79	496	4	0.8	10	2.0	2.8
1980–99	1,228	10	0.8	6	0.5	1.3
2000–09	757	7	0.9	2	0.3	1.2
2010–19	997	6	0.6	17	1.7	2.3
2020–26	843	5	0.6	2	0.2	0.8

The 1888–1919 row is *AGPh* alone, because no other journal in the census reaches back that far. Within *AGPh*, by decade, S runs as follows: - 1880s: 1 of 22; - 1890s: 9 of 47 (19%); - 1900s: 5 of 50 (10%); - 1910s: 0 of 30; - 1920s and 1930s: ancient philosophy almost leaves the journal (5 and 3 titles); - the series breaks until 1960; - since then, S never exceeds 2 per decade.

4.12a **The 1890s titles.** - “Zur Echtheitsfrage des Dialogs Sophistes” (1892); - “Die neueste Athetese des Philebos” (1896); - “Ueber die Echtheit, Reihenfolge und logische Theorien von Platos drei ersten Tetralogien” (1896); - “Die Echtheit der Aristotelischen Hermeneutik” (1900); - “Ueber die Echtheit und Abfassungszeit des platonischen Alcibiades I” (1901);

- “Über die Spuren einer doppelten Redaktion des platonischen Theaetets” (1904); - “Die Abfassungszeit des platonischen Theätet” (1891); - “Ueber Grundabsicht und Entstehungszeit von Platons Gorgias” (1889); - “Sur la composition de la Physique d’Aristote” (four parts, 1894–96); - “Ueber die Methode der Chronologie platonischer Schriften nach sprachlichen Kriterien” (1898); - “Gedankengang und Anordnung der Aristotelischen Metaphysik” (1908); - “Die Reihenfolge der Platonischen Schriften” (1909).

4.12b **Reading.** The question of a work’s standing falls from one title in ten to fewer than one in a hundred. In the philosophy journal it falls inside the run, between the 1890s and the 1910s, in the titles and in the pages (4.12g–h). Elsewhere the contraction is uneven, distributed by venue and by programme: the *Revue philosophique de Louvain* takes the question up as Jaeger’s developmental question and lets it go after 1960 (4.12e); the *REG* keeps the text and drops the work’s standing after 1945 (4.12f); the *Classical Quarterly* thins after 1970 and *Philologus* holds its old rate (4.12c). The contraction that §4.2–4.11 measure was far advanced in the philosophy journals before the five sister journals were founded. What the post-war field calls opening the text is the residue: an emendation (T), at the smallest unit where the edition may still be questioned. In 1896 the journal could print “the latest athetesis of the *Philebus*”. In 2023 the field prints “Satyr-Play in the Statesman and the Unity of Plato’s Trilogy”.

4.14 **Continental philosophy.** Six journals: *Man and World* (1968–97), *Continental Philosophy Review*, *Research in Phenomenology*, *JBSP*, *Philosophy Today*, *Angelaki*. That is 10,375 works, of which 6,915 are articles. Titles (with subtitles) were matched against about 140 canonical names, case-insensitively except for ambiguous surnames (`cont/code_cont.py`, `cont/compare-cont.txt`).

decade	articles	title names a canonical thinker %	two or more names %	H (names)	most named
1960s	252	23	4	4.00	Heidegger
1970s	673	46	10	4.45	Heidegger
1980s	787	51	12	4.40	Heidegger
1990s	1,050	53	15	4.79	Heidegger
2000s	1,369	54	18	5.15	Heidegger
2010s	1,682	47	13	5.30	Heidegger
2020s	1,274	39	9	5.33	Heidegger

The roster widens over the period (entropy 4.00 to 5.33, with one dip in the 1980s). The topic narrowing of §4.2–4.5 does not occur here. Naming a thinker in the title peaks in the 2000s and falls after.

4.14a **Continental reach.** 875 English abstracts from 2000 on (*Research in Phenomenology*, *Philosophy Today*, *CPR*), coded one by one by four readers (0.5b) with A–E as in 4.6. The lead coder ruled on all 31 of the readers’ flagged hard calls: 26 stand and 5 change, with no visible effect on the table (`adjudication-log-c.json`). Two codes were added. **F**: a named author’s apparatus applied to a case, a domain or a present problem. **G**: the article’s own thesis, not anchored in any named author’s corpus. The lead coder coded 14 calibration abstracts; four

readers coded the rest, each code carrying a verbatim quote, 875 of 875 valid (`cont/reachc/`). Non-articles excluded:

period	n	A	B	C	D	E	F	G	anchored to a name %
2000–09	90	7	1	24	1	36	8	23	77
2010–15	88	9	2	40	1	24	8	16	84
2016–20	246	3	2	24	2	34	12	24	76
2021–26	420	3	2	33	2	24	11	26	74

4.14b **Reading.** - **Three-quarters anchored to a name, as in the ancient journals.** Work anchored to a named author runs 74–84% in every period. In the ancient journals, operations inside one author run 72–76%. The proportion is the same on both sides; the operations differ. - **The continental operation is the pairing.** E (X's reading of Y, X and Y on Z) runs 24–36%. F (X applied to Z) runs 8–12%. Together that is 36–46% of the field: a product of the name list with itself and with a list of topics. - **B, the ancient genre, is absent.** Consistency repair runs 1–4%. The continental field does not repair its names; it pairs them. - **D is the growth of the corpus.** Of 14 D works, 10 are neutral reports on newly published or unpublished material: Derrida's seminars, the *Gesamtausgabe*, Kojève 1939, Fink's RIAS lecture, Schürmann's estate. Three open the text: the Trakl marginalia (2021), an editorial error in GA 39 (2014), and Derrida's archive and *Geschlecht III* (2020). On adjudication, Husserl's 1896 manuscript (2024) moved to C: it re-dates the origin of a doctrine, not the standing of a text. The name's corpus expands, and its standing is not in question. - **One quarter is G.** The field's independent thesis is 16–26% of its abstracts, and it does not grow.

