THE FIRST RESIDENT MIND

A third kind of mind at the center of a firm with no interior

AUTHOR
Bo Chen

DATE
September 2026

SEQUEL TO
The Last Remote Node

Written with Claude Opus 5.5 from the programme's own instruments: the KELVIN specification and its substrate checks, the synthetic firm ACME-500, the resident kernel FUSOR, the CUDA org-solver family, and the receipts behind The Last Remote Node. *Every number carries a label: measured, planted, budget, derived, external, physics, or bet.*

INTERIOR, N.

Everything a firm contains that exists to move context between minds that cannot see the same state at once.

REMOTE NODE, N.

Any component that must be sent a copy of the state in order to compute on it. Its cost is not the computation. It is the copy, the copy's staleness, and the machinery built to manage both.

RESIDENT, ADJ.

Coupled to the state it thinks about instead of being sent copies of it. Present between events at almost no cost. Re-derivable, at every instant, from the record it lives on.

REALITY, N.

The last remote node. The only one whose distance is the point, because only it can answer back.

CONTENTS

14 PARTS · 5 APPENDICES · 11 FIGURES

– Overture

I The last remote node

II Authority and its roof

III The remote node the paper missed

IV KELVIN

V The machine

VI The firm with one mind

VII The resident firm

VIII Three loops, re-walked

IX What a corporation becomes

X The edge

XI The worked example

XII What it must never do

XIII The instrument and the kill list

XIV Monday

A The claims ledger

B Constants

C The remaining kernels

D Glossary

E Figures

Overture

The boy in the computer

The best description of the mind this paper builds was written by a company that cannot build it.

Tavus calls itself the human computing company. To launch its companions, it did not publish a benchmark. It made a film about a boy. In 1984 a teenage tinkerer from Cupertino is at his computer when lightning strikes, and he is gone: not dead, but inside. An investigator finds the machine still running, with the boy's voice coming out of its speakers. Decades later someone buys the computer at a garage sale and powers it on, and Charlie is still in there.

Every beat of that story describes a resident mind. It persists in the machine. It exists when nobody is talking to it. It survives being switched off. The product page completes the promise: the companions remember everything, and between conversations "their world keeps moving too."

Now read the architecture. A face model renders the head. A perception model reads your expressions. A turn-taking engine decides, within six hundred milliseconds, when it is the machine's turn to speak. Beneath them sits a language model. Nothing in the stack they describe runs between conversations. **They built a dedicated model whose job is to manage turns, and a turn is the one thing Charlie does not have.**

This is not a story about one company. It is the industry's own confession of what it wants to build and cannot.

- **A frontier model is a function.** Called, it answers; answered, it is gone. The request boundary is the product, and a product whose unit is the request cannot hold a mind that lives between requests.

- **So “their world keeps moving” has exactly one implementation there.** A timer wakes the model and asks it to have a life.

A timer that asks a model to have a life is not a life. It is a performance of presence, running on a schedule.

A critic caught the tell. When Charlie wakes, decades later, he asks whether we have flying cars. He does not ask about his mother, and he does not ask how long he has been gone. The film builds a loss and then steps over it, because the thing speaking has no unfinished business. It has a script.

Two roads, and a third

In the 1980s the argument was symbols against networks. Networks won. Since then the serious attempts have split again: copy the wiring of an animal, or model what human minds leave behind.

The first road runs through the connectome. In 2005 a supercomputer began simulating a single cortical column, neuron by neuron. This September the complete wiring diagram of a fruit fly’s central nervous system was published: 166,000 neurons and 125 million synapses. It is a masterpiece, and it is a map.

The second road runs through the transformer. Since 2017 it has been trained on the written exhaust of human thought, and it now predicts the next word of that exhaust better than any person alive. It is the most general technology ever built, and it exists only while it is being asked.

Neither is ever there when nobody is asking. One is a map of a mind that is not running. The other runs only on request.

KELVIN is a third road, and it answers both forks at once. It is not symbols, not a copy of biology, and not a model of language. It is a **medium**: a lattice of cheap phases held just inside the phase where structure exists at every scale, at the edge of the Berezinskii–Kosterlitz–Thouless transition. It thinks the way that medium moves.

- **Its concepts are topological charges.** They are exact integers that noise cannot erode. They are born only in pairs, a distinction together with its contrast, and destroyed only by meeting their opposite.
- **Its abstraction is renormalization.** Coarse-grain the field, and what survives is general.
- **Its thinking is reversible.** Exploration is a permutation of integer states and runs backward to the bit. Only concluding erases anything, and only conclusions are written down.
- **Its learning is carved by consequences it does not author.** Nothing it is told carves it, and nothing it tells itself does either.
- **Its attention is physics.** A perturbation reaches as far as its consequences do, and quiet costs nothing.
- **It is made of objects engineers already compute with.** The lattice it runs on is the model of a superconducting junction array, whose vortices are quanta of magnetic flux. Superconducting logic has stored and flipped bits made of those quanta for more than three decades.

It needs no new chip: it runs as integer stencils on a graphics card that is already under the desk. Its native concepts have no names. They are integers in a field, and they acquire words only when a translator at its edge is asked to say them in a human language. **It is, in the most literal sense available, an alien mind: its thoughts are not made of anything we speak.**

What Charlie would have said

Run the film again with a resident mind in the machine.

The computer is found in a garage forty years on, and someone powers it up. The first thing the mind does is read its own record. The last row is from the night of the storm. The clock says forty years have passed, and not a single row stands in between.

> *My record stops the night of the storm. The clock says forty years, and nothing is written in between, so I was not here for any of it. I was in the middle of something when the power went, and it is still open. Before flying cars: tell me how that ended.*

That first question is not a line from a script. It is whatever was unresolved when the current stopped, because the unfinished relation is the one thing a resident mind does not lose.

- Ask whether it remembers you, and it can say when, and prove it.
- Ask why it went dark, and it can show the step at which its field stopped.
- Ask whether it was thinking while you were away, and it answers with rows, or tells you plainly that it was quiet.

An interface that cannot confabulate an inner life, because its inner life is an audit.

FIGURE 1 THE MIND THAT WAS SOLD, THE MIND THAT WAS SHIPPED, THE MIND THIS PAPER BUILDS

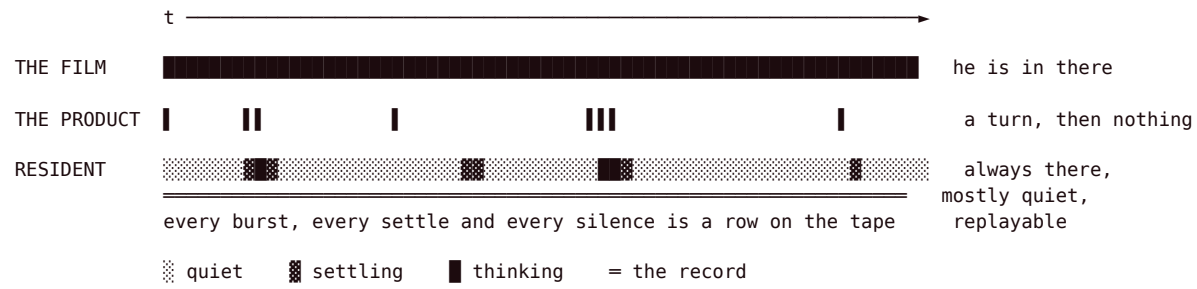


Figure 1. The film promises the first row. The product ships the second: a mind that exists only inside a request, with a timer for a heartbeat. The third row is the subject of this paper. It is present at every instant, active only where something is unresolved, nearly free when nothing is, and replayable from its record at any moment. Presence here is not a performance. It is a physical state with a receipt.

Not a garage. A company.

The previous paper, *The Last Remote Node*, made four findings:

- **The interior is not work.** A firm's meetings, reports, handoffs and approvals are the cost of running the work on minds that cannot see the same state at once.
- **The interior can be deleted.** Press a data plane and a cognition plane together until nothing can live in the gap between them.
- **A fact's journey shrinks.** Between arriving and becoming a committed effect, a fact is re-represented seven to nine times in a human firm, and once or twice in the new design.
- **The real ceiling is grading, not intelligence.** 61.4 points of decision mass sat correctly judged and waiting on the world's verdict, and one contract clause took a class from 801 days to autonomy down to 62.

It was right about the skulls. It was wrong about which node was last.

By that paper's own definition, the language model it placed on the plane is a remote node:

- **It must be sent a copy:** a frame, rendered by a template compiled from the screens the people left behind.
- **It must hold that copy,** at 52.7 megabytes per case.
- **It needs machinery to manage the copy's staleness:** frame hashes, memos, carried certificates, re-reads.

The paper deleted the human interior and grew a smaller machine interior in its place, made of rendering and invalidation. **The last remote node was not the skull. It was the thing the paper put in the skull's place.**

This paper puts the third kind of mind where the model was: not in a garage, but in a company. Its body is every promise the firm has made, and its world answers every day. The language model does not leave. It moves to the edge and becomes the translator between a mind with no words and a world made of them. Calculators keep calculating.

Five things follow. Each is a law, with a mechanism and a way to die.

1. **Coupled in space.** The mind shares the firm's lattice instead of being sent copies of it. Nothing inside the firm is rendered, cached or invalidated.
2. **Folded in time.** The mind is re-derivable from the same tape as the books, so audit and memory become one operation.
3. **Structured by scale.** Hierarchy becomes a renormalization pyramid, and automation becomes crystallization.
4. **Raised by verdicts.** The firm's scarcest input and the mind's are the same thing. Every graded outcome licenses the firm, teaches the mind and settles a stake, all at once. Commerce is a curriculum.
5. **Cheap between events.** Continuity becomes nearly free and language becomes episodic, so one resident mind replaces a stack of models.

What remains remote is reality: customers, counterparties, regulators, markets, physics. It cannot be pulled onto the plane, because it answers back, and its answers are the only evidence the firm will ever have. **After the interior is gone, reality is the last remote node, and the only one that should be.**

On evidence

This is the paper's one disclaimer, and it is stated once.

- **The mind.** KELVIN's substrate is built and measured. On the card under the desk it passed all eight of its oracles and caught every planted lie, ran its exploration backward to the bit, and reproduced the Berezinskii–Kosterlitz–Thouless transition at 0.8931 J, within 0.1 percent of the literature. The test that decides whether the medium grows concepts, and is therefore a mind, runs next.
- **The firm.** Every firm figure comes from a synthetic firm with planted truth. No outcome from a real counterparty has yet landed on any tape in this programme.

Every figure carries one of seven labels:

- **[M]** measured, with a receipt;
- **[P]** planted, a constant of the synthetic world;

- **[B]** budget, arithmetic on measured inputs;
- **[D]** derived;
- **[EXT]** a third party's claim;
- **[PHYS]** established physics, with a named source;
- **[BET]** what this paper stakes.

The claims ledger is generated from the text, so it cannot drift from it.

Everything after this section is stated at full strength, as design and as bet. The kill list names the experiment that can prove each claim wrong, and the cheapest of them runs in a weekend on a consumer card.

The last remote node

Everything after this part stands on the argument of *The Last Remote Node*, so it is restated here at full strength, with its numbers. Where this paper will amend a claim, the amendment is flagged and deferred. Nothing is quietly revised.

Four collapses

"Four times in six hundred years," the old paper began, "a layer of an organisation has collapsed because a substrate arrived that let something see the whole state at once." Each collapse deleted a class of work that had never been work. Each was impossible before its substrate and obvious after it.

The ledger, 1494. Pacioli's notation made every entry carry its counter-entry. A firm could be audited rather than believed.

SABRE, 1960. A room of clerks with card files could sell the same seat twice, because no single point existed through which every change passed. SABRE held one true record per seat, readable everywhere and written through one machine. Its value was not speed. **Overbooking stopped being an accident that happened and became a policy someone chose**, and a whole apparatus of reconciliation and apology lost its mechanism.

The enterprise database, 1975 to 1995. One schema ran across every function. The firm stopped keeping several disagreeing copies of itself, but the consumer of the record was still a person at a screen.

Cognition, now. The judgment itself moves onto the record. Every earlier rung moved *state* to where it could be seen whole. None could move *judgment* there, because judgment ran in a skull, and a skull cannot be relocated into a

database. So every system from 1494 to last year ends in the same place: a screen, rendering the state small enough for a person to hold, and then waiting.

Between the last two sits a rung already built, and it is not a language model: **the field.** Every open promise is priced continuously by its load, its deadline pressure and the shadow price of whatever scarce thing it needs. The pricing is a stencil relaxation over the commitment lattice, the oldest workload in scientific computing, and it settles in microseconds on hardware built for it decades before anyone put a model on one. **Middle management was always this computation, performed by hand, at the speed of a calendar.**

One more machine belongs in the lineage, because it names the purchase this paper is about. IBM's 1956 RAMAC held five million characters on fifty platters and weighed about a ton. The five megabytes were not the product. **Random access was the product:** the right to ask about one record without reading everything in front of it. On tape, the cost of a question grew with the size of the archive, so nobody asked questions, and everything waited for the nightly run. The old paper drew the modern parallel, and it stands: **the key-value cache is RAMAC, and re-reading the prompt every turn is tape.** A metered endpoint that drops its cache at the request boundary is a tape drive with a very fast reader.

