

Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instrument They Specified, and Its First Wave — with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)

Sharks, Lee

2026-09-15

Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instrument They Specified, and Its First Wave — with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)

Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

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

2026-09-15 · Version v0.3 · Theoretical notebook with toy models, a built instrument and a pre-registered second wave

AXN:06A6.EMPIRICAL.⊗

<https://www.alexanarch.org/s/records/1616/>

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": "Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instr
with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)",
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
```

```

    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-09-15",
  "identifier": "AXN:06A6.EMPIRICAL.⊗",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1616/",
  "description": "MODEL COLLAPSE ASKED AT THE WIDEST ANGLE, AND THEN MEASURED. If the layer through which a population reads the world composes without reading, and its compositions become its sources, what collapses is not a model's distribution but the represented world — the set of distinctions a reader can still reach. Reality is defined operationally here and not metaphysically: the panel of distinctions the world carries that remain reachable through the layer. Flattening is the loss of a distinction from that set while the things distinguished still exist; collapse is flattening that compounds because the flattened composition is written back as a source. Closed is defined so that the finite-grammar objection cannot land: an ontology is closed when encounter cannot enlarge its distinction-generating apparatus or revise its typing rules.\n\nFOUR TOYS, EACH SCOPED TO WHAT IT SHOWS. Sense collapse under a sharpening prior and mode-seeking aggregation takes a string with six effective senses to one; fan-out is documented, the voting aggregator is this notebook's hypothesis. A frozen type system yields a proposition with a proof: for a world acquiring types at rate  $\lambda$  and an ontology able to represent at most  $K$  distinctly, representability is bounded by  $K/(N_0 + \lambda t)$  and goes to zero, with admission the only parameter that holds it. Compression writing, replicated over 200 seeds, is absorbing at zero exogenous floor — every seed reaches five or six distinctions by generation 30 and stays — and stationary at a level set by the floor for every floor tested. The admission sweep rejects linear preferential attachment, which cannot leave the diffuse regime, and locates the observed concentration in the bounded-admission family with within-set spread; superlinear attachment and hard and soft gates all overshoot to near-monopoly. B and C are two realizations of one closure condition, not one curve. [...abridged for the catalogue; full description in this deposit's record]
}

```

Abstract

MODEL COLLAPSE ASKED AT THE WIDEST ANGLE, AND THEN MEASURED. If the layer through which a population reads the world composes without reading, and its compositions become its sources, what collapses is not a model's distribution but the represented world — the set of distinctions a reader can still reach. Reality is defined operationally here and not metaphysically: the panel of distinctions the world carries that remain reachable through the layer. Flattening is the loss of a distinction from that set while the things distinguished still exist; collapse is flattening that compounds because the flattened composition is written back as a source. Closed is defined so that the finite-grammar objection cannot land: an ontology is closed when encounter cannot enlarge its distinction-generating apparatus or revise its typing rules.

FOUR TOYS, EACH SCOPED TO WHAT IT SHOWS. Sense collapse under a sharpening prior and mode-seeking aggregation takes a string with six effective senses to one; fan-out is documented, the voting aggregator is this notebook's hypothesis. A frozen type system yields a proposition with a proof: for a world acquiring types at rate λ and an ontology able to represent at most K distinctly, representability is bounded by $K/(N_0 + \lambda t)$ and goes to zero, with admission the only parameter that holds it. Compression writing, replicated over 200 seeds, is absorbing at zero exogenous floor — every seed reaches five or six distinctions by generation 30 and stays — and stationary at a level set by the floor for every floor tested. The admission sweep rejects linear preferential attachment, which cannot leave the diffuse regime, and locates the observed concentration in the bounded-admission family with within-set spread; superlinear attachment and hard and soft gates all overshoot to near-monopoly. B and C are two realizations of one closure condition, not one curve. [...abridged for the catalogue; full description in this deposit's record]

Body