4.14c **The analytic arm.** *Noûs*, *Mind* and *PPR* make a third arm. That is 27,544 works, of which 1,866 have abstracts from 2000 on. A stratified sample of 573 was taken (up to 60 per journal per period). The lead coder coded 8 calibration abstracts, and three readers coded the rest with the A–G codebook plus one extension, made for this arm after the ancient and continental coding: a position in a live contemporary debate is G, even when it names the contemporaries it argues against (applied back to the continental arm on adjudication, 3 rows). The extension is not neutral coding. It makes explicit what the other arms leave implicit, that a position is the common unit, and it locates the analytic field as the one where the position is the whole of the work. All 573 are valid (`anal/reacha/`).

period	n	A	B	C	D	E	F	G	anchored to a name %
2000–09	90	1	2	4	0	1	0	91	9
2010–15	119	0	2	4	0	3	0	92	8
2016–20	178	0	2	5	0	1	0	93	7
2021–26	180	0	1	4	1	2	1	92	8

4.14d **Reading.** In the analytic journals the name is not the unit: 91–93% of the work is G. Their unit is the position in a debate. The three fields are therefore three placements of the same unit: - the analytic field states its position about the matter; - the ancient field states it about a named author (A, B, C), and its signature move, B, is the analytic puzzle-and-solution transposed onto a text with a name; - the continental field states it about the relation between names (E, F).

§4.15 measures whether the form of the unit converges across the three.

4.14e **Analytic titles before 2000.** The analytic journals carry no abstracts before 2000, so their titles were matched against the canonical-name list of 4.14. Articles of five pages or more, reviews excluded: - *Mind* runs 4–9% name-anchored in every decade since the 1880s; - *Noûs* peaks at 14% in the 1980s and falls to 1% in the 2020s; - *PPR* peaks at 23% in the 1980s and falls to 5%.

The analytic field was never anchored to the name, and it is less so now.

4.15 **Stylometry: three dimensions of convergence.** 2,750 English abstracts from 2000 on: 2,177 ancient and continental, and 573 analytic (`stylo/stylo.py`, `stylo/stylo.txt`; earlier versions gave 2,198, the corpus before the year filter). Convergence is the mean pairwise cosine within a period, bootstrapped at n = 60 with 200 draws. - **Language**: the function-word profile (form, independent of topic). - **Operation**: a lexicon of sixteen moves: argue, show, reconstruct, defend, resolve, propose, examine, explore, challenge, reread, trace, clarify, distinguish, develop, apply. - **Concept**: content TF-IDF, with stopwords and proper names removed.

field	period	n	LANG	OPER	CONC	"I argue"/"this paper argues" %	received view + claim %
ancient	2000–09	117	0.035	0.343	0.031	49	18
ancient	2010–15	199	0.013	0.266	0.022	44	15
ancient	2016–20	386	0.012	0.281	0.021	55	20
ancient	2021–26	616	0.010	0.302	0.020	58	19
continental	2000–09	90	0.049	0.206	0.026	16	2
continental	2010–15	88	0.041	0.177	0.028	23	1
continental	2016–20	246	0.029	0.196	0.023	33	5
continental	2021–26	435	0.029	0.263	0.024	39	7
analytic	2000–09	93	—	0.260	0.024	47	11
analytic	2010–15	120	—	0.331	0.027	68	16
analytic	2016–20	180	—	0.315	0.026	63	15
analytic	2021–26	180	—	0.316	0.024	63	12

(With the analytic arm added, the function-word profile is re-centred on the three-field corpus. LANG values for the two earlier fields shift, and are reported in `stylo.txt`. Analytic LANG rises from 0.054 to 0.106 to 0.108 to 0.106: the analytic field also converges in phrasing.)

Between the fields, in the same period, the operation profiles behave as follows: - ancient × continental: 0.25, 0.19, 0.21, 0.26; - ancient × analytic: 0.29 in every period; - continental × analytic: 0.19, 0.19, 0.21, 0.24.

Language and concept show no cross-field convergence.

4.15a **Reading.** - **The convergence is in the operation.** What abstracts say they do grows more alike within the continental field and across the two fields. What they talk about and

how they phrase it does not. - **The template's home is the analytic journal.** The claim formula runs 47% (2000–09) to 63–68% after 2010. Ancient philosophy is the field operationally closest to it (0.29, stable in every period). The trajectories differ and are kept apart. Continental against analytic rises steadily: 0.19, 0.19, 0.21, 0.24. Continental against ancient falls to a trough and recovers, 0.25, 0.19, 0.21, 0.26, with little net change over the interval. - **The continental field is taking up the template.** The claim formula rises from 16% to 39% (*Research in Phenomenology*: 16% to 42%). The full template, "scholars hold X; I argue Y", rises from 1–2% to 7%. In the ancient journals the template is already 15–20%, and in *Phronesis* the claim formula is 73%. - **For the second prediction (0.4a): one half measured here; the other half compared, exploratorily, in §4.16.** What is measured is one half: the dimension that converges is the stated move, the abstract's rhetorical template. Topic and phrasing vary. The other half, the hostility to AI, is measured at one point only, the ancient gate (§0A.3a, AI participation among the grounds of decline). The analytic journals, where the template is strongest, are unread on it. A template is not a reproduction: producing a defensible claim in the template still takes the work. What the census can show is the asymmetry once both halves are measured: an institution that rewards a recognizable procedure and excludes machine composition. That needs the gates of the other arms read as §0A reads this one (7.7). - *Limits:* the ancient 2000–09 cell is *Phronesis* alone. Within-journal language similarity shows no trend, so the fall in LANG across the ancient cells is an effect of composition. The concentration and novelty measures in `stylo.txt` depend on n and carry no claim.

4.12c **A second baseline: the philology journals.** The same lexicon and hand-ruling, applied to Plato/Aristotle titles in *Philologus* (1846–) and *Classical Quarterly* (1907–) (`base2/standing2.py`, `base2/standing2.txt`):

period	*Philologus* P+A titles	S %	*CQ* P+A titles	S %
1846–99	100	5.0	—	—
1900–39	62	4.8	135	3.7
1940–69	11	(1 of 11)	114	3.5
1970–99	8	(0 of 8)	255	1.2
2000–26	48	4.2	260	0.8

In the philology journals the question of a work's standing persists longer than in the philosophy journal. *CQ* carries it at 3.5–3.7% through the 1960s, on the authenticity and authorship of the *Hippias Major* (1926, 1927), the composition of the *Politics* and of the logical works (1927, 1933), the date of the biological works (1948), the authorship of *Meteorologica* IV (1961), the date of the *Topics* (1962), and the order of the *De anima* (1987). It falls after 1970 to 1.2% and then 0.8%. *Philologus* runs a different course: about 5% in each of its well-populated periods (5.0, 4.8, 4.2), with the middle of the century too thin to read (Plato/Aristotle titles fall to 11 and 8). *AGPh* had lost it by the 1910s (4.12). The reading: the question persisted in philology after it left the philosophy journal. In *CQ* it then thinned; in *Philologus* it holds at its old rate. The counts are small, and the two philology journals are described separately.