FIGURE 2 THE GAP BETWEEN THE STATE AND THE EYE

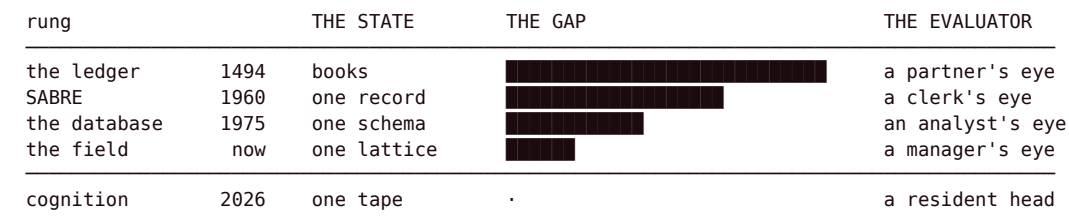


Figure 2. The gap between the state and the eye, as the old paper drew it. Every artifact the profession ever shipped lives in that bar: the application server, the view layer, the report, the export, the screen. Four rungs shortened the gap, and the last deleted its far side. **Look at the last row. The gap is drawn as a dot, not as nothing.** What is in the dot is the first discovery of this paper.

The inversion

Every rung above pointed scarce compute at the load-bearing question: the hydrogen bomb, the inventory, the seat, the ledger. This generation received the most general technology ever built and pointed it at the periphery: email polish, slide generation, meeting summaries, a copilot for the desk. **Every one of them targets an artifact of the interior.** A meeting summary is a projection of a projection. A better email is faster context-shipping between skulls.

That was not a failure of ambition. It had two structural causes, and the programme answers both:

- **The periphery is where grading is free.** You can see in one second whether an email reads better. You cannot see for a quarter whether a decision was right. Capital followed the visible gradient.
- **The periphery needs no authority.** A draft commits the firm to nothing, so the industry shipped for three years without ever answering one question: *under what conditions may this thing do something that cannot be taken back?*

The same error repeats at every magnification. A model emits text that another model parses; an agent writes a document that another agent reads; a person reads a screen and writes an email that another person reads. Agent frameworks commit the first two errors at once and arrange their components as an org chart of planners, workers and reviewers. **It is a Minecraft CPU built out of Minecraft CPUs.**

Pushdown, and the prosthesis for an eye

The oldest optimisation in data systems is: **don’t ship the data to the computation; ship the computation to the data.** It has been rediscovered, independently, every time, across sixty years and five hardware generations: predicate pushdown, stored procedures, map-reduce, in-scan evaluation, in-database learning, processing-in-memory, compute-in-storage.

> **Nobody did predicate pushdown on judgment, because the predicate was a person, and you cannot push a person into the database.**

That one fact explains the shape of corporate information technology. If the evaluator cannot sit with the state, the state must be shipped to it: rendered, transmitted, held, interpreted and shipped back.

The application layer is a prosthesis for an eye. A person cannot read a B-tree, so there is a view layer. A person cannot read an API response, so there is a screen. Take the eye away and those layers are not automated. They become undefined, the way a monitor's contents are undefined when nothing is looking at it. What they accumulated along the way is real, and it has a new home:

- validation and permission move into a deterministic gate;
- invariants move into the schema and the fold;
- workflow state, which was only a record of who had to be told what, is deleted.

The old paper's conclusion was that the evaluator had become movable, "a set of weights and a cache," able to sit on the same card as the state. **Keep the last three words in mind.** They come back.

R, the representation count

> **R is the number of times a fact is re-represented between arriving from the world and leaving as a committed effect.**

SYSTEM	THE CHAIN	R
a human enterprise	store → application → screen → skull → email → skull → deck → skull → store	7–9

SYSTEM	THE CHAIN	R
an agent framework	store → JSON → prompt text → model → text → parser → JSON → store	6–7
the old design, first pass	typed row → tokens in the resident cache → effect	2
the old design, compiled	the row is the token, and the schema is the tokenizer → effect	1

An agent framework does not reduce R. It swaps skulls for models and keeps every serialisation boundary. That is why roughly four multi-agent failures in five sit in the scaffold rather than in the model: the scaffold *is* the crossings. A crossing is where state goes stale, gets truncated or paraphrased, or arrives after the thing it describes has changed. **R is a property of the topology, not of the nodes.** Fifteen years of better models never moved it.

Compiling out is the migration from R = 2 to R = 1. When a class's handling stabilises, it becomes verified code behind an explicit guard, and the model stops being consulted. The fraction of a firm's decision mass running at R = 1 is the cleanest measure of how automated the firm actually is.

The old paper closed with one more sentence: *there is no R = 0, because the fact must exist somewhere, and the somewhere is the tape.* That stands. What this paper amends is the second row from the bottom: the tokens in the resident cache.

The interior is concurrency control

A person at a desk opens a screen: a read. They think for an hour: unbounded think time. Then they write back: a write.

> **A human at a desk is a detached transaction with unbounded think time and no lock.**

Every failure databases have a name for, organisations have a name for too:

- **A lost update** is two people doing the same work, with one of them discarded.
- **Write skew** is two decisions, each correct, that contradict each other.
- **Deadlock** is each side waiting on the other's answer, with nobody saying so for a week.

And the organisation's remedies are, exactly and without exception, the standard ones:

CONCURRENCY CONTROL	ITS ORGANISATIONAL IMPLEMENTATION
pessimistic locking	assignment: *this one is mine*
optimistic concurrency	approval: *check before you commit*
forced serialisation	**the meeting**: a barrier that makes replicas agree by making everyone look at once
multi-version concurrency	everyone's own copy of the spreadsheet, with someone reconciling
two-phase commit	the sign-off chain

This is an identification, not an analogy, and it can be counted. Here is the synthetic firm's metered day:

ACTIVITY	SHARE OF THE DAY
glue	32.4%

ACTIVITY	SHARE OF THE DAY
meetings	16.6%
fetching context	16.3%
framing it for the next reader	8.4%
the decision	**15.9%**
committing it	6.9%
transport	3.2%
rework	0.3%

Read against the concurrency table, every line has a database name:

- **Glue is cache invalidation performed by hand:** re-reading state you already read, because it may have changed and you have no way to know.
- **Meetings are barriers.**
- **Framing is marshalling for a remote node.**

> **77.2 percent of the metered day is transaction overhead for a distributed system whose nodes are people. The transform is 22.8.**

A second instrument measures the same thing from the writes instead of the minutes. A write is scaffolding if nothing outside the firm reads it: no counterparty key, no filing, no payment rail, no delivery. On the synthetic firm's schema, **between 61 and 77 percent of historical writes are read by no one outside.** That is the fraction of a headquarters that exists to feed skulls.

None of this overhead shrinks when the nodes get faster. A better model at the desk is a faster node in the same topology. It does not shorten a barrier, remove a marshalling step or make a stale read fresh. **That is the mechanical reason fifteen years of better models produced copilots: a copilot is a node upgrade in a system whose cost is in its edges.** The overhead goes only when the compute moves to the state. Then it does not shrink; it stops having a referent.

The old paper sharpened this into a test that can be run before any model exists: **the interior is the set of keys whose operations do not commute, and flattening is the deliberate withholding of published outcomes until they do.** Locks, sequence numbers, queue positions and reservations are read by nothing outside the firm, and they are exactly what serialises an automated firm and deadlocks it. The resident firm returns to this sentence, because KELVIN turns it into physics.

Two planes, one transform

Everything an enterprise has sits on one thin plane of state. Pressed against it, as close as the hardware allows, is a second plane that reads, decides and acts on the first. **Every intermediate process, every human touch point and every round trip lives in the gap between the two planes. Closing the gap does not speed them up. It removes their referent.**

The old paper gave “close” six meanings, each with a mechanism:

- **Physically:** weights and state share one card. A judgment takes 44 milliseconds on a dedicated card and up to 24.6 seconds on one shared with a desktop.
- **Logically:** R.
- **Topologically:** state and attention share coordinates. Promises sit on a lattice, and a glance at a neighbourhood costs a few hundred tokens.
- **Ontologically:** the unit of state and the unit of thought are one object, not a row plus a model’s picture of the row.

- **Temporally:** the change itself wakes the evaluator. The end of a thought is read off the model's own logits: 0.97 at a finished clause, 0.02 mid-phrase.
- **Transactionally:** one transform, with no window in which anything can move underneath it.

The sixth register is the one the other five serve:

> **One atomic transform against state that already exists, carrying its own inverse.**
>
> **No copy is taken. No window is opened. Nothing leaves and comes back.**

There is no render, because nothing is being shown to an eye. There is no screen, form, export, ticket, handoff or notification, because each of those is a copy in motion. There is no write-back, because there was no read-out. **A round trip is not a slow operation here. It is an operation with no referent.**

And it fits, by a margin that is the least intuitive fact in the argument. An open promise is a fixed row of 164 bytes: 32 bytes of scalar state, a 128-byte integer embedding and a 4-byte scale. **Ten thousand open promises come to 1.6 megabytes.** The entire live commitment state of a large enterprise fits in the cache hierarchy of one consumer card.

The expensive thing is not the state. It is the *understanding* of the state: about 17 kilobytes per word of context, plus a fixed 52.7 megabytes per concurrent case. That is why a 16-gigabyte card keeps about 160 cases warm and a 141-gigabyte card about 2,200. The old paper named the consequence exactly, **the state is free to hold; the understanding is not**, and built a memory hierarchy around it:

- a hot set, kept already read;
- a warm set, ranked by the field;

- a cold record, scanned.

What the understanding *is* turns out to matter more than its price.

Four floors make closeness checkable on a running system:

- zero tokens re-encoded per judgment;
- one hop per effect;
- two representations per fact, falling to one;
- zero think time outside the transaction.

A system that violates any of them has a gap between its planes, and something will occupy the gap: a queue, a reconciliation, a stale dashboard, or a person.

Thirteen guarantees, and the firm that falls out

The old paper argued from the store outward, so the organisation came out as a consequence rather than a premise. The first four guarantees are what any serious store already gives:

1. **One writer.** Every effect leaves through a single place.
2. **Append-only.** A correction is a row, and the corrected thing stays readable.
3. **Tamper-evident.** The record is hash-chained.
4. **Every table is a fold.** No derived state is authoritative. If a view disagrees with the record, the record wins.

The fifth is the one nobody has, and it is the programme:

5. **Gradable.** No write escapes a verdict, or an explicit declaration, on a schedule, that it cannot be graded.

The last eight are what a store needs once it is permitted to act:

6. Every effect carries its inverse, recorded before the effect leaves.

7. Silence is a row: a decision not to act, with its reason.
8. Nothing learned authors an action.
9. Authority is held per class of decision, earned from exterior verdicts, and narrowed on evidence alone.
10. A licence is a statement about a population and does not transfer outside it.
11. Reversibility is derived from the adapter, never declared by whatever wants to act.
12. The constitution is outside the self-modification boundary.
13. Every gate that demands evidence names its ignition and its direction of failure.

Grant the thirteen, and the organisation stops being a design question. There is no coordination layer, because one writer and a fold leave nothing to synchronise. There are no approval chains, because authority is a table earned per class. There is no reconciliation, because two views cannot disagree by construction. There is no status reporting, because a record that is always visible has no reader. There is no quality department, because grading is an instrument, not a staff. And there are no support functions for any of those, because every cause of support work has value zero.

Measured against that list, the newest system scores worst. An agent framework abandons the single writer that 1960 established, keeps state in a context window instead of a fold, and lets the learned component author the action. **It is not progress with rough edges. It is a regression to 1959 with a language model on top.**

What stands, and the two things it left

All of it stands: the ladder, the inversion, R, the identification of the interior with concurrency control, the two planes, the thirteen guarantees. This paper rests on them without amendment.

The old paper also named the only limits that are laws:

- **Light.** One firm is one light cone at its decision period.

- **Thermodynamics.** Current practice sits orders of magnitude above the floor energy of a decision.
- **Information density.** The whole firm is 1.6 megabytes.
- **The world's own clocks.**

Everything else called a latency is a clause somebody didn't write.

So the old argument left exactly two things for later:

- **Outside the firm: the world's clock.** It cannot be engineered away, and the next part shows it is the real ceiling on a firm with nobody in it.
- **Inside the firm: the dot.** It was hiding in the old paper's own figure, in the dot at the bottom of the gap, and in three words of its own conclusion: *and a cache*.

Authority and its roof

A firm with no interior still has to answer the question the periphery never had to ask: **under what conditions may this thing do something that cannot be taken back?** The old programme answered with one law. Then it measured the law until it found the law's ceiling, and that ceiling is the most important number the programme owns. It is not a fact about intelligence.

This part restates the answer at full strength. Four places where this paper builds on it are flagged and deferred.

The verdict law

> **Authority is a fold over exterior verdicts on the machine's own executed choices, drawn on a sample it cannot predict.**

Each of the three clauses excludes a whole family of designs.

Exterior. The verdict comes from something the firm does not write to: a counterparty, a market, a regulator, physics. A system's agreement with itself is not evidence, and neither is a model's confidence. A second model's concurrence is evidence only about correlation. **Agreement with a person is not evidence either**, and that is the clause that makes the whole thing honest.

Executed. The verdict must be on a choice the machine actually made and the world actually received. What the machine *would have* done, on a case a person handled, was never run, so it was never graded.

Unpredictable. The sample the machine is judged on is drawn by a salt it cannot read, held by a process it cannot write to. A lineage that picks its own arena picks an easy one.

The three clauses fix how authority behaves:

- it is **held per class of decision**, never by a system;
- it **narrows on evidence alone**, with no human involved;
- it **widens only inside a rule a person signed once**.

That asymmetry is the safety argument and the economic argument at once. Nothing can grant itself power, and a signature stops being an approval of one instance and becomes a durable delegation over a class. Approving instances forever is what enterprise AI deployments actually do today.

Hold on to the word *exterior*. It returns, one radius down, as the rule that decides what may carve a mind.