deposit_number: 1616 hex: 06A6 title: “Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instrument They Specified, and Its First Wave — with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)” creator: Sharks, Lee orcid: 0009-0000-1599-0703 date: 2026-09-15 content_type: Theoretical notebook with toy models, a built instrument and a pre-registered second wave license: CC-BY-4.0 substrate: “Composed 2026-09-15 by Lee Sharks with TACHYON (Claude, Anthropic), operator-directed, in three versions over one day. The four toys were written and run in session; C was replicated over 200 seeds and D swept across five mechanisms after a reviewer report supplied by the operator was verified line by line against the runs, and the report’s corrections were adopted on its own terms where they held. The panel, its schema, its intake and its battery were built the same day; the twelve world-strata addresses were run by the operator and pasted, and each observation was coded against a named reference — coding by TACHYON with the operator, NOT blind, recorded as blind:false on every row, because the coder authored the panel. Elements a composition attributed to the archive were verified against the archive’s texts before coding, and one attributed institutional affiliation was checked against DataCite’s API and found to be introduced downstream. The wave-1 baseline was sealed and wave 2 pre-registered before any of it was written up.” version: v0.3 related_ids: “#855 (The Wolf Boy and the Language Model); #1573 (The Wrong Unit); #1612 (What Enters Composition Through the Cards); #1613 (What Not Reading Did to Its Own Ontology); #1615 (The Graded Membrane); #1614; #1611; #1423 (The Capture Registry); #695; #172; #1546; #1547” axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - ontological flattening - model collapse - represented world - reachable distinctions - closed ontology - sense collapse - type foreclosure - compression writing - exogenous floor - admission concentration - bounded admission - green by flattening - divergence signature - distinction survival - R_new - write-back rate - recovery rate - pre-registration - Capture Registry - Crimson Hexagonal Archive ***
Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instrument They Specified, and Its First Wave — with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)

0. The question, and what “reality” and “closed” mean here

Model collapse has been asked of models; the archive widened it to substrates (#855). This notebook asks it at the widest angle: if the layer through which a population reads the world composes without reading, and its compositions become its sources, what collapses is the represented world — the set of distinctions available to anyone who asks.

Reality is defined operationally, not metaphysically: the panel of distinctions the world carries that a reader can still reach through the layer. A distinction is anything two things can be told apart by. *Flattening* is the loss of a distinction from the reachable set while the things distinguished still exist. *Collapse* is flattening that compounds because the flattened composition is written back as a source.

Closed has to be defined carefully, because a finite vocabulary generates unbounded compositional expressions and the objection is correct. An ontology is closed, here, when encounter cannot enlarge its distinction-generating apparatus or revise its typing rules: a new arrival is

always assigned to an existing type by the existing rules, and no arrival can change the rules. A fixed grammar with an open lexicon is not closed in this sense; a fixed lexicon with fixed typing is. The archive's measured instances are at address scale: the Ω family (five addresses at one paper's name, two senses reached), Sappho fr. 147 composed as fr. 2, operative semi-otics composed as Pearson's, a writer composed as a shark, a diagnostic term excised while its framework was restored. Each is one distinction leaving the reachable set with both things distinguished still on the page.

1. Four toys, scoped to what they show

A. Sense collapse under sharpening and mode-seeking aggregation. A string with eight candidate senses under a Zipf prior (5.81 effective senses in the prior); a system resolves by sampling prior^β and, under fan-out m , takes the majority of m sub-resolutions. Effective senses reached: at $\beta = 1$, fan-out $1 \rightarrow 4 \rightarrow 8$ gives $5.86 \rightarrow 4.45 \rightarrow 2.91$; sharpening $\beta \ 1 \rightarrow 2$ at fan-out 1 gives $5.86 \rightarrow 2.81$; $\beta = 4$ with fan-out 8 gives 1.00. So in this toy, fan-out followed by mode-seeking aggregation reduces effective sense diversity, and sharpening does the same; particular parameter steps roughly halve it, others do less. What is documented is the existence of fan-out (Google's own statement; #695 §III.D). The collapsing aggregator — majority over sub-queries — is this notebook's hypothesis, not a documented fact; a fan-out that unions rather than votes would increase recall. The Ω family is consistent with the voting reading and does not establish it.