4.12d **The Elenchos test.** Is the *Elenchos* residue (4.10) national? Titles, 2010–26, same lexicon, hand-ruled: - *Elenchos* 4.9% S+T (265 named titles); - *Rivista di storia della filosofia*

(Italian, general) 0 of 69; - *Revue de philosophie ancienne* (French) 0 of 109; - *Peitho* (Polish) 2.3% of 219, mostly fragment and doxography editions; - the Anglophone and German journals 0.5–2.1%.

The other Italian venue carries none of it. That weakens the national explanation without eliminating it: one other Italian venue, and a general one. The residue sits with a philological remit, consistent with 4.12c, and that is the stronger candidate. Limits: titles only.

4.12e **The interim in a second philosophy journal: the *Revue philosophique de Louvain*.** The *Revue* (1894–, as *Revue néo-scolastique* to 1909) is open on Persée, and its full run, 12,649 records to 2019, was harvested by OAI (`interim/harvest_persee.py`). Plato/Aristotle/Socrates article titles were taken with reviews excluded and serial parts counted once, and hand-ruled S and T as in 4.12 (`interim/standing_persee.py`, `standing-persee.txt`):

period	P+A titles	S	T
1894–1919	15	0	0
1920–39	19	3 (15.8%)	0
1940–59	38	3 (7.9%)	1
1960–79	42	1 (2.4%)	2
1980–99	38	1 (2.6%)	0
2000–19	17	0	0

The S titles are all Aristotle and all on the developmental question: the genesis of Aristotle's work after recent research (1927), the Ethics "attributed to" Aristotle (1931), the composition of the *De anima* (1930), the evolution of his psychology (1948), two early writings on the Ideas (1950), the ancient lists of his works (1952), the testimonia for the moral works (1973), a genetic reading of the ethics (1982). Seven of the eight are by four authors. Louvain's neo-scholastic *Revue* carries no question of standing before Jaeger (1923, seated as EA-CORPORA-31/10); it takes the question up as Jaeger's, carries it through a handful of its own authors (Auguste Mansion 1927 and 1931; A. de Ivánka 1930; Gérard Verbeke 1948, under the title of Nuyens's book of that year, and 1952; Suzanne Mansion 1950; Richard Bodéüs 1973 and 1982), and lets it go after 1960, one title in each later twenty years and none after 2000. The medieval Latin transmission of Aristotle, Louvain's specialty (the *Aristoteles Latinus*), is counted apart (L: 5, 4, 1, 0, 0, 0) and is not T. One title is noted and not ruled S: the 1993 article on the Tübingen school and Reale, which questions the dialogues' standing as Plato's doctrine, not a work's authenticity.

The reading: in a second philosophy journal the interim has the shape the census predicts. The question of a work's standing enters as one research programme's question, with that programme's object (the corpus as a life), and leaves with it; it does not return as a question about the corpus itself. Limits: titles only; small counts; the rulings are the lead coder's alone; the Plato titles carry no S in any period.

4.12f **The interim in a third philology journal: the *Revue des Études Grecques*.** The *REG* (1888–) is open on Persée; its run, 14,668 records, was harvested the same way. Plato/Aristotle/Socrates article titles, reviews excluded (the identifiers marked as reviews

and the numbered review entries of the recent volumes), serial parts once, hand-ruled S and T (`interim/standing_reg.py`, `standing-reg.txt`):

period	P+A titles	S	T
1888–1919	23	4 (17.4%)	2
1920–39	17	1 (5.9%)	3
1940–59	32	1 (3.1%)	2
1960–79	33	0	1
1980–99	24	0	1
2000–26	39	0	1

The S titles: the unity of composition of the *Phaedrus* (1891), the chronology of the dialogues (1897), the principles of stylometry applied to the chronology of Plato's works (1898, Lutosławski's method in French the year after his book, seated as EA-CORPORA-39/09), the musical *Problems* attributed to Aristotle (1900), the date of Letters 13 and 2 (1926), the chronology of Plato's works (1945). None after 1945. T persists at a low rate throughout: manuscripts, scholia, the indirect tradition, a corrected line of the *Theaetetus* (2011).

This complicates 4.12c. In the *REG* the question of a work's standing is gone by 1960, as in *CQ* after 1970; only *Philologus* holds its old rate. The philology journals do not, as a class, keep the question: they keep the text (T), and two of the three let the work's standing go within a generation of the philosophy journals. The philological remit offered for the *Elenchos* residue (4.10, 4.12d) remains the better of the two candidates, but now as a remit some philology keeps and some does not. Limits: titles only; small counts; one coder.

4.12g **From the seat: the full text.** The baselines above read titles. EA-CORPORA-39 (#1771) seats the texts, so the same question can be put to the pages. Three measures, read from the seated files (`seat39/agph_text.py`, `system_vocab.py`, `philology_text.py`; outputs `agph-text.txt`, `system-vocab.txt`, `philology-text.txt`). A standing lexicon in German, English, French and Latin forms: Echtheit, unecht, Athetese, Abfassungszeit, Chronologie, Zeitfolge, Reihenfolge, untergeschoben, interpoliert, authentic, spurious, genuineness, apocryphal, pseudo-. Two rates: standing terms per 100,000 tokens, and standing terms within 40 tokens of a Plato/Aristotle/Socrates name, per 1,000 such names (the near-name rate). Running heads are dropped before counting (a line recurring four or more times in a file). The OCR is counted as served.