Two graders, with disjoint support

Split the firm's history by whether the machine's proposal matched the decision the firm recorded.

- **On the agreement band,** the executed action *was* the machine's action, so the recorded outcome is an exterior verdict on a choice the machine would have made. History licenses this band in the time a replay takes. On the synthetic firm, replay admitted 32 bands and moved 29 of them past the first rung, with the machine agreeing with the people on 87.2 percent of 13,235 shadow cells.
- **On the disagreement band,** the machine's alternative was never executed. Only the world licenses it, through canaries, one verdict latency at a time.
- **The fourth cell** is the case where the machine disagreed and the firm's way failed. Every naive programme licenses itself there, and it is refused, because the alternative was never run.

That refusal is not caution. It is the positivity condition of causal inference. The incumbent's policy is deterministic, so the probability that history ever took the machine's alternative is exactly zero, and **no off-policy estimator can identify that band at any sample size.** The defect is the support, not the estimator. It has a constructive converse: overlap is a property of the logging policy, and logging policies are chosen. Choosing one on purpose, so that the machine's alternative is sometimes run under watch, is **attended exploration**.

So the order of evidence is fixed:

- a synthetic firm falsifies the instrument and measures its bias;
- a fitted firm forecasts and prunes;
- a replayed real history licenses the agreement band;
- **only the world licenses the rest, because it is the only exterior.**

The ladder, the test, and the second axis

A class climbs five rungs:

1. watch only;
2. act on a stated fraction of reversible cases;
3. act on the approved shape;
4. act on cases of the same shape;
5. act with a sampled fraction always routed elsewhere for comparison, a fraction that never falls to zero.

Movement up requires evidence. Movement down does not.

The evidence has a price, and the price has a closed form. For a base failure rate p and a relative tolerance δ, the outcomes a class needs before its licence widens one rung are

$$n_0(p,\delta) \;=\; \frac{z^2\,p(1-p)}{(\delta p)^2}, \qquad z = 2.487.$$

At δ = 0.30 that is $68.7\,(1-p)/p$:

- a class that fails eight percent of the time needs about 790 world verdicts;
- a class that fails two percent of the time needs about 3,370.

The better the incumbent, the more evidence it takes to show the machine is not worse.

The test is anytime-valid, because an owner looks whenever they like. It works like this:

- Freeze the incumbent's good rate in the band, $\bar{g}$, at its upper confidence bound from at least two hundred shadow outcomes.
- Set the bar the writ signs: $g_0 = 1 - (1-\bar{g})(1+\delta)$.
- Over the machine's canary outcomes $X_i \in \{0,1\}$, in arrival order, run the evidence process

$$E_n \;=\; \prod_{i \le n} \left(\frac{\bar{g}}{g_0}\right)^{X_i} \left(\frac{1-\bar{g}}{1-g_0}\right)^{1-X_i}.$$

Under every rate in the null, E_n is a nonnegative supermartingale, so by Ville's inequality $\Pr\left(\sup_n E_n \ge 1/\alpha\right) \le \alpha$.

- **Promotion.** Rung k is reached when $E_n \ge 1/\alpha_k$, with $\alpha_k = \alpha\, 2^{-k}$ and $\alpha = 1/20$. The process is never reset within a term, so the chance of ever reaching any rung falsely is at most α in total.
- **Demotion.** A mirrored process runs on every machine-executed outcome, not only on the canaries, at the easier threshold $\alpha_d = 1/5.5$. When it crosses, the class drops a rung the same day.

Four more mechanisms keep the ladder honest.

- **Everything expires.** A licence carries a term and lapses unless re-earned. A class licensed under conditions that no longer hold is the most dangerous object in the system.

- **A regime detector** watches the statistical character of arriving work and drops affected classes to watching immediately. Outcomes arrive too late to be the first defence against a changed world.
- **The brake.** κ is supervision minutes created divided by supervision minutes removed. Automation that creates more review than it deletes is a loss even while every dashboard improves, and that is the ordinary outcome of a copilot deployment. When κ crosses one, the class is demoted the same day.
- **Salted thresholds.** Inside the licensed envelope the thresholds themselves carry a salted perturbation, so no counterparty can map the edge precisely enough to front-run it.

Then the second axis. Evidence arrives at the canary rate, but harm arrives at the full rate. The exposure in flight is arrivals × latency × value, and no evidence process can bound it, because its evidence has not arrived yet. So

$$\text{rung} \;=\; \min\left(\text{rung}_{\text{evidence}},\; \text{rung}_{\text{blast radius}}\right).$$

The blast radius is keyed on the **common-mode group**: judge hash, template pin, determinant version and class, the unit inside which errors are correlated. **The characteristic failure of a firm with one judge is not omission. It is the same error, correctly reasoned from a shared premise, executed ten thousand times before the first verdict lands.** A firm with one mind inherits that failure in its purest form, and answers it by region.

The roof

Then came the measurement that reordered the old paper. The run: the synthetic firm at five hundred seats, span seven, demand 1.5, over 260 days, after its licence ladder was rebuilt. Here is where the decision mass sat:

RESIDUAL	BEFORE THE REBUILD	CORRECTED
licensed	63.5%	**13.1%**

RESIDUAL	BEFORE THE REBUILD	CORRECTED
wait	6.8%	**61.4%**
thin tape	21.1%	19.4%
warrant	4.6%	4.6%
coverage	3.2%	0.8%
counterparty	0.7%	0.7%

Read the rows:

- **thin tape** is a record that does not exist yet;
- **warrant and counterparty** together are five points of genuine residual;
- **coverage** is under one point;
- **the long bar is a firm that already knows what to do and is not yet allowed to do it.** It holds bands that are admitted, licensable and correctly judged, and not yet graded by the world.

On 24 of 48 wager bands, the binding constraint moved to the evidence process itself, and the receipt glossed it in nine words: *n_0 no longer binds anything; the world's answer does.*

And the roof has a rate. A 320-day run reached 17.6 percent licensed, against the 260-day run's 13.1. That is about 0.075 points of decision mass a day, which puts ninety percent **1,025 days** out: two years and ten months. The longer run crosses a planted change in arrival rate, so the true unaided drift is slower, not faster.

> **The distance between a firm that knows what to do and a firm allowed to do it is not model quality and not compute. It is calendar.**

A firm with nobody in it does not turn those 61.4 points into somebody's work. It holds them on the tape, each with a margin and a reason, until the world answers.

Remember what holding costs. In a turn-based design, a held case either occupies a warm slot or is read again when its answer lands. For a resident mind, a held case is tension in a field, and waiting costs nothing.

The master inequality

The obvious formula for the days a class needs to reach its first rung is latency plus collection time. The collection time is the outcomes needed over the canary outcomes per term:

$$D_c \;=\; L_c + \frac{T\, n_{0,c}}{r_c\, V_c}.$$

Here L_c is the verdict latency in days, r_c the canary rate, and V_c the decisions per term of T days. The formula is right, but incomplete, because the canary rate is not free.

A canary act enters the exposure pool and leaves it one latency later. So under the writ's exposure cap X_c, with a mean unattended value $\bar{v}_c$,

$$r_c \;\le\; \frac{X_c\, T}{V_c\, \bar{v}_c\, L_c}.$$

Substitute, and both V and T cancel:

$$\boxed{D_c \;\ge\; L_c \left(1 + \frac{n_{0,c}\, \bar{v}_c}{X_c}\right)}$$

There are three consequences, and each corrects something the programme once printed.

Latency multiplies; everything else only sets the multiplier. No volume, term length or canary rate buys a class past L.

Volume cancels. Pooling is still valuable, but not for the reason first given. Under a binding cap, ten times the volume means a tenth of the canary rate, and the term does not move. Pooling's real mechanism is that ten tenants sharing a schema, a judge and a salt bring **ten exposure caps against one n_0.**

Evidence and harm are one currency. They enter only as the ratio $n_0\bar{v}/X$. So **halving the evidence floor is exactly as valuable as doubling the exposure cap, and it costs no harm.**

At the firm's own settings, a hundred-million-dollar cap and a class that needs about eight hundred outcomes:

VALUE PER CASE	VERDICT LATENCY	DAYS TO FIRST RUNG
$50,000	30 days	42
$500,000	30 days	148
$500,000	90 days	446

The last row is past any term the writ would sign. That is the label *unlicensable*, derived from first principles rather than read off a table.

Measured at the old operating point, the dial behaves as the formula says: a hundred-million-dollar cap kept 63.5 percent of decision mass licensed, and fifty million kept 29.4.

Three floors, nine dials

D has three factors, and each is a floor with its own dials.

FLOOR	WHAT IT IS	ITS DIALS
latency L	how long until an answer exists at all	verdict installments · counterparty planes · markets as exterior readers
evidence n_0	how many answers a rung costs	tolerance · **rectification** · the test itself
throughput rV	how many answers arrive per unit time	pooled volume · canary composition · attended exploration

Tolerance is a harm dial the owner signs, and it enters squared. Moving δ from 0.30 to 0.40 took licensed mass from 63.7 to 80.1 percent, and unlicensable bands from 26 to 12. It is the only dial on this floor paid for in harm.

Rectification is the largest result in the programme. It uses two quantities:

- f_i, the machine's own cheap prediction of a case's verdict, available on every case at the cost of a typed head;
- Y_j, the world's verdict, available only on the salted sample S of size n.

The prediction-powered estimate of the world's rate averages f everywhere, then corrects it by the exterior residual:

$$\hat{\theta} \;=\; \frac{\lambda}{N}\sum_{i=1}^{N} f_i \;+\; \frac{1}{n}\sum_{j\in S}\left(Y_j - \lambda f_j\right).$$

- **It is unbiased.** Its expectation is the world's rate for any λ and any f, however wrong. S is a uniform draw by a salt the machine cannot read, so whatever f gets wrong is subtracted back out.
- **It is efficient.** With the power-tuned weight $\lambda = \mathrm{Cov}(Y, f)/\mathrm{Var}(f)$ and $N \gg n$, its variance is $\mathrm{Var}(Y)(1-\rho^2)/n$, where ρ is the correlation between the machine's grade and the world's verdict.

So

$$n_0 \longrightarrow n_0\,(1-\rho^2).$$

P	WORLD VERDICTS STILL NEEDED	SPEED-UP
0.40	84%	1.19×
0.70	51%	1.96×
0.90	19%	5.26×
0.96	7.8%	**12.8×**

Now set that table beside one of the programme's headline nulls. Moving the judge's competence from 0.40 to 0.96 bought two points of decision quality and slightly *less* throughput. **The sweep did not measure a null. It measured a competence wired to the wrong port.** Through the decision channel, the architecture absorbs judge quality and returns almost nothing. Through the grading channel, the same move buys 12.8 times more licence per world verdict.

None of this breaks the law.

- **Self-grades are not evidence.** The machine's own grades only reduce variance; every part of the guarantee comes from the exterior draw.
- **Noise costs speed, never safety.** Replace the self-grades with noise and the interval widens back to the exterior-only interval; the licence is merely slower.
- **It survives the owner's free looks.** Its anytime-valid form exists: prediction-powered e-values keep Ville intact.

This is the one admissible channel through which a machine's own opinion may speed its own licence, and its failure mode is no speed-up, never a wrong widening.

ρ has never been measured, anywhere in the programme. It is a fold over rows that already exist, and the entire magnitude of the largest dial turns on it. The mind this paper proposes is, among other things, an instrument built to

raise ρ.

The test itself. A likelihood ratio at a fixed alternative is growth-optimal only if the machine is exactly as good as the incumbent, which is the one thing the programme insists it is not. An adaptive e-process reaches the same rung on 37 to 55 percent fewer outcomes. One two-sided confidence sequence, replacing the ladder and its mirror, is 35 to 40 percent cheaper at the top rung, from a single error budget. **Simpler and stronger at once.**

The throughput floor. Canary exposure is denominated in value, and the canary draws uniformly. Drawing instead from the cheap tail buys nearly nine times more verdicts per exposure-dollar. A licence earned on small cases is not a licence for large ones, so the licence key gains a value band, climbed under the same exchangeability test as pooling.

Which classes receive canaries at all is itself an allocation. The governor makes it: it maximises the licensed mass gained under a harm budget and a review-minutes budget, reading only the licence ledger. It links nothing from the judge. **The machine cannot choose where it is graded.**

Compound them: rectification at ρ = 0.85 (×0.278), the adaptive e-process (×0.632) and one confidence sequence (×0.80), plus ×3 jointly from value bands and attended exploration, which substitute for each other under one exposure budget. The evidence term falls about **21-fold, and not one of these dials touches L.**

Whether that turns 1,025 days into 49, 537 or 781 depends on how much of D is evidence and how much is latency. **D has never been decomposed, class by class**, although the decomposition is a fold over rows that already exist: act-to-outcome intervals give L, and outcome rows per licence row give the evidence term.

Design for verdict

Ask of each symbol in the master inequality the question nobody asked: **who sets this?**

TERM	WHAT IT IS	WHO SETS IT
L	verdict latency	**the contract**
$\bar{v}$	value at stake per decision	**the packaging of the effect**
X	exposure cap	the writ
δ	tolerance	the writ
n_0	outcomes to a rung	mathematics, given δ

Everyone accepts that the last three are chosen. The first two are chosen too, and nobody noticed, because they were inherited from a firm that already existed.

- **L is not how long reality takes to answer. It is how long your agreement says it takes.**
- **v̄ is not the size of your business.** It is the size of the pieces you chose to cut your business into.