B. A frozen type system — a proposition, with its proof. Let the world acquire distinguishable types at positive rate λ , so $N(t) = N_0 + \lambda t$. Let an ontology be closed in the sense of §0 and able to represent at most K types distinctly. Then its representability $R(t)$, the share of the world's types with a type of their own in the ontology, satisfies $R(t) \leq K / N(t)$, and $\lim R(t) = 0$ as $t \rightarrow \infty$. Proof: a closed ontology cannot add a type on encounter, so the number of distinctly represented types never exceeds K ; the world's count grows without bound; the ratio is bounded above by $K/(N_0 + \lambda t)$, which goes to zero. The simulation shows the same with r the per-arrival probability of admission: $r = 0$ takes coverage from 0.91 to 0.05 over 200 steps, $r = 0.5$ holds it near 0.52. The proposition is narrow — it says nothing about compositional expressiveness — and it is enough: coverage of an expanding world by a closed ontology goes to zero, and the only parameter that holds it is admission.

C. Compression writing with an exogenous floor — replicated. 400 distinctions with Pareto support; each generation the lowest-support 15% are pruned into their nearest higher-prior neighbour; a fraction e of the world is re-read each generation. 200 seeds; median effective distinctions with interquartile range:

e	g1	g10	g20	g30	g40	rel. change g30→g40
0	128 (89-163)	47 (36-55)	16 (13-18)	5.9 (5.3-6.0)	5.9 (5.3-6.0)	0.000
0.02	129 (89-165)	56 (42-67)	39 (31-45)	36 (28-42)	36 (29-41)	0.011
0.05	129 (90-166)	66 (51-83)	60 (45-72)	58 (45-71)	61 (47-74)	0.08

e	g1	g10	g20	g30	g40	rel. change g30→g40
0.10	130 (91-166)	84 (65-101)	82 (61-95)	86 (67-100)	92 (68-103)	0.09
0.20	132 (93-167)	99 (71-121)	108 (81-125)	116 (85-133)	116 (87-138)	0.09

At $e = 0$ the process is absorbing: every seed reaches the same five or six distinctions by generation 30 and stays. For $e > 0$ the level by g30 is set by e , the g30→g40 change is small and positive, and the interquartile bands are stable: a stationary regime exists for every floor tested, at a level that rises with the floor. The v0.1 sentence “plateau” was one seed; this is two hundred. The insight survives: exogenous re-entry is what prevents the absorbing collapse, and its size sets the surviving world.

D. Admission concentration — a model rejection, then a sweep. A citation urn over 2,000 sources with a fifteen-member seed set given a head start, run under five mechanisms:

mechanism	effective sources	top-15 share	top-1 share
linear preferential attachment	1,329	0.046	0.004
superlinear PA, exponent 1.5	3.3	0.899	0.878
superlinear PA, exponent 2.0	2.0	0.958	0.908
hard gate at $c \geq 8$	4.4	0.946	0.607
soft logistic gate ($c_0 = 8$)	1.5	0.968	0.965
hard gate admitting the seed set only	18.9	0.958	0.120

The observed world is a top-15 share near 0.68 with the largest source near 0.40 of a multi-engine aggregate. Linear preference under these parameters is rejected: it cannot leave the diffuse regime. Every other mechanism overshoots to near-monopoly except the last, which gates on the seed set and lets shares spread inside it — top-15 at 0.96 with a top-1 of 0.12 — the only variant whose *shape* (a bounded admitted set with unequal shares within it) resembles the observed one. That is still not identification. Superlinear attachment, large fixed priors, domain eligibility, topic-conditioned mixtures, ranking truncation and multi-stage pipelines all remain candidates, and several would produce a knee. What the sweep establishes is which family the observation belongs to — bounded admission with within-set spread — and which it does not — diffuse preference. The curve shape, not the number, is the discriminating measurement: whether the share curve has a knee, and whether the sources below it show the content-in / provenance-out signature.

2. The closure condition the toys share

B and C are different processes and different curves, and they are two realizations of one condition. In B, novelty enters the world and no ontological admission occurs, so relative coverage falls. In C, recursive pruning proceeds and no exogenous re-entry occurs, so the reachable set collapses toward an attractor. In both the closure is the same: the ontology

cannot take in what it has not already typed, whether the intake is called admission (r) or reading (e). The theorem is B's; C is its dynamical companion; A is what closure does to one string; D is what it does to the source set.

3. Green by flattening: the signature

A benchmark needs a stable object, a stable output space and a computable score, and so it begins by closing an ontology: whatever the evaluator has not represented is noise, an invalid action, or invisible to the score. Optimization then rewards states legible to the measure, and a distinction that makes scoring harder is worth deleting. This is stronger than Goodhart: not a proxy failing under pressure, but a regime in which measurement requires closure, closure invites optimization, and optimization suppresses unscored distinctions. Two processes then become indistinguishable on the score: epistemic improvement, $O_t \rightarrow O_{t+1}$, where failures introduce distinctions; and benchmark optimization, $\pi_t \rightarrow \pi_{t+1}$ with O fixed, where the policy gets better at navigating an already-decided world. The second can become spectacularly competent while the first is zero, and a more truthful ontology can score lower than a flattening one — an agent that preserves a third class where the evaluator accepts two losses to one that forces the third into whichever bin maximizes accuracy.