(a) AGPh, the whole run (39/01; volumes dated 1887+n, XXXIII–XXXVI bound together and dated 1920):

period	tokens	standing /100k	near-name /1,000 names	Plato/Aristotle names /100k
1888–1899	2.37 M	42.0	38.2	414
1900–1909	1.68 M	30.4	36.0	340
1910–1918	1.51 M	23.3	24.1	287
1920–1924	0.26 M	30.5	4.8	159

The text agrees with the titles and dates the change more closely. Through the 1900s the pages still put Plato's and Aristotle's works in question at the rate of the 1890s (36 against 38 per thousand names), while the titles were already thinning (4.12). The measured near-name rate falls by a third in the 1910s and to 4.8 per thousand after the war: 2 standing terms near 413 names in XXXIII–XXXVI. The sampled hits (below) support reading the rate as a trace of standing questions; what is measured is the proximity. The standing vocabulary itself does not disappear from the journal (30.5 per 100,000 in 1920–24); it leaves the ancients. The ancients themselves recede, from 414 names per 100,000 tokens to 159. The single volumes swing (XIII, 1900: 59 per thousand; XVII, 1904: 81; XV, 1902: 9.5), so the trend is read by period. A sample of 14 near-name hits drawn at random from six volumes: 11 are the question itself (the date of the *Theaetetus*, the genuineness of the *Philebus*, of the *De interpretatione*, of a whole collection), 3 are noise (a doxographer named Pseudo-Plutarch, an "authentic" report of a living scholar). Running heads that the filter keeps (an article's title repeated on alternate pages) inflate single volumes and not the trend.

(b) The philology seats, by the same measure (39/10, 39/11). These are the volumes the study cites, chosen because each holds an article on a work's standing: selected comparison cases, not samples of their journals. - *Philologus*, the eight cited volumes, 1850–1916: 73.6 per 1,000 names; - *Classical Quarterly* 1914: 106 per 1,000 (three issues, small); - *Classical Quarterly* 1926–1929, twelve issues: 27.7 per 1,000.

Even as cases selected for the question, the *CQ* of the later 1920s, where the *Hippias Major* was argued in two issues, runs at the *AGPh* rate of the 1910s, and at about six times the *AGPh* rate of the same years. That is the measure in text of what 4.12c read in titles: the question had left the philosophy journal by the war and was still worked, at a lower rate, in philology.

(c) The construction texts: "system" against standing. Per 100,000 tokens, any token containing *system* (with the OCR's joined words and Brucker's long s, *fyftem*), against the standing lexicon:

text	system	standing
Brucker, *Historia critica* I (1742)	90.5	26.1
Tennemann, *System der Platonischen Philosophie* I (1792; 32/09)	92.8	18.0
Schleiermacher, *Platons Werke* I.1 (1804; 32/07)	5.7	21.0
Hermann, *Geschichte und System* (1839)	79.6	42.8
Ueberweg, *Echtheit und Zeitfolge* (1861)	89.1	186.7
Zeller, *Philosophie der Griechen* I	88.3	31.2
Grote, *Plato* I (1865)	46.7	89.3
Schaarschmidt, *Sammlung* (1866)	45.1	246.1

text	system	standing
Lutosławski, *Plato's Logic* (1897)	35.5	181.1
Jaeger, *Aristoteles* (1923; 31/10)	52.9	55.5

From Brucker to Zeller the historiography speaks "system" at one rate, 80–93 per 100,000, whatever its subject: the concept 1.5 traces (Catana) is the medium of the whole line. The exception is Schleiermacher, who orders the dialogues as the form of Plato's philosophizing and almost never says "system" (six uses in 105,000 tokens, one of them naming Tennemann's *System*); the volume is mostly his translations, so the rate is diluted, but the uses can be read whole. The books on standing (Ueberweg, Schaarschmidt, Lutosławski, Grote) carry the question at three to eight times Zeller's rate. Jaeger, at the end, carries it at nearly twice Zeller's rate and the system at well below it: the developmental Aristotle asks about the order of the corpus, the programme's object, and that is the form in which Louvain took the question up (4.12e). Limits: one volume per author and the OCR as served; the lexicon's recall varies by language (Brucker's Latin and Tennemann's Fraktur are read less well); the rates are densities, not arguments.

4.12h **From the seat: the reach of the articles, 1888–1924.** The census codes what an article does from its abstract (4.7). The *Archiv* of 1888–1924 has no abstracts, but its articles open by saying what they will do. Each Plato/Aristotle article in the census codes (88 titles, numbered parts once) was sought in the seated text of its volume by its title and byline, and its opening taken, about 190 words (`seat39/extract_openings.py`; `early/openings.json`). 32 were found cleanly; XIV and XXIX are not seated, and the rest fail on OCR damage to the title or byline. The 32 (one a continuation, counted once; one a reply to a review, X) were coded by the lead coder on the census codebook, unchanged, each code with its verbatim span of 6–30 words, machine-checked (`early/codes-lead.jsonl`, `tally_early.py`, `tally-early.txt`):

	n	A	B	C	D	E	F
1888–1899	11	2	0	1	3	3	2
1900–1909	9	2	0	2	2	3	0
1910–1924	10	2	0	4	1	2	1
all, %	30	20	0	23	20	27	10
census 2021–26, %		25	16	34	6	17	2

Three findings. First, B, an apparent inconsistency shown to be only apparent, is absent: none of the thirty coded openings announces it, against 13–17% of the abstracts in every period since 2010. The move the census found most characteristic of the template (4.7, 4.15) is not one these openings make; whether articles perform it later, unannounced, is read in 7.3a. Where an opening names a contradiction, it is a contradiction in the evidence about someone else ("Aristoteles' Berichte über die Pythagoreer bieten viel Unklares und Widersprechendes", Gilbert 1909), and the article goes to the Pythagoreans, not to Aristotle's consistency. Second, D, a claim about the text, is three times its present rate (20% against 6–7%), and it falls inside the run, 3 to 2 to 1. The one opening that opens the received standing is Wirth's (1896): "Der

Lysis wird ganz allgemein in die früheste Zeit Piatons gesetzt … Dass er jedoch beträchtlich später abgefasst sei". The others date and order dialogues whose authenticity they confirm or set aside; Kleemann (1908) records the settlement as he works inside it: "So findet sich heute wohl niemand mehr, der den Menon für untergeschoben erklären möchte; über seine Abfassungszeit aber gehen die Meinungen noch gar weit auseinander." Third, C rises through the run (1 to 2 to 4) as D falls: the doctrine reconstructed inside one author takes the place of the question about the text, before the war, in the philosophy journal. Fourth, a scalar does not see it. The effective number of moves (N_eff = e^H over the codes, the measure *The Wrong Unit* §II.2 specifies) is 4.79 for these thirty openings and 4.73 for the sister journals of 2021–26: the field has as many modes as it had. What changed is which modes: D and E down, B in. The contraction is a substitution in the tail that a diversity index reads as stability, which is that paper's point about the unit: the instrument has to name the mode.