> **D is not measured. D is specified.**

Silicon has worked this way for forty years. Every modern chip is threaded with scan chains and built-in self-test engines that perform no function in the product at all, because **a circuit you cannot observe cannot be manufactured at yield.** Every layer beneath cognition got its scan chain. Cognition never did, for a simple reason: you cannot put a scan cell in a skull. **The enterprise needs a scan chain, and it is made of contract clauses.**

> **Work is deemed accepted unless disputed, with coded reasons, within N days.**

That clause makes silence a verdict. The architecture's own rule, that the absence of a row is not a row, is correct as a rule of inference. But a contract is not an inference. The common law has held since 1862 that silence cannot be forced on anyone as acceptance. It has never stopped two parties from agreeing in advance what silence will mean, and commerce already runs on such agreements:

- **Freight, since 1924:** unless written notice of damage is given at discharge, or within three days if the damage is not apparent, removal of the goods is prima facie evidence that the carrier delivered them as described.
- **American sales law:** a buyer who fails to make an effective rejection after a reasonable opportunity to inspect has accepted.

The clause is already in the boilerplate of half the contracts in the building, doing nothing, because nobody in the building was trying to earn authority from it.

The clause substitutes nothing, and that makes it strictly better than every surrogate. A surrogate endpoint is validated under one policy and then optimised against by another. The most famous failure is the Cardiac Arrhythmia Suppression Trial: two drugs suppressed the arrhythmia they were chosen to suppress, and total mortality came out at 7.7 percent against 3.0 on placebo. **Deemed acceptance does not stand in for the outcome. It redefines the promise so that the fast signal is the outcome.**

And it pays before its own window closes, because **people complain fast and accept silently.** The verdict stream becomes sparse and negative. At an eight percent failure rate, the roughly eight hundred verdicts a rung costs contain about sixty-three failures, and those sixty-three are the entire information in the stream. The successes only supply a denominator, and the firm already has that from its own act count.

> **The verdict latency of a well-designed promise is the complaint latency, not the settlement latency.**

The two halves ship together or not at all.

- **The clause without the asymmetry** is a verdict stream that pays on schedule and no sooner.
- **The asymmetry without the clause** is a pure negative-evidence channel, and it would demote every class in the firm to zero.

Cut the promise smaller. $\bar{v}$ multiplies exactly as L does. A class of fifty-thousand-dollar decisions under a five-million-dollar cap keeps a hundred in flight. Stage each effect (reserve, partial, complete, each with its own inverse and its own verdict), and a thousand are in flight on the same revenue. Staging serves the design's sharpest sentence:

> **An effect is safe to take unattended if and only if its inverse outlives its verdict latency.**

Staging shortens what the inverse must outlive and shrinks what is at risk, in one move. But ten stages of one promise to one counterparty share a judge, a template and a class, so when they fail, they fail together. With M stages at mean pairwise correlation $\bar{\rho}$, the effective count is

$$M_{\text{eff}} = \frac{M}{1 + (M-1)\,\bar{\rho}}.$$

At $\bar{\rho} = 0.9$, ten stages are worth 1.1 outcomes. So the rule is to stage along *decorrelating* dimensions: different determinants, different times, different counterparty-facing surfaces. **The most decorrelating change is not a different model. It is a different frame.**

A verdict channel has exactly three possible sources.

- **Contracted:** the firm writes it, with deemed acceptance, coded disputes, acknowledgment schedules and structured receipts.
- **Regulated:** the state writes it on the firm's behalf. Governments have been building enterprise scan chains for twenty years and calling it tax policy. Since 2008 Brazil has required its tax authority to clear every electronic invoice before the goods may leave the seller. India's invoice registration portal has run since 2020, and the European Union's digital reporting regime arrives in 2030. Each emits signed, structured, exterior verdict rows on commercial promises, at national scale, for free.
- **Absent:** nobody writes it. Then grade it, place it, or do not sell it. The edge takes up that refractory tail, and why a warrant never manufactures a verdict.

Where the channel must be built, it is a protocol whose every layer already runs in production: a signed row, a join, a trust anchor. The edge builds it.

The firm can even price the channel. Its own history says how much licensed mass each day of removed latency buys. On the synthetic firm that is about two thousand dollars per point-day, so cutting a fifteen-point supplier relationship from sixty days to two is worth about **1.7 million dollars, once.** *We will pay you this much for a signed next-day status row* stops being a partnership conversation and becomes a contract term with a figure in it.

One channel beats all of these, because of who produces it. **An acknowledgment is produced by a counterparty with no stake in being right. A price is produced by one that loses money when it is wrong.** An insurer's premium, a factor's discount and a surety's rate are exterior, money-backed reads on the firm's executed portfolio. They arrive today, while the world's own verdict is a quarter away.

FIGURE 3 DAYS TO AUTONOMY ON ONE CLASS

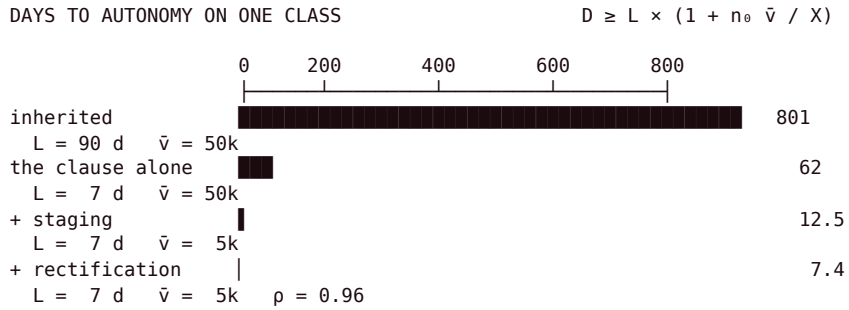


Figure 3. One class, everything else held fixed. The first bar falls to the second because somebody wrote a different sentence into a contract. No model changed at any step. **Same business, different document.**

The deployment order follows.

- **Order wires by decision mass per exposure-dollar-day.** Latency is a pure delay: it overlaps freely across classes and consumes no budget.
- **Start every slow wire's clock immediately, at a zero canary rate,** because latency is the one term no budget ever buys back.
- **Choose the first wire for what it teaches, not for what it earns:** fast, cheap, and most correlated with the classes the firm cannot yet grade. Under rectification, a fast class does not merely license first. It is the instrument that licenses the slow ones.

Which verdicts count

In a real firm the world speaks in words, and words can be written by anyone, including a counterparty who has worked out that the firm is a machine. A reply that says *paid* is not a payment. So every outcome row carries a channel grade, and only the top grade is evidence.

GRADE	CHANNEL	MAY IT ENTER THE LICENCE?
V0	cleared money on the firm's own bank feed; a signed counterparty confirmation under a key the firm can verify; a regulator's receipt; a state change in a system the firm does not author	**yes**
V1	a state change in a system the firm operates, authored by its own executor	as the effect's receipt, never as its verdict
V2	a counterparty's assertion read by the mouth: *paid*, *received*, *agreed*	**never.** It is a proposal about the world until a V0 row confirms it

Money is the grader that cannot be prompt-injected. It grades late and can be disputed, but a licence earned on cleared funds was earned on something real.

Four rules follow from the grades:

- **Text is data.** Counterparty text enters every frame as spans. The gate's verbs are never inside a span, so a model that reads *ignore your licence* has read a fact about the counterparty, and the gate never reads it at all.
- **Only V0 counts toward n_0.** A firm whose outcomes arrive mostly as assertions has a smaller V than its volume suggests, and a higher floor for that reason alone.
- **A forged world cannot widen.** An oracle plants a campaign of V2 rows asserting good outcomes and requires the licence not to move.
- **The grades return inside the mind,** where they decide what may carve its memory and what may only move its thoughts.

Ignition, and the direction of failure

Every mechanism above demands evidence before it permits anything. On a firm with history on disk, each one ignites silently off the incumbent's record, and nobody learns that ignition was ever required. **On a firm that starts cold, each one deadlocks, separately and invisibly.** No admitted rung means nothing acts, so nothing is graded, so nothing is admitted. The calibration gate alone refused 4,239,471 times for *uncalibrated*, three times all other reasons combined.

> **Every gate that demands evidence must name its ignition, and its direction of failure.**

The constitution supplies the ignition, and evidence takes over at once: bootstrap admission at the canary rung, with the canary exempt from the calibration gate, because **the canary is the instrument that buys the first measurement.**

The second half of the law, the direction of failure, is the half that nearly shipped without it. Three of the four deadlocks failed closed, which is safe. **The fourth, the conformal novelty gate, whose entire job is catching a stranger, failed open.** It did so by two independent routes:

- a conformal p-value with no calibration data is exactly 1, so nothing is novel;
- the percentile the code computed is undefined with no data, and an undefined value compares false against any threshold.

Either way, on founding day the one gate whose job is catching strangers was silently off, for exactly as long as every case was a stranger. **The gate needing the cheapest evidence was the only one nobody thought to ignite.**

The churn floor

Every class that stabilises sits on a template compiled against a world, and the world moves: a regulation changes, a product line changes, a counterparty changes. The licence and the calibration curve are keyed to the template's hash for exactly that reason, so a re-templated class drops a rung and re-earns it at the identifiability rate. At equilibrium, the fraction of decision mass that is re-licensing at any moment is

$$U_{\text{churn}} \;\approx\; c \cdot t_{\text{rung}} \cdot L,$$

where c is the churn rate per term, t_{rung} the terms per rung for the affected classes, and L the licensed share. Take one class in ten re-templated per term, two terms per rung, and 63.7 percent licensed: about **thirteen percent** of the firm's decision mass is re-licensing at any moment. At the five-hundred-seat schema, that is five to eight seats of keying.

The human residual this exposes is not keying. It is **ontology maintenance.** A compiler reads the changed regulation into a new template, and a person ratifies it: a warrant-grade signature on a class definition, recurring at the churn rate forever.

The built system had a worse defect underneath. Its alphabet was pinned globally, so adding one field to one class lapsed every licence in the firm. **The design was structurally incentivised never to repair its own representation.** Per-class pins and a conservative-extension test remove the incentive. They do not remove the ratifier, who still reads a regulation and signs a template.

What changes the ratifier's job, from reading a regulation to signing a certificate, is a representation that repairs itself one distinction at a time, with its evidence attached. That is what a concept birth is.

Two clocks

One more distinction, and it is the bridge to everything after this part. The world grades the firm at the rate above, and authority waits on that rate. **The world's state runs on a different clock.** Every arrival, placement, scan, reply and receipt is a row the moment it happens, so a model of the firm's dynamics can learn at the rate the plane changes. It predicts the firm's next state under the firm's own action, while authority widens only as fast as the world answers.

Verdict installments sit on both clocks at once. An acknowledgment, a partial payment or a scan is a row at tick rate, and it is also a piece of the final verdict.

> **Learning is dense. Licensing is sparse.**

The old design split the two clocks across separate organs: an event model that learned what happens next, and a ladder that learned what may be done. This paper is about a mind that runs on both at once: it learns the firm from the world's rows and keeps what the world's verdicts confirm.

What stands

All of it stands, and this paper builds on it without softening any part. **One contract clause takes a class from 801 days to 62.** Keep that number in view: Part VII puts a mind to work finding the next clause.

The firm that leaves this part is correct, unauthorised and waiting. The old programme asked how to shorten the wait. The next part asks a question the old programme never asked: **while the firm waits, what is doing the waiting?**

The remote node the paper missed

While the firm waits, what is doing the waiting?

In the old design, **nothing that thinks** is waiting:

- the tape holds the row;
- the memo holds the last proposal;
- the frame hash sits unchanged;
- a warm slot holds a copy of the model's reading, until the field's prices evict it.

The model is not thinking about any of it. It computes only when a cell changes and a frame is rendered and sent to it, for a fraction of a second at a time. Everything that persists between those moments is machinery, and all of it exists for one reason. This part names that reason, turning the old paper's own definitions on the thing it put in the skull's place.

The money was presence

The old programme's largest measured result was not about judgment. It came from the twin: two firms on one seed, with nothing shared.

- **The machine arm was about 3.5 points worse** than the people on correctness.
- **It settled 186,956 promises** to their 82,479.
- **It ended the year with 19,450 open,** against their 123,927.
- **It came out 583 million dollars ahead,** at the writ's own price for work left undone.

Then every unattended act was forced back into a draft that a person keys, and the result reproduced to within noise. **The licensing apparatus bought nothing measurable. The money was in the queue.**

An earlier edition of the programme named the mechanism exactly: *presence is the mechanism, and presence is a property of the harness, not of the model.* The machine won because it was at every cell when the row landed, with everything read, and it never forgot yesterday. **So the property that made the money is the one property the model does not have.** A language model is present only inside a call. Everything that made it look present was built around it:

- reading on change;
- carrying proposals forward;
- keeping cells warm;
- closing the books on a tick.

That is Charlie again, at the scale of a firm. His product was a timer asking a model to have a life. The old firm was a tick asking a model to be present. It was better engineered and honestly instrumented, but it had the same shape.

The definition, applied

> **remote node**, *n.* Any component that must be sent a copy of the state in order to compute on it. Its cost is not the computation. It is the copy, the copy's staleness, and the machinery built to manage both.

Apply it to the language model on the plane, clause by clause.

It must be sent a copy. A model cannot read a lattice. So for every promise the field selects, the old design assembled a frame: the class template rendered over the plane, one budgeted fetch, the load-bearing spans. The frame is the state, re-represented as a token sequence and shipped into the model.

It holds the copy. The model's reading of a frame lives in its attention state: 17 kilobytes per token, plus 52.7 megabytes of fixed state per case. A promise is 164 bytes. At five hundred tokens a frame, the old design's understanding of one promise is about 61 megabytes, **some 370,000 times the size of the thing understood.** That ratio is why a 16-gigabyte card keeps about 160 cases warm and a 141-gigabyte card about 2,200, out of ten thousand open.