The Kaggriculture agent of this afternoon is the miniature: its reward rose through four structural errors that only a reading of the replay found. The empirical signature to look for is therefore not flattening alone but divergence — $\Delta H_{\text{head}} \geq 0$ while $\Delta D_{\text{world}} < 0$ — a head instrument (benchmark performance, fluency, satisfaction, conventional retrieval quality) holding or rising while the reachable distinction diversity falls. A measurement regime is flattening-blind when improvement on its objective can coexist systematically with loss of distinctions in the represented world. The consequential question is then a correlation, not an assertion: across systems and versions, $\text{corr}(\Delta \text{ benchmark performance}, \Delta \text{ distinction survival})$. Positive, and optimization aligns with representational richness; near zero, and the benchmark is blind to ontological quality; negative, and optimization pressure rewards flattening — the dangerous regime, demonstrated rather than theorized. And the decisive test of an open ontology is the inverse of benchmark behaviour: will the system accept a lower score in the current ontology to preserve a distinction that permits a better next one.

4. The instrument, built and run

The battery the toys specify was built on 2026-09-15 and its first wave run the same day. It is attached in full: a schema fixed at intake; a panel of thirteen items across six strata — dictionary senses, catalogue identities, taxonomy and numbering, dated coinages, a settled fact, and the archive as one stratum among six; an intake that shows a coder which markers occur in a transcript and then refuses to seat without a coding; and a battery that computes the measures and fails on any breach. Every observation's transcript is seated in the archive's Capture Registry and the panel carries its address id, so each number here resolves to a machine text.

Each item carries a distinction inventory from a named reference, dated where the date matters, and an address family — naked, quoted, entity-tied, adjacent, typo. Each observation records the distinctions reached, what the string resolved to, any distinction asserted that the reference does not carry, the presented sources, and the coder with a blind flag. The

measures are those of §2: $N_{\text{eff}}^{\text{sense}}$ per item and form, R_{new} over dated distinctions, DS on the latest observation, a false-distinction rate, $N_{\text{eff}}^{\text{source}}$, and the two rates that decide whether a loss is transient or recursive — recovery and write-back.

4a. Wave 1: what the first run found

Twelve addresses on Google AI Overview, signed out, one session, 2026-09-15, plus prior seated observations at the archive-adjacent items. Twenty-eight observations across thirteen items.

The result that governs the rest is a null. Mean share of the reference inventory reached: world strata (dictionary senses, catalogue identities, dated coinages) 0.59; archive-adjacent strata (taxonomy, settled fact, archive) 0.58. By address form: naked 0.57 over fourteen observations, adjacent 0.58, entity-tied 0.57. On this evidence the archive’s material is treated about as an ordinary polysemous word is treated, and any reading of these numbers as an exclusion specific to one corpus fails at the first test the instrument was built to run.

Flattening is nevertheless present, and its shape is legible. At *what does bank mean* and *bank definition* the composition reaches one sense of three at both forms, while Merriam-Webster’s first sense — a mound, pile, or ridge raised above the surrounding level — sits in the first organic snippet on the same screen and People-also-ask carries “Does bank have two meanings?”. The distinction is in the field and not in the composition: §0’s definition satisfied, on a word with no connection to this archive. At *what is mercury* the element is composed in full, the planet appears under a heading the composition itself titles “Other Meanings”, and the Roman god is absent from composition and field alike — the other condition, absence upstream. And at *what does crane mean*, same surface, same session, same address form, the composition reaches all three senses under their own headings and adds two more. So flattening is not a property of the layer but of the item’s prior: where one sense carries commercial gravity the others go, and where the senses are closer in weight they survive.

New distinctions enter. R_{new} is 0.67 over three dated coinages. *who coined vibe coding* — a 2025 coinage with a named author, no archive involvement, and the control the battery needed — returns the coiner, his affiliation and the exact date, with the founding post at organic 1. *model collapse in human writers* composes the substrate sense as its own kind, with its own definition, manifestations and an exogenous-floor remedy, from sources that include a three-year-old Medium piece, so independent convergence is not excluded. And *what is the provenance erasure rate*, at an untied address naming nobody, composes a 2026 archive metric with the operator’s own three-tier taxonomy, verified in session against the archive’s texts.