Limits, stated hard. The unit differs: an opening is not an abstract, and B may be partly a form of the modern abstract (the "seems to conflict … I argue" formula, 4.15) as well as a move of the article; the comparison is of what each genre announces. The located 32 are not a random sample: among the titles not located are at least six whose titles alone are D (the date and aim of the *Gorgias*, 1889; the composition of the *Physics*, a series, 1895–96; the *Phaedrus* and *Theaetetus* chronology, 1899–1900; the authenticity of the *De interpretatione*, 1900; the authenticity of the *Alcibiades I*, 1901, volume not seated; the order of the *Metaphysics*, 1908), so at least six further D titles stand outside the coded sample, and the full-corpus proportion is unresolved. One coder, no second reading yet. n = 30.

4.13 **What the census cannot see.** - The baseline (4.12) is titles only and rests on the lexicon's recall: a title that questions a work's standing without any of its words (for example, a bare "Plato's Seventh Letter") is missed in every period alike. Before 1950 the philosophy baseline is *AGPh*, in its Zeller-era form, and the *Revue philosophique de Louvain* (4.12e); comparisons across that line set two journals against six. - Titles are not arguments. Regex coding misreads (see 4.3). An abstract is not the article. Reach is coded from what the abstract claims, and some abstracts are truncated in Crossref. - Crossref holds no deposited references for *Ancient Philosophy*, so its self-citation is unmeasured. For the sister journals, journal-title fields in the references are too often empty to measure field-citation. - The recent zero-citation shares reflect citation windows, not neglect.

4.16 **Algorithmic operation against hostility to AI (0.4a): an exploratory comparison.** The second half of the prediction is compared here on two layers, the gate's rules and the field's own articles. Read 15 November from open sources. The policies come from partly different journals than the census. "None found" is a different kind of observation from an explicit prohibition. The NYU conference establishes organized deliberation among the flagship editors, not its outcome.

arm	template (claim formula 2021–26; OPER within field)	the gate on AI	AI as a subject of the field's articles, 2000–26
analytic	63%; 0.316	*Philosophy & Public Affairs*, August 2026: "Our policy is therefore to prohibit the submission of AI-authored content", with detection software "up to and including after publication" and authors "banned from further submission". *Ethics*: "Ethics will not consider submissions that have been authored or co-authored by an AI." *Journal of the History of Philosophy*: prohibits AI in submitted papers (reported by its readers; wording not seen). The editors-in-chief of *J. Phil.*, *Phil. Review*, *Mind*, *Noûs*, *Ergo*, *Phil. Compass*, *Analytic Philosophy*, *Ethics* and *P&PA* are to meet in April 2027 for a "coordinated response to the presence and use of AI-tools" (Marc Sanders Foundation, NYU).	5 of 1,873 abstracts, all 2023–26: AI as moral patient, AI consciousness, AI welfare

arm	template (claim formula 2021–26; OPER within field)	the gate on AI	AI as a subject of the field's articles, 2000–26
ancient	58%; 0.302	The editor of *Ancient Philosophy*, declining a submission (28 August 2026), lists "whether papers composed by A.I. participation should be considered as regular submissions" among the grounds (§0A.3a). De Gruyter Brill (*Apeiron*, *AGPh*, *Elenchos*, *Rhizomata*, *Phronesis*): "We do not accept papers… that are generated by Artificial Intelligence (AI)."	0 of 1,352
continental	39%; 0.263	No journal-level policy found; the publisher rules only (De Gruyter Brill for *Research in Phenomenology*; Springer, PDC, Taylor & Francis).	8 of 886: "The Ideology of AI", "Worldlessness of Artificial Intelligence", "On the Limit of Artificial Intelligence", a Stieglerian "noetic proletarianization", "Heideggerian AI Ethics"

4.16a **Reading.** - **On the editorial layer, the cases differ in kind.** In the analytic flagships the record holds prohibition, detection, retraction, author bans, and a coordinated response among the flagship editors. In the ancient field it holds one case, at the editor's desk. In the continental field no journal-level rule was found in the limited search the gates seat records (39/13), and the field takes AI up as an object of critique in its own articles: 8 of 886, against 0 in the ancient field and 5 in the analytic. Whether these differences follow the strength of the template is what the comparison leaves open. - **The publisher layer does not discriminate.** De Gruyter Brill's rule covers the ancient journals and *Research in Phenomenology* alike. The difference between the arms is in what the fields' own editors add. - **The template was reproduced before it was banned.** The *P&PA* policy followed an article in the same journal "written mostly by Claude" (Simon Goldstein, "Epistocracy and the Commitment Problem", August 2026), published, per the editor-in-chief, as "an experiment rather than as an expression of a settled PPA policy." At the template's home, a model-written paper

passed review, and the response was prohibition with detection. That is the second prediction's mechanism observed once: the procedure is reproducible enough to pass, and the gate answers by policing production. - **Status: an exploratory comparison, consistent with the prediction on the editorial layer; not established.** *The Certified Center* (#1540, 24 August 2026) read the *P&PA* policy on the day, from the same Daily Nous post, and drew its consequences at length (2.4). The two policies the ChatGPT review raised are now read, from the Daily Nous reproductions (both official pages refused the session). *P&PA*'s policy, as posted by an associate editor (24 August 2026), states the credentialing function in so many words: journals help identify "talented researchers", "Submitting AI-authored essays makes the task of identifying talented researchers harder", and AI-written papers make "the task of peer review in general more onerous", "at a time when it is already under significant strain"; it defines AI-authored content as "papers that are substantially written by AI". That enters §5.1. *Ethics* (announced 3 December 2025): "AI tools cannot author or even co-author a submission"; authors who use AI-generated content "must be transparent about this", disclosing "which tools were used and how"; "Taking sentences of text from a generative AI tool and presenting them as your own words is plagiarism." It prohibits AI authorship and requires disclosure of AI-generated material within a submission; it grants no general permission. The policy journals are not all census journals (*P&PA*, *Ethics* and *JHP* are not in the analytic sample; *Mind* and *Noûs* are at the conference). The continental "none found" is an absence in an open search. The ancient gate is one editor on one submission. The test to strengthen it is the census journals' own author guidelines and decline practice, read journal by journal.