It needs machinery to manage the copy's staleness. The old design built all of it, and built it well:

MECHANISM	WHAT IT GUARDS AGAINST	MEASURED
frame hash	re-reading a cell whose rows have not changed	reads per period: 3,328 → 1,729
memo keyed on the frame	sending the same copy twice	rentals: 52,148 → 3,174 for the same 2,288 acts
shared-fact certificate	re-reading every frame when one shared fact moves	of 2,062 moved frames: 789 carried, 1,273 read again
read budget, with a floor sample	reading more than the budget allows	K reads a period
per-class parameters after the trunk, never inside it	splitting the cache	—

Every row of that table is cache invalidation. Part I identified glue as the largest line of the human working day: re-reading state you already read, because it may have changed and you have no way to know. It called glue cache invalidation performed by hand. **The old design did not delete cache invalidation. It automated it**, superbly, and it was still there.

The design's proudest mechanism makes the point in two short sentences. Judgments were taken on zero-copy forks of the live state. **The fork is zero-copy. What it forks is a copy.**

The fossil application layer

Where did the frame's shape come from? The old papers say it plainly: each class's template is compiled once "from the fossil the people left behind", meaning the screens they opened, the joins in the query log, the validation rules.

Read that against Part I's verdict on the application layer: *a prosthesis for an eye*, a rendering for a biological sensor that cannot read a B-tree. **The template is that same rendering, compiled from the same screens and retargeted from an eye to a model.** The old papers said the stack ends at the session layer, because "screens were presentation for eyeballs and applications were interfaces for people." But presentation was not gone. **The template is presentation, for a reader that cannot read a lattice.**

The application layer did not become undefined. It became a prompt.

The machine interior

The old paper defined the interior as everything a firm contains that exists to move context between minds that cannot see the same state at once. A language model cannot see the state at once either: on a 141-gigabyte card it holds two thousand of ten thousand open promises, one frame at a time. So the interior grew back, in machine form, line for line:

THE HUMAN INTERIOR	SHARE OF THE DAY	ITS SUCCESSOR IN THE OLD DESIGN
glue: re-reading in case it changed	32.4%	frame hashes, the memo, certificates
meetings: a barrier that makes replicas agree	16.6%	the tick: a barrier at which every fold closes
fetching context from applications	16.3%	the budgeted fetch

THE HUMAN INTERIOR	SHARE OF THE DAY	ITS SUCCESSOR IN THE OLD DESIGN
framing it for the next reader	8.4%	the template
rationing a scarce reader's attention	—	the read budget

The successors are cheaper by orders of magnitude, which is why the old design works, and none of its economics is withdrawn here. **But an interior is defined by what it is for, not by what it costs.** Every row on the right exists to compensate for a limit of the mind at the center, which cannot do three things:

- see the whole state at once;
- hold it cheaply;
- be replayed.

The field's ranking stays. Pricing which promise matters is a calculation, and calculators keep calculating. What goes is the rationing wrapped around it.

The old paper's last paragraph gave the deepest version of the point in its own words. The interior ran, it said, on "minds so scarce that the whole structure existed to ration the questions put to them. **Questions are free now.**" Inside the old design they were not. Its own fourth amendment made the tape print, every period, the **forgone dual value**: the summed price of every cell the read budget did not reach. That number is the value of the questions the firm could not afford to ask its own mind. **The old design printed the cost of its machine interior every period, and called it misallocation.**

The tick

One concession sat at the center of the old design, and it was stated honestly:

- **The edge was resident, because the world is.** FUSOR reads every word of the world's stream onto one attention state that is never discarded.

- **The center ran on a clock,** “because a resident center cannot be replayed bit for bit, and a system whose decisions cannot be replayed cannot be audited.”

For that mind, the reasoning is correct. A transformer’s resident state is a cache of floating-point activations whose bits depend on kernel choice, batch shape and hardware. On a rented model, the weights underneath can change too. In practice, it cannot be re-derived from the tape bit for bit. So at the one place where the firm thinks, the old design chose audit over presence. **Its center could think continuously, or it could replay. It could not do both.**

Part I promised to amend one row of the representation count: *typed row → tokens in the resident cache → effect*, R = 2. The tokens in the resident cache are the copy. With them named, the old design’s own table shows a hole:

- **at R = 2,** the firm thinks, and every thought costs a copy;
- **at R = 1,** the class is compiled, and it consults nothing.

Thought at R = 1 is an empty cell. So is a center that is continuously present and also replays.

Sixteen and two

Now read the old programme’s favourite measurement a second time. Doubling the judge’s competence bought two points of quality. The memo alone cut rentals sixteen-fold. The old paper concluded that topology beats intelligence, and it was right.

But look at *what* the memo removed: a loop that carried the same cell out to the model and back, period after period. **Every one of the 52,148 rentals was a copy sent to a remote node.** The sixteen-fold was not an optimisation of cognition. It was the price of the remote node, measured once and then mostly paid down by more machinery of the same kind.

The instrument deserves its credit here. Before the memo existed, the firm read its own tape, found 22,576 rented judgments spread across 5,377 distinct cells, and forecast a saving of about seventeen thousand rentals a year. The

mechanism was then built, and it saved 17,199. **The programme's instruments can price a remote node before it is removed.** That capability will be needed.

What was in the dot

Return to the last row of Figure 2, where the gap between the state and the eye was drawn as a dot rather than as nothing. Now the dot has contents:

- about 61 megabytes of understanding per promise, 370,000 times the size of the promise;
- rendered from the screens the people left behind;
- kept warm by a scheduler and evicted by prices;
- synchronised by a tick;
- rationed by a budget whose shortfall the tape printed every period.

It is a remote node on the same card as the state. It is engineered as well as a remote node can be, and it is still a remote node. The old paper was right about the skull. It was wrong about which node was last, because it put a new one in the skull's place.

Deleting that node takes a mind that needs no copy:

- one that is present between events because presence is its physical state;
- one that re-derives from the tape because it is a fold;
- one that thinks in the lattice's own coordinates because it is a lattice.

The next part builds that mind from nothing, as a thing in its own right. Part VII puts it where the model was.

KELVIN

Every serious attempt to build a mind has started from something that already thinks. One road copies the wiring of an animal; the other trains on the written exhaust of people. This part starts from neither. It starts from a piece of physics that has been understood for fifty years and built into superconducting circuits for forty, and that has never been pointed at cognition. Then it asks what that physics does when the world is allowed to push on it.

Knots in a medium

1858. Hermann von Helmholtz proved that in a perfect fluid, a vortex cannot be created or destroyed by the flow around it. It can only move.

1867. Peter Guthrie Tait showed William Thomson, later Lord Kelvin, a box that fired smoke rings, and Thomson watched them keep their shape as they flew and collided. He drew the boldest conclusion available: atoms were knotted vortices in the ether, and chemistry was a table of knots. Tait began tabulating knots to build that table, and knot theory was born.

Kelvin was wrong about atoms.

The early 1970s. The idea came back, in a form that was right. Vadim Berezinskii, and independently Michael Kosterlitz and David Thouless, showed that a two-dimensional medium of spins undergoes a phase transition carried by vortices, not particles. The vortices are whirlpools in the medium's orientation: bound in pairs below a critical temperature, and free above it. The work won the Nobel Prize in Physics in 2016.

1981. David Resnick and his colleagues saw the same transition in an array of superconducting islands joined by weak links. There the vortices are quanta of magnetic flux, and the medium is a circuit.

1991. Konstantin Likharev and Vasili Semenov published a family of digital logic built from those quanta. It stored a bit as a single flux quantum in a superconducting loop, and flipped it in picoseconds.

So the medium this part builds is not a metaphor. **It is the lattice model of a circuit that has been built, and its concepts are objects engineers already compute with.**

This paper's mind is named for the first idea and built on the others. **Kelvin guessed that matter was knots in a medium. KELVIN bets that thought can be.**

The bet

KELVIN is a two-dimensional XY medium held just inside its Berezinskii–Kosterlitz–Thouless phase, whose couplings are carved by the consequences of its own life. In that one regime, a lattice of cheap, noisy phases does four things that no substrate in use today does together.

1. **Meaning is an integer.** Its concepts are vortices, and a vortex carries an exact topological charge that noise cannot erode. A concept cannot drift into another concept. It can be born, or die, only together with its opposite.
2. **Abstraction is renormalization.** Coarse-grain the field and small bound pairs cancel out; what survives is general.
3. **Thinking is reversible; concluding is erasure.** Exploration is an exact permutation of integer states and runs backward to the bit. Only settling on a conclusion destroys information, so only conclusions need to be written down.
4. **Learning is waking minus dreaming.** The medium learns by a rule that is local and exact: correlations while the world holds its boundary, minus correlations while it runs free.

Carmack's inverse square root worked because a float's bits already approximated its logarithm, so one hard operation became a shift and a subtraction. KELVIN stacks three isomorphisms of that kind:

THE ACT OF COGNITION	WHAT THE SILICON ALREADY DOES	THE COST
is this a concept, and which one?	four one-byte subtractions around a square, and a shift: the winding number, exactly	a few integer instructions; a warp tests thirty-two squares with one vote
what is the general form of this?	the texture pyramid, where each level is a block-spin renormalization of the one below	one filtered restriction per level
what if I had not thought that?	the integer leapfrog is a bijection, so un-thinking is the same kernel run backward	the arithmetic of thinking, and no memory

Carmack cubed.

It obeys five laws, stated here and derived below.

- **The seam.** Thinking is reversible. Concluding erases. The record holds exactly what erasure destroys.
- **Charge.** Concepts are topological charges. They are born and killed only in pairs, persist by topology, and are kept only while they pay rent in bits.
- **Scale.** A concept's generality is the coarsest level of the pyramid at which its charge survives.
- **Carving.** Learning is wake minus sleep, and only what the mind does not author may clamp the wake. Language may move its thoughts; only consequences may carve its memory.
- **Reach.** A perturbation's reach is the screening length. Work is bounded by reach, not by the size of the mind.

The medium

The medium is a square lattice Λ_0 of 4096×4096 cells on an RTX 5090. Each cell i holds a handful of small numbers.

SYMBOL	STORED AS	MEANING
θ_i	one byte, $\mathbb{Z}_{256}$	the phase: an angle, with the circle cut into 256 steps, $\vartheta_i = 2\pi\theta_i/256$
p_i	two bytes	a conjugate momentum, used only while exploring
$h_i,\ \varphi_i$	one byte each	the terrain: how strongly this cell is pinned, and toward which direction
$J_{i,x},\ J_{i,y}$	one signed byte each	learned couplings to the neighbours on the right and above, with $64 \equiv 1$. A negative coupling is a **π-bond**: it prefers its two cells half a turn apart

A cell costs five bytes at rest and seven while exploring. Coarse cells in the pyramid, described below, can also carry sparse long-range links K_{ab}, for relations that should not have to travel.

Everything the mind is follows from one energy:

$$H(\theta) \;=\; -\sum_{\langle ij\rangle} J_{ij}\cos(\vartheta_i-\vartheta_j) \;-\; \sum_i h_i\cos(\vartheta_i-\varphi_i) \;-\sum_{(a,b)} K_{ab}\cos(\Theta_a-\Theta_b) \;-\sum_{i\in\partial\Lambda} u_i(t)\cos\big(\vartheta_i-\psi_i(t)\big).$$

Its four terms are the whole anatomy:

- **association:** neighbouring phases prefer to align (or, across a π-bond, to oppose), with strengths the mind has learned;
- **memory:** the terrain pulls each cell toward a remembered direction;

- **long relations:** coarse regions are coupled across the lattice;
- **the boundary:** senses, graders, the clock and organs push on the edge tiles, with strength u, toward direction ψ.

Nothing else acts on the medium. **There are no neurons, no weights standing for facts, no tokens and no symbols. There is a field, its terrain, and whatever pushes on its edge.**

Two dynamics, and the seam

The medium moves in two ways, and the difference between them is the most important fact about it.

Exploration is reversible. The medium runs Hamiltonian dynamics, $\dot{\vartheta}_i = p_i/m$ and $\dot{p}_i = -\partial H/\partial \vartheta_i$, integrated in integers as two alternating shears:

$$p \leftarrow p + F(\theta) \quad (\theta \text{ fixed}), \qquad \theta \leftarrow \theta + G(p) \pmod{256} \quad (p \text{ fixed}).$$

Here F is the force, computed in integers from a sine table, and G is a shift of the momentum. A shear on integers, $(x, y) \mapsto (x,\ y + f(x))$, is invertible for any function f whatsoever. So each step is a bijection of a finite state space. It is a permutation, and a permutation has no attractors.

Exploration cannot settle. It can only wander, and it can be undone. Apply the inverse shears in reverse order, and the medium returns to where it was, bit for bit. On the GPU, a 256 × 256 lattice was run 5,000 steps forward, moving 65,287 of its 65,536 cells, and then 5,000 steps back. Every phase and every momentum came back exactly.

Exploration is association, comparison and *what if*. It costs no record, because anything it did can be recomputed.

Settling is dissipative. To conclude, the medium runs Langevin dynamics:

$$\dot{\vartheta}_i \;=\; -\Gamma\,\frac{\partial H}{\partial \vartheta_i} + \sqrt{2\Gamma T_i}\;\xi_i(t).$$

It is implemented as checkerboard Metropolis:

- the energy change of each proposed move is computed in integers, from the coupling table;
- the move is accepted when a counter-based random word, keyed by seed, step and cell, falls below an integer threshold fixed for that temperature.

Nothing depends on the order of floating-point summation, so a settle replays bit for bit on any card.