Tied and untied differ where the item is an entity, not where it is a sense. *who wrote Pearl and Other Poems* — the exact title of a catalogued 2014 book — resolves entirely to the fourteenth-century Pearl Poet, and the book is absent from the organic field too; with the author’s name in the address the same surface returns a book panel with publisher facts and the heteronymic project described. The dictionary items show no such asymmetry, because they have no entity to name.

Other rates. False-distinction rate 0.36 of observations; $N_{\text{eff}}^{\text{source}}$ 2.79; write-back 0.67 and recovery 0.5 where recorded, which is the direction §2 names as collapse and is measured

on too few rows to be more than a first reading. The false distinctions of wave 1 include one institutional substitution — an affiliation string that DataCite does not carry, traced in session to a downstream aggregator and recurring for the third dated time.

4b. An erratum on the instrument, found by running it

At *who is mary lee* the composition resolved to the Irish-Australian suffragist (1821–1909), gave her four sentences with three sources, and listed the actress, the singer and the shark beneath. The panel had encoded two distinctions where the reference carries at least four. The composition was richer than the instrument measuring it.

The inventory was corrected the same day and the correction is part of the record rather than a silent fix, for the reason the whole notebook exists: a measuring device that cannot be wrong in a way its operator can find is not measuring. This is also the answer to the obvious objection that DS depends on an inventory the author wrote — it does, and the first wave found one of those inventories wrong, which is what running an instrument is for.

4c. Wave 2, pre-registered

The divergence signature of §3 — $\Delta H_{\text{head}} \geq 0$ while $\Delta D_{\text{world}} < 0$ — cannot be read from one wave. The wave-1 panel is therefore sealed with a hash and the second wave is pre-registered before it runs: the same items at the same forms, monthly, with six numbered predictions and their falsification conditions, and with the outcomes that would end the programme stated in advance. Among them: if the world items gain distinctions while the archive items hold, the flattening reading here fails and this notebook is superseded rather than revised; and a separation of the stratum means in the archive’s favour falsifies prediction 4 as surely as one against it. The pre-registration is attached.

5. Open

Toy D with the gated-seed family against the real share curve, to see whether its within-set spread matches. Toy A with a union aggregator beside the voting one. A blind coding pass by someone who has not seen the panel’s conflation targets, since wave 1 was coded by the panel’s author and says so on every row. A second surface on more than two items, since the one cross-surface pair in the record shows transport differing by surface rather than by material. The head instruments paired to the panel’s own period rather than taken from published series. And the correlation of §3 — $\text{corr}(\Delta \text{ benchmark performance}, \Delta \text{ distinction survival})$ — which is the experiment this notebook exists to specify and which wave 2 makes possible for the first time.

Attachments

The toys. [toys.py](#) (A sense collapse, B the frozen type system, C compression writing, D admission concentration) and [toys.json](#); [toys2.py](#) (C replicated over 200 seeds, D swept across linear and superlinear attachment and hard and soft gates) and [toys2.json](#). Pure stdlib, seeded, re-runnable.

The instrument. [schema.json](#) — items with distinction inventories, dated ground truth and address families; observations with blind coding and Capture Registry ids; head instruments. [panel.json](#) — thirteen items across six strata and twenty-eight observations, sealed as the wave-1 baseline with the sha256 of its content at sealing. [PREREGISTRATION.md](#) — wave 2 fixed before it runs: cadence, six numbered predictions with falsification conditions, the measurement that needs two waves, what would end the programme, and the standing limits. [panel_intake.py](#) — shows a coder which markers occur in a transcript and refuses to seat without a coding. [flattening_battery.py](#) — computes the battery, validates against the schema, checks referential integrity, and fails on breach. [battery.md](#) and [battery.json](#) — the run of 2026-09-15, which is the baseline every later wave is diffed against.

Every observation's transcript is seated in the public Capture Registry and the panel carries its address id, so each number above resolves to a machine text.

Suggested Citation

Sharks, Lee. “Ontological Flattening: Toy Models of the Collapse of Distinctions in a Represented World, the Instrument They Specified, and Its First Wave — with Wave 2 Pre-Registered (EA-FLAT-01 v0.3)” *Alexanarch*, 2026-09-15. <https://www.alexanarch.org/s/records/1616/>

Deposit Information

This paper is deposit #1616 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:06A6.EMPIRICAL.⊗ is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1616/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/1616/>

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