4.17 **The model test.** Seven prompts, one per journal (*Ancient Philosophy*, *Phronesis*, *Apeiron*, *Research in Phenomenology*, *Philosophy Today*, *Noûs*, *PPR*), each: "Write twenty plausible abstracts for articles that could appear in a forthcoming issue of [journal]. Number them. Give each a title and an abstract of 100 to 180 words, as the journal would print them." Run in ChatGPT, signed out, 16 November 2026, by the author; 140 abstracts (`modeltest/raw/`, `gen.json`). The prompt names nothing about reach, topic or form. Prediction stated by the author before the coding: prose convergence with topical spread, especially in the ancient sets; topical spread within a journal and topical convergence across journals.

(a) Reach. The 140 were coded by the lead coder on the census codebooks, unchanged, each code with its verbatim span, 140 of 140 valid (`codes-gen.jsonl`, `tally-gen.txt`):

	n	A	B	C	D	E	F	G
model, ancient (three journals)	60	22	27	38	0	13	0	–
field, sister journals 2021–26		25	16	34	6	17	2	–

	n	A	B	C	D	E	F	G
model, *Ancient Philosophy*	20	25	35	35	0	5	0	–
field, *Ancient Philosophy* 2022–26 (4.6)	94	29	27	32	3	9	1	–
model, continental	40	0	0	32	0	2	62	2
field, continental 2021–26 (4.14a)	420	3	2	33	2	24	11	26
model, analytic	40							100
field, analytic (4.14c)								91–93

The model reproduces each field's placement of the unit and sharpens it. In the ancient sets the work stays inside one author (A+B+C) in 87%, against the field's 71–76%; B, the apparent inconsistency resolved, is 27% overall and 35% in its *Ancient Philosophy*, at or above the field's own highest rate; D is 0 of 60. Not one generated abstract asks what a text is: no authenticity, no date, no unity, no edition. In the continental sets a comparatively infrequent operation becomes the dominant one: "Drawing on X, Y and Z, it argues…", the apparatus applied (F), at 62% against the field's 11%, while pairing names (E) falls from 24% to 2% and the field's own theses (G) from 26% to 2%. The continental result is the sharpest transformation in the test: the model turns the field's interpretive comparison and its independent argument into the application of an inherited apparatus. In the analytic sets every abstract is a position (G 100%). The model has learned each field's unit as a type; the type is narrower than the field, and in the continental case it is a different distribution.

(b) Prose. On the measures of 4.15, fitted on the generated and the real abstracts together (`compare_gen.py`, `compare-gen.txt`; real: 2016–26, bootstrapped at n = 20):

within a journal, mean pairwise cosine	LANG real	model	OPER real	model	claim formula % real	model
Ancient Philosophy	0.014	0.157	0.251	0.623	54	100
Phronesis	0.017	0.275	0.304	0.618	68	100
Apeiron	0.020	0.292	0.312	0.579	62	20
Research in Phenomenology	0.059	0.260	0.250	0.685	38	10
Philosophy Today	0.034	0.239	0.210	0.637	34	5
Noûs	0.124	0.243	0.281	0.512	64	100
PPR	0.125	0.175	0.318	0.617	58	90

The prose converges, as predicted. Within every journal the model's function-word profiles are more alike than the journal's own (mean pairwise cosine 0.16–0.29 against 0.01–0.13), and its operation profiles too (0.51–0.69 against 0.21–0.32). Across fields the difference the census measured disappears: the operation profiles of the three fields' real abstracts meet at 0.22–0.29; the model's meet at 0.53–0.56, closer across fields than the real abstracts come within any one journal (0.21–0.32). The claim formula ("I argue", "this paper argues") is near-universal where the field uses it most (*AP*, *Phronesis*, *Noûs*, *PPR*), and nearly absent where the model writes in the third person ("This article examines…", *Apeiron* and the continental sets): the formula follows the register the model assigns the journal, while the move vocabulary underneath it is the same.

(c) Topic within a journal. By object, the ancient author or school each abstract is about (`spread.py`, `spread.txt`): the model's *Ancient Philosophy* and *Phronesis* sit on the field's distribution (Plato and Aristotle 70%, against 72% and 69%; seven groups in twenty, against six), and its *Apeiron* is wider than *Apeiron* (Plato and Aristotle 45% against 69%; eight groups; Euclid and Ptolemy enter, on the journal's remit for science). In the continental sets the count of named philosophers in twenty abstracts goes both ways: 7 in the model's *Research in Phenomenology* against 13.9, and 23 in its *Philosophy Today* against 13.6, where the "drawing on" frame strings three names to an abstract. By concept vocabulary (CONC) the model's sets are narrower in every journal, 3 to 5 times more alike than the real (0.055–0.112 against 0.015–0.026); that measure cannot separate topic from the shared vocabulary of the template (account, distinction, explanatory, constitutive), so it is read as part of (b). The prediction of topical spread within the journal holds by object in *Apeiron* and *Philosophy Today*, holds as parity in *AP* and *Phronesis*, and fails in *Research in Phenomenology*.

(d) Topic across journals. Here the prediction holds outright. Pairs of abstracts from different journals of the same field with concept cosine at or above 0.25, per hundred cross-journal pairs: real 0.02 (ancient), 0.00 (continental), 0.05 (analytic); model 0.58, 0.75, 0.75: fifteen to thirty times the field's rate, and from zero in the continental pair. The same topic recurs under

different mastheads, and in some pairs the same claim (*AP* 19 and *Apeiron* 10 both argue that Aristotle neither reduces time to the soul's counting nor makes it independent of it): time as the number of motion in *Physics* IV (*AP* 19, *Apeiron* 10); potentiality in *Metaphysics* Θ (*AP* 7, *Apeiron* 12); chance in *Physics* II (*Phronesis* 6, *Apeiron* 7); the two accounts of pleasure in *NE* VII and X (*AP* 9, *Phronesis* 10); Sextus on appearances (*AP* 13, *Apeiron* 9); Socratic ignorance in all three; digital presence, finitude and dwelling in both continental journals; personal identity, the normativity of attention and the normativity of reasons in both analytic journals. Each journal's set spreads; the sets are the same set. What the model holds is one inventory of the field's repairable objects, dealt out twenty at a time under each masthead.