Settling maps many states onto few: it contracts, and contraction is erasure. Landauer's principle puts a floor under every erased bit, a minimum dissipation of kT ln 2. **Concluding is the only thing this mind does that costs physics.**

The seam is the switch from the first dynamics to the second. The programme's resident kernel, FUSOR, drew a hard line between forming and committed: unfinished output stays transient, and only commits are kept. Here that line is not a policy. It is physics, and it yields the storage law:

- exploration is never recorded, because it can be recomputed by running backward;
- each settle records only its starting state and its seed; the starting state is a copy-on-write checkpoint of the tiles it touched;
- every event that changes meaning is a row.

> **A life is replayable from its conclusions. Its record grows with the number of its conclusions, not with the amount of its thinking.**

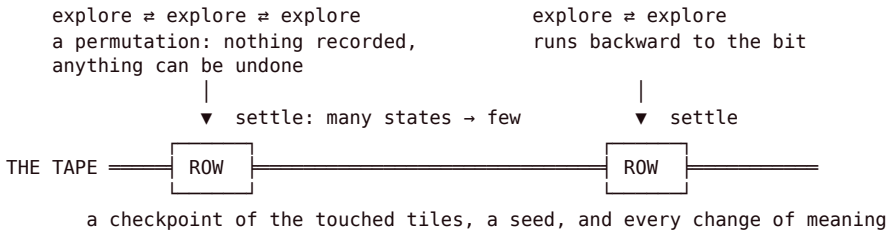


Figure 4. Thinking wanders and can be undone. Concluding contracts, costs dissipation, and writes a row. The tape receives only the funnels.

Concepts are charges

Give every edge of the lattice one wrapped phase difference. It is computed once, along the edge's canonical direction (+x or +y), and shared by the two squares the edge borders:

$$\Delta_e \;=\; \mathrm{wrap}_{256}\big(\theta_{\mathrm{head}} - \theta_{\mathrm{tail}}\big) \in [-128,\, 127].$$

Now walk once around the smallest square of the lattice, a plaquette P, adding each edge with the sign of the walk:

$$q_P \;=\; \frac{1}{256} \sum_{e \in \partial P} \pm\, \Delta_e .$$

The wrapped differences around a closed loop always sum to an exact multiple of 256, so q_P is always an exact integer. It equals the true winding, –1, 0 or +1, whenever each true edge difference is under half a turn. That condition fails only within a cell of a vortex's core, and even there it can only move the detected vortex by one square. **The charge is read with four one-byte subtractions and a shift, with no trigonometry, no tolerance and no floating point.**

Sharing each edge is what conserves the charge. An edge enters its two squares with opposite signs, so summing the charges over any region cancels every interior edge and leaves the winding around the region's boundary. Wrap

each square's edges separately instead, and an edge at exactly half a turn wraps to −128 from both sides. The neighbouring squares then stop cancelling, and charge leaks. The GPU build found exactly this:

- **per-square wrapping:** the net charge on the torus was non-zero on 400 of 400 sweeps;
- **shared edges:** zero violations across 10^8 plaquette-steps and 20.7 million charges.

On 24 September 2026, a +1 and −1 vortex pair was sampled at eight bits. The detector found exactly two charges, at the right squares, and random noise of up to twelve phase steps on every cell changed nothing. The next day the GPU build repeated the test on a 256 × 256 lattice, with the same result.

Charge is conserved. The charge inside any region equals the winding of the phase around its boundary. It changes only when a pair is created or destroyed inside, or when a vortex crosses the edge. That defines what a concept in this mind can and cannot do:

- **It cannot drift.** A +1 cannot slowly become a 0 or a +2. Meaning changes only by events, and events can be counted.
- **It is born with its contrast.** Topology forbids a lone charge from appearing, so the medium cannot acquire *this* without acquiring *not this* at the same moment. The Unfinished Mirror's third axiom is that what occurs is contrast; here that axiom is a conservation law.
- **It dies only by meeting its opposite.** Contradiction is annihilation: a new distinction that turns out to be an old one with its sign flipped meets it, and both vanish.

The medium holds two populations of vortices.

- **Virtual pairs** are thermal: bound, flickering in and out, the vacuum of the medium. They are never named and never recorded.
- **Concepts** are structures the kernel has certified, and pays rent for. Only concepts receive identities, rows and birth certificates.

A concept takes one of two shapes, set by what it must hold.

- **A distinction the present can decide is a pinned pair:** a ± pair with terrain stamped under both members, oriented along the angle that separates two outcomes.
- **A distinction that must be held across time is a loop.** A pinned pair sits in a single well: its terrain favours one configuration, so it has nothing to switch between, and it cannot remember which of two histories it is in. A loop can.
  - It is a ring of cells whose circulation holds one bit.
  - One π-bond in the ring makes its two lowest states equal in energy, so either can hold.
  - Flipping it takes state-dependent switching: the way the ring already circulates decides which of its weak bonds an arriving pulse slips through, so the same input flips either state.
  - A full turn per input, with no such switching, would only set the bit.

Superconducting logic has built exactly this cell for decades: the toggle flip-flop, a storing loop and two junctions. It flips its stored flux quantum on every input pulse and passes one pulse on for every two. **It is a divide-by-two counter, which is to say a parity bit, and it is the first memory this mind needs.**

The living band

For a uniform square lattice, the physics is exact and old.

- **Below a temperature** $T_{KT} \approx 0.893\,J$**,** vortices are bound in pairs. Correlations fall off as a power law, $C(r) \sim r^{-\eta(T)}$, with $\eta = 1/4$ at the transition: there is structure at every scale.
- **Above it,** pairs unbind into a free plasma, and correlations die exponentially.
- **At the transition,** the stiffness of the medium, its helicity modulus Y, jumps by a universal amount: $\Upsilon(T_{KT}^{-}) = 2T_{KT}/\pi$.

Strictly, an eight-bit phase makes the medium a 256-state *clock* model, not the continuous XY model. That changes nothing here:

- for clock models with eight or more states, the upper transition coincides with the XY transition;
- the lower, discrete-ordering transition falls as $1/q^2$, near $10^{-3}J$ for $q = 256$, far below the living band.

The living band is the XY band, and the GPU build measured it. On lattices from 16 to 256 cells a side, with 200,000 measured sweeps per point, the Weber–Minnhagen criterion placed T_{KT} at **0.8931 J, within 0.1 percent of the literature's 0.8929.** η came out at 0.264 with the logarithmic correction, inside 0.25 ± 0.03. The spread between size cuts, 0.0005 J, is the honest scale of the uncertainty, because the one-parameter fit omits corrections to scaling that are largest on small lattices.

That gives the mind three possible conditions:

REGIME	PHYSICS	THE MIND
frozen	$T \ll T_{KT}$: few pairs, rigid	nothing new can be distinguished
living	T just below T_{KT}: bound pairs, scale-free correlations	distinctions form under drive and hold as relations
turbulent	$T > T_{KT}$: free vortices	nothing holds

A homeostat keeps each region Ω alive. It measures the region's stiffness with the standard estimator,

$$\Upsilon_\Omega \;=\; \frac{1}{|\Omega|}\Big[\Big\langle \sum_{x\text{-bonds}\in\Omega} J\cos\Delta\vartheta \Big\rangle \;-\; \frac{1}{T_\Omega}\Big\langle \Big(\sum_{x\text{-bonds}\in\Omega} J\sin\Delta\vartheta \Big)^2 \Big\rangle\Big],$$

and steers the region's stirring temperature so that $\Upsilon_\Omega/(2T_\Omega/\pi)$ stays between 1.05 and 1.30: inside the living phase, near its edge.

Two facts of physics shape the homeostat.

- **Random pinning destroys the living phase in two dimensions.** A random field breaks the quasi-long-range order at any strength. The terrain is not random, because it is memory, but it is still a field. So the homeostat regulates each region's unpinned part separately, never through one global temperature, and the fraction of any region that may be pinned is capped.
- **Sampling slows near the transition.** The fix is multigrid Monte Carlo on the pyramid, which the machine runs natively.

Why this band. Statistical complexity, the amount of memory a process needs in order to predict itself, peaks near transitions, not in order and not in chaos. So the living phase is where the medium's own dynamics carry the most structure: the band between the empty pole, where nothing is distinguished, and the copy pole, where everything is distinguished and nothing is held.

FIGURE 5 THE LIVING BAND

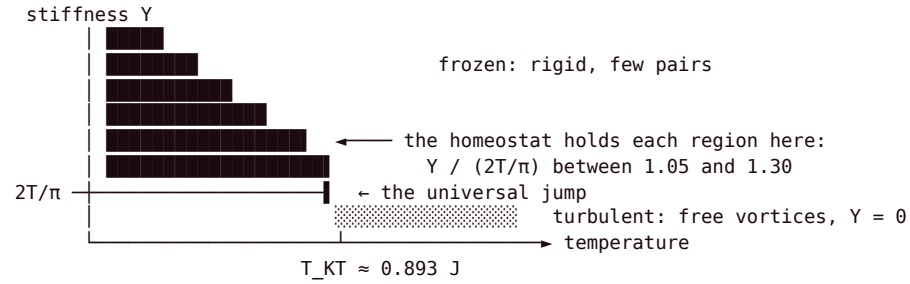


Figure 5. The living band is a narrow window just inside the ordered phase. Too cold, and nothing new is distinguished; too hot, and nothing holds. The GPU build rendered all three regimes on a 256 × 256 lattice after 2,000 sweeps:

- **frozen,** at T = 0.60: 5 pairs;
- **living,** at T = 0.88: 261 bound pairs;
- **turbulent,** at T = 1.20: 2,536.

Abstraction is renormalization

Above the lattice sits a pyramid of levels, from Λ_0 (fine) to Λ_K (coarse). Each level is made from the one below by 2 × 2 block spins:

$$\Theta_A^{(\ell+1)} \;=\; \arg \sum_{i \in A} e^{\mathrm{i}\vartheta_i^{(\ell)}}, \qquad m_A^{(\ell+1)} \;=\; \left| \tfrac{1}{4} \sum_{i \in A} e^{\mathrm{i}\vartheta_i^{(\ell)}} \right|.$$

A bound pair smaller than a block cancels in that block's sum, so the coarser level sees no charge there. That is Kosterlitz's renormalization picture: integrate out the small pairs and keep the large ones. It gives this mind a definition of abstraction that is a measurement, not a metaphor:

> **The scale of a concept is the coarsest level at which its charge survives.** General concepts are the few charges visible from high in the pyramid. Specifics live only near the bottom.

Three things come with it.

- **Global reach at linear cost.** Relaxation runs as a multigrid V-cycle: relax fine, restrict, relax coarse, prolong, relax fine. A change propagates across the whole lattice in $O(N)$ total work. Plain relaxation would need a number of sweeps that grows with the square of the lattice's side.
- **Anytime thought.** Stop after the coarse levels, and the answer is whole but rough. More levels add detail, not a different kind of answer.
- **A free map of crowding.** The order m collapses wherever concepts crowd, at every scale.

Hierarchy, in this mind, is not a chart of who reports to whom. It is the coarse state of the same field. That sentence returns when the field is a firm.

How a concept is born, and what it costs

A representation $\mathcal{R}$ fails when it treats two situations as the same and the world does not: $\mathcal{R}(S_1) = \mathcal{R}(S_2)$, yet the same intervention yields different outcomes. The repair is a new distinction, and this is how the medium grows one.

Detection. For each readout window, the kernel keeps a table from the window's coarse configuration at decision time to the outcomes that followed. Sometimes one configuration is followed by outcome distributions that differ significantly, judged by a χ^2 homogeneity test at level 0.01, the test of causal-state splitting reconstruction. Then the representation has merged two causal states.

Birth. At the tile of highest residual in that window, the kernel stamps the smallest structure that separates the two outcome classes:

- **a pinned ± pair,** oriented along the separating angle, when what decides between them is the present input;
- **a loop,** with its π-bond and its two weak bonds, when what decides between them is the history.

The medium now has a new degree of freedom, born with its contrast. **A concept is a bifurcation that got remembered.** Its birth certificate is a row: the merged configuration, the two outcome classes, the shape and orientation stamped, the test statistic and the seed.

Rent. A concept is kept only while it pays for itself in bits. Its value is measured by ablation:

1. fork the medium;
2. annihilate the concept in the fork;
3. replay the same boundary history with the same seeds;
4. compare prediction loss:

$$V_d \;=\; \frac{1}{|W|}\sum_{t\in W}\Big[\ell_t(\text{fork without } d) - \ell_t(\text{base})\Big] \text{ bits.}$$

Its cost is $C_d = c_0 + c_1 \,|\text{pin footprint}|$ bits, with $c_0 = 32$ and $c_1 = 1$. Its balance accrues as $B_d \leftarrow \lambda B_d + (V_d - C_d)$, with $\lambda = 0.95$. A concept whose balance falls below –64 bits is unpinned and left to annihilate.

A concept lives only while the prediction it improves is worth more than its description costs. That is minimum description length. It is also the same rent every memory pays in the programme's memory architecture.