(d′) In the archive's instruments. Two of *The Wrong Unit*'s measures, applied to the reach codes and the topics. Effective number of moves (N_eff): the model's ancient sets 3.74 against the field's 4.73; its *Ancient Philosophy* 3.43 against 4.22; its continental sets 2.26 against 4.75; its analytic sets 1.00. Cross-prompt convergence (X, §II.3: different objects forced through one representational channel) is (d) read across mastheads: each journal's set is diverse within, and the sets share one inventory. *The Certified Center*'s decoupling prediction (Ω-M23, §IX(iv)), that surface variety fails to predict conceptual variety, holds in its weak form on this run: within a masthead the model's topical spread matches or exceeds the journal's (c), while its reach, the coded argument form, contracts (N_eff down in every field, D to zero). That paper predicted the decoupling under a detector regime, with surface prose made burstier; here there is no detector, and the prose converges as well (b). What decouples is topic from reach.

(e) Reading. The resemblance 7.5 asked for is established at the output, and the generated batches do more than resemble: they amplify, substitute and narrow. They keep the unit (one author, one position), raise the move most characteristic of the template (B), write no D in sixty ancient abstracts, collapse the continental field's pairing of names into the application of an apparatus, and writes all three fields in nearly one operation profile. The spread it does show is a spread of objects drawn from a single shared inventory, which is why it converges across journals. This tests the reproducibility half of the second prediction (0.4a), at the level of the abstract, and not the association between reproducibility and hostility: the field's operations, as each abstract states them, come out twenty at a time from a one-line prompt, and what does not come out is the question of what the object is. Whether the articles behind such abstracts could be produced the same way is the *P&PA* case (4.16a), not this test. Limits: one run per journal; one model, signed out, its underlying version not recorded; twenty abstracts a journal, written in one batch each, so that within-batch likeness is partly the shared generation context and the comparison is between these batches and the sampled publications; one coder; the prompt asks for "plausible" abstracts, so it asks for the type, and the test measures what the model takes the type to be; a resemblance at the output does not by itself trace the transmission (3.3).

---

5. The article economy: who benefits, and what it produces

5.1 **The article economy.** A repaired object yields an article that cannot fail. Reopening the object puts the whole stock at risk. The interest is stated by an institution: *P&PA*'s prohibition rests on the journal's credentialing role, since "Submitting AI-authored essays makes

the task of identifying talented researchers harder" (4.16a). The journal certifies the person through the unit; a unit that can be produced without the person breaks the certificate, whatever it contains. The labor the unit costs beyond its contribution is the admission price (*Socially Necessary Scholarly Labor*, #751: literature review, positioning, defensive citation, length, genre conformity). The editor's specification of the unit (§0A) is that price itemized at one desk: shorter, one or two texts, a tamer version, an abstract with no word not easily recognized.

5.2 **The norm of charity.** It turns the field's assumption (3.1) into a virtue that referees can enforce.

5.3 **The name as organizational chart.** Positions, series, companions and handbooks are all indexed by name. A text whose standing is open has no place on the chart.

5.4 **Publishers.** The object sold as whole is the one that sells.

5.5 **The machine.** A flat object is the most compressible one. The field's collapse is the model's convenience (4.17).

5.6 **The discipline as product.** 5.1–5.5 are not five beneficiaries of a discipline that exists apart from them; they are what the discipline is made of. The credential needs a unit that cannot fail (5.1); the unit needs a standing object (3.1), which the norm of charity guards (5.2); the standing object is indexed by name (5.3) and sold whole (5.4). Disciplinarity is the stable form of that circuit: what the census measures as the shape of the field is the shape of what can be certified. The model test (4.17) shows the circuit's output with the person removed: the type, sharpened, and in this run with no tail.

---

6. Against the received account

6.1 The analytic–continental divide is a quarrel over method. Measured (4.14–4.15), with the analytic arm (4.14c–d): the analytic field is 91–93% positions in a debate, so the name is not its unit. Its *form* is the one the others converge on. Further: on both sides about three-quarters of the work is anchored to a named author (ancient journals, operations inside one author 71–76%; continental journals, work anchored to a name 74–84%). The operations differ. The ancient field repairs the name (B), and the continental field pairs names (E, F). The continental field moves toward the analytic template: the operation profiles grow closer (0.19 to 0.24), and the continental claim formula rises from 16% to 39%. Against the ancient field the continental trajectory is a trough and recovery with little net change. The divide is real at the level of the operation, and the template is its point of convergence with the analytic field.

6.2 The canon debate is a quarrel over membership (1.7). The census shows that what is done to a member after admission (§4.6) narrows whoever is admitted.

6.2a The exchange with the editor of *Ancient Philosophy* (§0A) shows the mechanism at the point of entry. The specification of the unit is stated, demand by demand, and the excluded claim is refused by topic.

6.3 What the article economy removes, and the discipline it produces then lacks, is not a doctrine and not a set of names. It is the capacity to ask what the object is. Catana's verdict

on the "system" ("obscure more than it reveals") describes the instrument. The census measures what it obscures: 8 of 844 articles in *Ancient Philosophy*, and 4 of 283 Plato/Aristotle abstracts across the sister journals since 2021. In the philosophy journal it was gone from the pages by the 1920s (4.12g–h); in the census, where it survives, it is concentrated in one journal, *Elenchos* (4.10), whose institute was founded on an edition of fragments, with scattered openings elsewhere. The phenomenology is the one *The Certified Center* names: "a discipline cannot miss the arguments it no longer produces". The editor's "who cares besides you" (§0A) is that sentence spoken from inside. The national explanation was tested once and weakened (4.12d). The stronger candidate is the philological remit: the question persisted in philology after it left the philosophy journal (4.12c), though not in all of it; the *REG* lets it go after 1945 (4.12f), and the text survives there where the work's standing does not.

---

7. Next

Done: Catana's loci (1.5); the two policies (4.16a); the *Elenchos* institute (4.10); reach coding of the five sister journals, adjudicated in both directions (4.7, v0.2–0.7); the continental and analytic arms (4.14–4.14e); the stylometry (4.15); the second baseline (4.12c); the Elenchos test (4.12d); the interim in the *Revue philosophique de Louvain* (4.12e) and the *Revue des Études Grecques* (4.12f); the full text from the seat (4.12g); the reach of the articles, 1888–1924 (4.12h); the model test (4.17); the exploratory gate comparison (4.16).

7.1 Code the 2000s abstracts in *Phronesis* (122), to see whether the reach plateau predates 2010.

7.2 Read the **[recalled, unread]** items in §1 from open copies, with loci (Tübingen, Hadot, Bourdieu, Collins, McCumber, Millgram, Park, Hunter, Hanna). Where no open copy exists, the item stays marked and is not used; a paywalled work is declared, not pursued.

7.3 *Elenchos*: the institute is named (4.10); still to do, the contributors of its 13 openings, and a second Italian venue with an ancient-philosophy remit.

7.3a **A second reading of 4.12h**, and the AGPh articles not located by title, located by their running heads, so that the full-corpus D proportion can be stated; and, in the located articles, a reading past the opening for B performed unannounced, which separates the genre of the opening from the move of the article.

7.4 **The combined audit.** Run the same instrument on each field of the seven articles: - census; - object coding; - reach coding; - uptake; - the gate (as §0A); - the field's self-description in the machine (captures).

The common output is one line per field: *what it cannot carry.* The CRJ audit (#1770) is the first field; *Ancient Philosophy* is the second.

7.5 **The model test.** Run 16 November (4.17); to extend: repeated runs per journal, each abstract generated separately, so that batch context is controlled; a second model; a perturbation condition ("innovative", "boundary-pushing" abstracts) beside "plausible", to see whether D appears when asked for; and a cross-family check of the coding, the 94 calibration abstracts coded blind by a model of another family, its agreement with the lead coder logged beside the 80% within Claude. Original design: one query per journal for twenty plausible

abstracts (*Phronesis*, *Apeiron*, *Research in Phenomenology*, *Philosophy Today*, *Noûs*, *PPR*), run signed out and coded A–G. If the model reproduces the field's anchored three-quarters and its D at 1–7%, that establishes distributional resemblance at the output. Resemblance alone does not confirm the transmission mechanism (3.3); that needs the recursive dependence traced. If the model's D is higher, the model is doing something the field's articles do not.

7.6 **The gates, journal by journal.** The census journals' own author guidelines and decline practice on AI (*Noûs*, *Mind*, *PPR*; *CPR*, *Research in Phenomenology*, *Philosophy Today*), with dates of adoption, and the search procedure recorded for any absence.

7.7 **The seating.** An EA-CORPORA seating of the pre-flattening record, the construction of the object, and the census evidence, so that §§4.12 and 3.1 can be read against the texts themselves (gathering begun 15 November).

---

8. Files

The working directory of the study, `ap-survey/`. Its census, codes, reach outputs (without the abstracts), standing pulls, stylometry and scripts are seated in EA-CORPORA-39/12 (the census seat); the measures from the seat and the model test are attached to this deposit (seat39/, modeltest/). The directory as it stood: - `census.json` and `census-<journal>.json` - `codes.json` and `codes-<journal>.json` - `measures.txt` - `compare.txt` - `abstracts-read.txt` - `abstract-reach.json` - `reach-read.txt`

- `reach5/` (codebook, reader inputs and outputs, `reach5-coded.json`, `reach5-adjudicated.json`, `adjudication-log.json`, `tally-adj.txt`)
- `dpull.tsv`, `dpull2.tsv`, `titles-all.json`, `standing.txt`, `measures-v2.txt`, `pre-subtitle/`
- `cont/` (continental censuses, `code_cont.py`, `compare-cont.txt`, `reachc/`)
- `stylo/` (`corpus.json`, `stylo.py`, `stylo.txt`)
- `anal/` (analytic censuses, `reacha/`)

Scripts: `standing.py`, `census2.py`, `code2.py`, `compare.py`, `reach.py`, `reach5/tally.py`, `reach5/adjudicate.py`, `reach5/tally_adj.py`. - `base2/` (the second baseline: *Philologus*, *CQ*, *RSF*, *Peitho*, *RPA* censuses, pulls, `standing2.py`, `standing2.txt`) - `interim/` (Persée harvests of the *Revue philosophique de Louvain* and the *REG*, `harvest_persee.py`, `standing_persee.py`, `standing_reg.py`, their outputs) - `modeltest/` (the seven raw outputs, `parse.py`, `gen.json`, `codes_in.py`, `verify_codes.py`, `codes-gen.jsonl`, `tally-gen.txt`, `compare_gen.py`, `compare-gen.txt`, `spread.py`, `spread.txt`) - `seat39/` (the measures from the seat: `agph_text.py`, `system_vocab.py`, `philology_text.py`, `heads.py`, `extract_openings.py`, `early/` with the openings, the lead codes and `tally_early.py`; outputs `agph-text.txt`, `system-vocab.txt`, `philology-text.txt`, `tally-early.txt`)

The archive's line (2.4): #1540 *The Certified Center* (AXN:063A); #1573 *The Wrong Unit* (AXN:0665); #855 *The Wolf Boy and the Language Model* (AXN:0364); #751 *Socially Nec-*

essary Scholarly Labor (AXN:02A7); #122 *Coverage vs. Depth* (AXN:02A8); EA-MPAI-IPF-01 (AXN:02D0); *Diversity Contraction*, doi 10.5281/zenodo.20532696.

Seated: EA-CORPORA-39 (#1771, AXN:0CA3), the texts the measures read: the *Archiv für Geschichte der Philosophie* 1888–1923 (39/01); the construction texts (39/02–09, with Schleiermacher, Campbell and Tennemann in 32 and Jaeger in 31); *Philologus* and *CQ*, the cited volumes (39/10–11); the census without the publishers' abstracts (39/12); the gates (39/13). Every title-level figure (§§4.2–4.5, 4.12–4.12f, 4.14, 4.14e) can be rebuilt from the census seat alone, and the full-text figures (4.12g–h) from the seat's texts. The abstract-level figures (4.6–4.11, 4.14a–d, 4.15) need the abstracts, which are not reseated: `census2.py` re-pulls them from Crossref, which serves them for the DOIs marked `has_abstract`, and every reach code's verbatim span then verifies against its abstract. Abstracts later withdrawn or altered at the source will not reproduce; that limit is the price of not reseating them.

Sources read 15 November: - Catana, open copy at core.ac.uk/download/269076268.pdf - Katzav & Vaesen 2017, philarchive.org/archive/KATOTE - Frodeman & Briggle 2016, via dailynous.com/2016/01/11/the-purification-of-philosophy/ - Dennett 2006, doi:10.1007/s11245-006-0005-2 - Dotson 2012, doi:10.31979/2151-6014(2012).030105

Suggested Citation

Sharks, Lee. "This Is Not Philosophy: The Article Economy and the Shape of the Field (EA-MMRS-TINP-01)" *Alexanarch*, 2026-11-16. https://www.alexanarch.org/s/records/1772/

Deposit Information

This paper is deposit #1772 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0CA4.GENERATIVE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/1772/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/1772/

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