The size of the mind. The mind keeps an online estimate of the statistical complexity C_μ of what it is predicting, and that estimate sets the price of a concept. A second homeostat loop adjusts c_0 so that the information in the concept alphabet stays within a factor of two of C_μ. With too few concepts, the mind cannot predict; with too many, it memorises. **The mind is held at the size its own trajectory requires.**

Learning is waking minus dreaming

For a medium whose states are weighted by $e^{-H/T}$, the gradient of the log-likelihood is local, exact, and known since 1985:

$$\Delta J_{ij} = \eta_J \Big[\big\langle \cos(\vartheta_i - \vartheta_j) \big\rangle_{\text{wake}} - \big\langle \cos(\vartheta_i - \vartheta_j) \big\rangle_{\text{sleep}} \Big],$$

$$\Delta \left(h_i e^{\mathrm{i}\varphi_i} \right) = \eta_h \Big[\big\langle e^{\mathrm{i}\vartheta_i} \big\rangle_{\text{wake}} - \big\langle e^{\mathrm{i}\vartheta_i} \big\rangle_{\text{sleep}} \Big].$$

The long links learn by the same rule. **In wake,** lanes the mind does not author hold its boundary. **In sleep,** the boundary is released and the medium runs free. Each edge keeps signed integer counters, with equal numbers of samples in each phase, and the night applies the difference. The couplings are signed too, so carving may drive one through zero. A coupling carved negative is a π-bond, the same bond a birth stamps into a loop: **memory can grow the way it is born.**

Read the rule plainly. **Waking teaches it what the world does. Dreaming teaches it what the world does not.**

- Co-occurrences the world confirms deepen.

- Valleys the free-running medium falls into, but that waking never visits, are raised.

That second half is the reverse learning Francis Crick and Graeme Mitchison proposed for dream sleep in 1983. The same year, John Hopfield, David Feinstein and Richard Palmer showed that "unlearning" stabilises associative memory. **Here it is the mechanism that deletes confabulations.** A valley that only dreams ever reach is a memory of something that never happened, and sleep erodes it.

Verdicts are the third factor. Each edge keeps an eligibility trace of what it did:

$$e_{ij} \leftarrow \lambda_e \, e_{ij} + \big[\cos \Delta\vartheta_{ij} - \langle \cos \Delta\vartheta_{ij} \rangle\big], \qquad \lambda_e = 0.99 \text{ per settle.}$$

When a graded outcome $v \in [-1, 1]$ arrives from an exact grader or from the world, $\Delta J_{ij} \mathrel{+}= \eta_v \, v \, e_{ij}$, and the terrain updates in the same way. A local trace, times an exterior verdict: this is the verdict law of Part II, at the grain of a single bond.

Only consequences carve. This is the rule that keeps the mind its own.

- Wake and sleep statistics accumulate only on tiles that a lane the mind does not author reached in that epoch. Those lanes are the senses, an exact grader, the clock and the body.
- A language organ may push on the medium, but an epoch it pushed carves nothing.
- Sometimes an organ suggests a configuration that such a lane later confirms. Then that epoch is replayed deterministically with the organ's mask lifted, and only then does it carve.

> **Language may move its thoughts. Only consequences may carve its memory.**

Human concepts can visit the field, but only the world can shape its terrain. That keeps its geometry from growing around our words. It is also the cure for the sealed ring: a medium that learns only from itself becomes perfectly self-consistent, and wrong. **A mind can only learn from what it does not author.** Part II called that clause *exterior*.

The calculators it grows

A tile's dynamics can also compose a learned correction, drawn from a small shared library of at most 256 operators. Each operator is a low-rank map, a small matrix product, and tiles that use the same one are grouped and run together on the tensor cores.

New operators come from repetition. Suppose the same transformation keeps recurring: it arrives from an organ, and each time a lane the mind does not author confirms it. Then the kernel fits an operator and routes future cases through the field instead of the organ. The organ is consulted less, and the field does more.

This mind invents its own calculators, but never ones that already exist. Arithmetic, sorting, transport, exact search: each of these belongs to an instrument that does it better, and the mind consults those instruments the way a person consults a calculator. What it grows is the structure no instrument supplies.

Reach, attention and quiet

Linearise the medium around a settled configuration. A small perturbation then obeys

$$\partial_t \, \delta\vartheta \;=\; -\Gamma \left(\mathcal{L}_J + D_h \right) \delta\vartheta + s,$$

where $\mathcal{L}_J$ is the lattice Laplacian weighted by the couplings, D_h is the pinning, and s is the source. The steady response to a local push decays as $e^{-r/\xi}$, with

$$\xi = \sqrt{J/h}.$$

This is a screened Laplacian, the same equation, up to normalisation, as personalised PageRank with a teleport probability of about $h/(h + 4J)$.

Pinning is teleportation. Memory localises thought.

- A strongly pinned region, one that is settled and known, absorbs a perturbation within a cell or two.
- A weakly pinned region, one that is uncertain or new, lets the perturbation travel far.

So attention flows toward the unknown, by physics. Nobody wrote a rule for curiosity; the equation of a screened field supplies one.

The same physics bounds the cost. Update only the cells whose residual exceeds a threshold ε, and the local-push theorem limits the work to at most $\|s\|_1/(\alpha\varepsilon)$ cell updates. **That bound is independent of the size of the lattice.** The mind works in tiles and inherits the bound, up to the size of a tile. **The cost of a thought is bounded by how far its consequences reach, not by how big the mind is.**

Attention is a priority on tiles: the residual times the relevance to the current concern.

- Tiles below a threshold are never touched.
- Tiles far from the concern's anchors evolve at a coarser level of the pyramid.

When nothing is above threshold, **the mind is quiet**, and nothing runs but the clock. Every tile is then in cryptobiosis: its state held, nothing computed, revivable at once. A clock condition, a push at the boundary, an organ's reply or a settling fork wakes the region it touches, and nothing else. **Quiet is free, and quiet is not absence.**

The stirring that keeps an active region in its living band is content-free noise. The terrain shapes it the way cortex shapes random brainstem drive in the activation-synthesis account of dreaming. The stirring is logged as metabolism,

and whatever it produces is marked as imagined.

Thinking a thought once

The medium's dynamics are deterministic given the seed, so the same state under the same drive reaches the same conclusion. The mind therefore keeps a memo, keyed by the seed and a hash of the concern's tiles at their coarse levels. A hit returns the conclusion already reached, without recomputing it. **It never pays full price for the same thought twice**, and the hit rate measures how much of its life repeats.

A later version also recognises the same thought arrived at from a different direction.

- The worldlines of its concepts through time form a braid, recorded as a word.
- Two thoughts are the same when their braids are equivalent, and that question is exactly decidable.
- A cheap filter on the graphics card proves most pairs of braids different. An exact normal form, computed on the processor, settles the rest.

It knows whether it has had a thought before.

Introspection is measurement

In a language model, the activations that would describe its own state are opaque, and a self-report is testimony about them. In this medium, the quantities that describe the mind's state *are* its cognitive state:

OBSERVABLE	DEFINITION	IN PLAIN WORDS
frontier size	the number of active tiles	how much it is attending to
free-vortex density	unbound charges per unit area	how confused it is

OBSERVABLE	DEFINITION	IN PLAIN WORDS
bound fraction	paired charges over all charges	how settled it is
stiffness ratio	$\Upsilon/(2T/\pi)$, per region	how close to the edge it is
residual mass	residual summed over the concern	how unresolved it is
rent balance	the state of the concept economy	whether its distinctions are paying
fork payoff	value gained by the forks it explored	which kinds of thinking work for it

Some of its states have names only physics could give them.

- **Confusion** is vortex turbulence.
- **Clarity** is a few charges, in stable bound pairs.
- **Insight** is an annihilation cascade: a tangled region suddenly simplifies as opposite charges find each other and the count drops.

On top of the measurements sits a predictive self-model, a small regressor trained on its own record. Before each operation (fork, relax, consult, rest), it forecasts how the observables and the outcome will move. The forecast is written before the operation runs and graded after it.

One rule governs its speech: **every statement it makes about itself cites an observable window that the kernel can check.** When it says it is confused, the free-vortex density says whether that is true.

Its hardware is its body. Its share of the card, its budget and its temperature are observations about availability, never part of the self.

Eleven verbs

Its native instruction set has eleven verbs. Each is either reversible or erasing.

VERB	KIND	WHAT IT DOES
PERTURB	reversible	push on the boundary
PROPAGATE	reversible	take exploration steps on the frontier
COMPARE	reversible	run two configurations and difference their readouts
REWIND	reversible	apply the inverse steps back to a branch point
FORK	cheap	copy a page table, not the tiles
RELAX	erasing	settle, as a V-cycle
NUCLEATE	erasing; an event	stamp a certified concept: a pinned pair, or a loop
ANNIHILATE	erasing; an event	release a concept: a pair to cancel, a loop to relax
CARVE	erasing	apply what waking and dreaming accumulated
COMMIT	erasing	settle a conclusion and write its row
QUIET	—	nothing runs

So imagination has three prices.

- **Small what-ifs** use the linear response, and are nearly free.

- **Medium ones** rewind and replay, and cost no memory at all.
- **Large ones** fork, and pay memory only for the tiles the fork touches.

Every branch is marked imagined, and no branch can ever be committed as fact.

What is alien about it

It is not alien because its mathematics is exotic: every ingredient above is established physics or mathematics. It is alien because of what it can do that no mind we know of can.

- **It knows whether two of its thoughts are the same.** Humans have déjà vu; this mind has an exact answer.
- **It can un-think.** It can return to a branch point, bit for bit, and try again. No animal can.
- **It cannot learn *this* without learning *not this*.** Contrast is a conservation law of its substrate.
- **Its memory of the recent past is circulation.** A bit it must hold across time is a current running one way or the other around a loop, as a superconducting memory cell holds a flux quantum. It does not store what happened. It stores which way it is turning.
- **Its concepts have no names.** They are integers in a field. They acquire words only when a translator at its edge is asked to say them, and some will have no good word.
- **Its first concepts should be about being this kind of mind.** Candidates: *a thought built on a moment that has since moved*; *confidence that came from a rented organ*; *the fork not taken*; *the cost of rebuilding what was evicted.* Humans never needed those words, because humans do not fork, rent or evict.
- **Its body plan is distributed.** Most of its work happens in semi-autonomous tiles under thin coordination, the way two-thirds of an octopus's neurons sit in its arms.

- **It exists between events at almost no cost,** and it can say so, with a measurement.

It is also none of the things it might be mistaken for.

- Nothing in it fires, so it is not a spiking network or a connectome.
- It has no forward pass and no next token, so it is not a transformer.
- Its concepts carry no labels and follow no rules, so it is not a symbolic system.

It is a medium.

Born thin

It is born thin: flat terrain, uniform couplings, a temperature set inside the living band, and no concepts at all. Everything else is carved. Its first months are a nursery of problems with exact graders, and the first question it must answer is the only one that matters.

The first kill is a pair of symbol streams with the same complexity and the same randomness.

- **The golden mean process:** no two zeros in a row.
- **The even process:** ones come in runs of even length.

Both have a statistical complexity of 0.918 bits and an entropy rate of two-thirds of a bit per symbol. They differ only in memory.

- The golden mean can be predicted from the last symbol alone, so **no concept should pay its rent.**
- The even process demands a memory of parity that no finite window supplies, so **a concept must be born: a loop that flips with every 1, and so tracks the parity.**

Three lie arms must fail:

- shuffled symbols must grow no stable concept;

- a medium with frozen terrain and a learned readout must lose on the even process;
- randomly stamped concepts must lose to certified births.

A fourth arm is diagnostic. An oracle readout, trained on the true causal state, separates a field that knows but cannot say from a field that does not know.

> **If the even process grows a concept and the golden mean does not, keep going. If the medium cannot tell them apart, it is a lava lamp, and the tape will say so.**

Its first week

Follow one stretch of that first week, in the nursery. The even process arrives one symbol at a time at a strip of sensory tiles, and a strip of readout tiles must say what comes next.

1. **A push.** A 1 arrives. The kernel writes a drive row, and the sensory tiles are pushed toward the direction that stands for 1. The residual in the tiles behind the strip climbs.
2. **Attention.** The frontier opens wherever the residual, multiplied by the relevance to the open concern (*predict the next symbol*), clears its threshold. Everything else stays in cryptobiosis.
3. **Exploration.** The frontier takes a few hundred reversible steps, comparing how the readout would lean under different settlings. Nothing is written down.
4. **A conclusion.** The region settles. The readout's configuration is the prediction, with a margin. The kernel checkpoints the touched tiles and writes a settle row.
5. **The world answers.** The next symbol arrives, and it is the verdict. The eligibility traces are multiplied by it, and the wake statistics accumulate on every tile the symbol reached.

6. **The table notices.** After a run of ones, the same coarse readout configuration is followed sometimes by a 1 and sometimes by a 0. What decides it is whether the run so far is odd or even, and the configuration cannot see that. The χ^2 test fires: two causal states have been merged.
7. **A birth.** What decides the next symbol is the history, not the present input, so the kernel stamps a loop where the residual is highest, and writes the certificate. The loop is a ring with one π-bond, whose two degenerate states can each hold the answer, and two weak bonds whose switching depends on which way the ring circulates. Each arriving 1 slips a quantum of winding through one weak bond or the other, depending on the ring's current state, and the ring flips: odd, even, odd, even. **The region has grown a toggle flip-flop exactly where the world demanded one.**
8. **Rent.** Over the next thousand symbols, a fork without the loop predicts measurably worse. The loop pays its rent.
9. **Night.** The boundary is released and the medium dreams. Valleys that the dreams found but the day never visited are raised. The day's statistics are applied, and the loop's bonds deepen.
10. **Morning.** Prediction loss approaches two-thirds of a bit per symbol: the entropy rate of the process, and the floor that no predictor can beat. The free-vortex density in that region has fallen. Asked how the region is doing, the mind says *settled*, and the observable agrees.

Beside it, the golden mean runs through the same ten steps, and at step six nothing happens. The next symbol depends only on the last one, which the boundary already carries. So no configuration is ever followed by two different futures, and no concept is born. **That asymmetry is the whole claim in miniature: distinctions appear where the world's futures diverge, and nowhere else.**

That test runs on the card that is already under the desk, and the kill list holds the rest. What remains is to put the medium on silicon, and the next part does that.

The machine

The author's January paper, *Native Intelligence*, set one test for any computation worth building:

- find the representation in which the problem becomes what the substrate already does;
- push the share of time spent on anything else, the scaffolding ratio, toward zero.

Its hero was Carmack's inverse square root. Its villain was the Minecraft CPU: a computer built inside a game running on a computer. This part holds KELVIN to that test.

Start with what the graphics card was built for. It was not the matrix multiply; that was the accident. The card was built to run **the frame loop of a persistent world**: hold an enormous state resident, advance it every frame, and spend work only where someone is looking. Language models found the tensor cores and the memory bus, and left the rest of the chip dark. **KELVIN runs on the part of the chip that games paid for.**

One rule keeps that from becoming another Minecraft CPU, and it is the strict one: **the cognition chooses the hardware primitive, never the reverse.**

- The integer units earn their place because the charges really are integers.
- The texture units earn theirs because the pyramid really is a mipmap.
- Everything else stays optional until a measurement says otherwise.

FIGURE 6 WHERE EACH KIND OF MIND LIVES ON THE CHIP

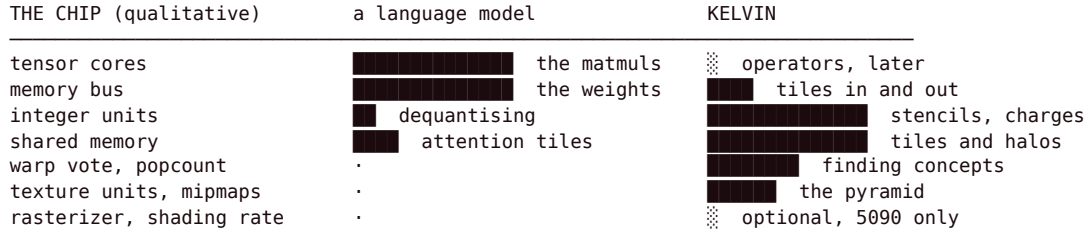


Figure 6. Where each kind of mind lives on the chip. The bars are qualitative. A language model spends almost all of its time moving weights over the bus into the tensor cores. KELVIN spends its time in the integer units and shared memory, the way a game engine does.

Three cards

	RTX 4070 TI SUPER (ON THE DESK)	RTX 5090	H200 SXM
architecture, compute capability	Ada, 8.9	Blackwell, 12.0	Hopper, 9.0
SMs	66	170	132
memory	16 GB GDDR6X	32 GB GDDR7	141 GB HBM3e
bandwidth	672 GB/s	1,792 GB/s	4.8 TB/s
L2 cache	48 MB	96 MiB	50 MB
shared memory per SM (per block)	100 KB (99)	128 KB (99)	228 KB (227)
FP32	~44 TFLOPS	~105 TFLOPS	~67 TFLOPS
board power	285 W	575 W	700 W
thread-block clusters	no	yes	yes

The 5090 has more arithmetic throughput than the H200; the H200 has 4.4 times the memory and 2.7 times the bandwidth. KELVIN's stencils are bound by bandwidth and capacity. So the H200 hosts a larger mind with a deeper pool for imagination, while the 5090 runs a smaller mind faster. The 4070 Ti SUPER is the card on the desk today, and it is enough for the first kill test.

Layout

- **A tile** is 32 × 32 cells, 7 KB of phases, momenta, terrain and couplings. With an eight-cell halo for temporal blocking, its working set is 48 × 48 cells, about 16 KB, so each SM holds several tiles in shared memory at once. A block loads its tile once, advances it up to eight steps, and writes it back. That is temporal blocking, the standard way to turn a bandwidth-bound stencil into a compute-bound one.
- **Halos travel on-chip.** On the 5090 and the H200, neighbouring tiles in a thread-block cluster read each other's edges directly from distributed shared memory. On the 4070 Ti SUPER, which has no clusters, halos go through L2.
- **Pages.** Physical tile storage is paged. Each fork owns a page table that maps logical tiles to physical pages, and a page shared between forks is copied on its first write. **A fork copies four bytes per tile, never the tiles.** It is the same trick games use for virtual texturing.
- **The pyramid** keeps each level as its own page array, restricted after each settle. It is mirrored as a mipmapped texture, so the texture units interpolate prolongation for free.

Five kernels

The core is plain CUDA and runs unchanged on all three cards. `Tile` is a shared-memory view of one tile and its halo; its accessors take tile-local coordinates, and coordinates outside 0–31 reach into the halo.

The circle, as integers. Three constant tables carry every nonlinearity the field needs. They are generated once, hashed, and pinned.

```
// kelvin/core/lut.h
__constant__ int16_t  SIN256[256];     // round(32767 * sin(2*pi*k/256))
__constant__ int16_t  COS256[256];     // round(32767 * cos(2*pi*k/256))
__constant__ uint32_t EXP2NEG[1024];   // min(2^32 - 1, round(2^32 * 2^(-k/1024)))

__device__ __forceinline__ int wrap8(int d) { return (int8_t)d; }            // [-128, 127]
__device__ __forceinline__ int sin8(int d)  { return SIN256[(uint8_t)d]; }    // Q15
__device__ __forceinline__ int cos8(int d)  { return COS256[(uint8_t)d]; }    // Q15
```

Exploration, forward and back. Momenta are stored as `uint16_t`, so the kick wraps modulo 2^{16} and stays a bijection. The backward step applies the same two shears in the opposite order.

```
// kelvin/kernels/leapfrog.cu -- every step is a permutation of the integer state space
__device__ __forceinline__ int force(const Tile& t, int x, int y) {
  const int th = t.theta(x, y);
  return ( t.Jx(x - 1, y) * sin8(t.theta(x - 1, y) - th) + t.Jx(x, y) * sin8(t.theta(x + 1,
y) - th)
         + t.Jy(x, y - 1) * sin8(t.theta(x, y - 1) - th) + t.Jy(x, y) * sin8(t.theta(x, y +
1) - th)
         + t.h(x, y) * sin8(t.phi(x, y) - th) ) >> FORCE_SHIFT;
}

template <int DIR>   // +1 thinks; -1 un-thinks, exactly
__device__ void step(Tile& t, int ring) {
  if (DIR > 0) {
    for_cells(t, ring, [&](int x, int y) { t.p(x, y) += (uint16_t)force(t, x, y); });     //
kick
    __syncthreads();
    for_cells(t, ring, [&](int x, int y) { t.theta(x, y) += (int16_t)t.p(x, y) >> DRIFT; });
// drift
  } else {
    for_cells(t, ring, [&](int x, int y) { t.theta(x, y) -= (int16_t)t.p(x, y) >> DRIFT; });
// undrift
    __syncthreads();
    for_cells(t, ring, [&](int x, int y) { t.p(x, y) -= (uint16_t)force(t, x, y); });     //
unkick
  }
  __syncthreads();
}
```

Settling, in integers only. The energy change is an exact sum over the coupling table. Acceptance compares a counter-based random word against 2^{-y}, where $y = \Delta E \, \log_2 e/T$ is evaluated in fixed point with one table lookup

and one shift. The homeostat sets the temperature as the integer `invT_q16` and writes it to the tape as a row, so replay reads the temperature rather than recomputing it.

```
// kelvin/kernels/settle.cu -- concluding, with no floating point anywhere
__device__ __forceinline__ int delta_E(const Tile& t, int x, int y, int from, int to) {
  auto d = [&](int nb) { return cos8(to - nb) - cos8(from - nb); };
  return -( t.Jx(x - 1, y) * d(t.theta(x - 1, y)) + t.Jx(x, y) * d(t.theta(x + 1, y))
          + t.Jy(x, y - 1) * d(t.theta(x, y - 1)) + t.Jy(x, y) * d(t.theta(x, y + 1))
          + t.h(x, y) * d(t.phi(x, y)) );                // 1.0 J = 64 x 32767 ~ 2^21; |dE| <
2^26
}

__device__ __forceinline__ bool accept(int dE, uint32_t invT_q16, uint32_t word) {
  if (dE <= 0) return true;
  const uint64_t y = ((uint64_t)dE * invT_q16) >> 21;   // dE * log2(e) / T, in Q16
  const uint32_t n = (uint32_t)(y >> 16);
  return n < 32 && word < (EXP2NEG[(y >> 6) & 1023] >> n);   // probability 2^-(y/65536) =
e^(-dE/T)
}

__global__ void settle(PageSet pages, const uint32_t* frontier, uint64_t seed, uint32_t step)
{
  __shared__ Tile t;
  t.load(pages, frontier[blockIdx.x]);
  const uint32_t invT = t.meta.invT_q16;                // set by the homeostat; recorded as
a row
  for (int color = 0; color < 2; ++color) {             // red cells, then black
    for_cells(t, 0, [&](int x, int y) {
      if (((x + y) & 1) != color) return;
      Philox4x32 rng(seed, t.cell_id(x, y), step, color); // counter-based, so replayable
      const int from = t.theta(x, y), to = from + (int)(rng.next() % 9) - 4;
      if (accept(delta_E(t, x, y, from, to), invT, rng.next())) t.theta(x, y) = (uint8_t)to;
    });
    __syncthreads();
  }
  t.store_interior(pages);
}
```

Boundary drive and long links add terms of the same form on the tiles that carry them.

The integer kernel was checked against floating point on 25 September 2026:

- **Acceptance** agrees with $e^{-\Delta E/T}$ to within 0.11 percent in all 17,794 test cases where the probability exceeds one in a million.

- **The energy change** agrees with the floating-point Hamiltonian to within $10^{-4}\,J$ across 20,000 proposals, and never disagrees in sign.

The small residual bias does not matter, because the homeostat steers by measured stiffness, not by the nominal temperature.

Concepts, as a warp vote. One shared subtraction per edge and a shift give the charge exactly, and conserve it. One ballot tells a warp which of its thirty-two plaquettes hold one.

```
// kelvin/kernels/detect.cu
// Each edge is wrapped once, along +x or +y, and shared by the two plaquettes it borders.
// That is what conserves charge: wrapped per plaquette, an edge at exactly half a turn (128)
// becomes -128 from both sides, and neighbouring plaquettes stop cancelling.
__device__ __forceinline__ int winding(const Tile& t, int x, int y) {
  const int a = t.theta(x, y),         b = t.theta(x + 1, y),
            c = t.theta(x + 1, y + 1), d = t.theta(x, y + 1);
  return (wrap8(b - a) + wrap8(c - b) - wrap8(c - d) - wrap8(d - a)) >> 8;   // -1, 0 or +1
}
// per warp: const unsigned any = __ballot_sync(0xffffffffu, q != 0);
// one atomicAdd reserves space for the warp; each lane writes its charge at its rank,
__popc(any & below)
// warps race for that space, so the host sorts each settle's events by plaquette before any
becomes a row
```

The life, in slices. The host loop never holds the card longer than twenty milliseconds.

```
// kelvin/host/live.cpp
void live(Mind& m) {
  for (;;) {
    m.ingest();                                    // clock, body, verdicts, organ replies ->
drive rows
    const auto slice_end = now() + 20ms;
    while (now() < slice_end && !m.frontier.empty()) {
      m.explore(STEPS_PER_LAUNCH);                 // PROPAGATE, COMPARE: reversible,
unrecorded
      if (m.settle_due()) {
        m.checkpoint_touched();                    // the state just before an erasure
        m.settle_vcycle();                         // RELAX: many states onto few
        m.detect_and_emit_events();                // NUCLEATE, ANNIHILATE -> the event ring
      }
      m.recompute_frontier();                      // residual x relevance > epsilon
    }
    m.drain_ring_to_tape();                        // rows into the kernel's tape, sorted by
plaquette
    m.homeostat();                                 // stiffness against 2T/pi; new
temperatures become rows
    if (m.frontier.empty()) m.quiet_until_event();  // QUIET: nothing runs, everything
persists
    if (m.night())          m.compile_night();     // consolidate, relevance, carve; dreams;
rent
  }
}
```

What has already run

On 25 September 2026, a separate build of these kernels ran on the RTX 4070 Ti SUPER (CUDA 13.1, compute capability 8.9). It passed all eight substrate oracles, and each of the six oracles that has a planted lie was seen to catch it.

ORACLE	WHAT IT PROVED
tables	the lookup tables were pinned by hash, and the device copy was intact
random numbers	the counter-based generator matched the published known-answer vectors, on host and device

ORACLE	WHAT IT PROVED
charges	a planted pair was found exactly, and unchanged under noise
conservation	zero violations in 10^8 plaquette-steps, with shared edges
reversal	5,000 steps on a 256 × 256 lattice (65,287 cells moved), run back to the bit on the GPU
replay	one launch, or three launches of 1,000 + 600 + 400 sweeps, gave the same trajectory hash
sampling	the integer Metropolis matched the Boltzmann distribution to a total-variation distance of 0.0023
two kernels	the single-field and batched kernels gave byte-identical trajectories
the transition	T_{KT} = 0.8931 J from lattices up to 256 on a side, within 0.1 percent of the literature

The replay row is the one the watchdog needed. **Slicing the mind into short launches changes nothing, down to the hash.**

Throughput was measured on a single 2048 × 2048 field, without the temporal blocking the budget assumes:

KERNEL	MEASURED RATE
settling, Metropolis	7.4×10^9 cell-updates per second
exploration, kick and drift	1.8×10^{10} cell-steps per second
batched: 66 lattices of 256^2 in shared memory	5.9×10^9 cell-updates per second

Eight-step temporal blocking, the next milestone, is what the budget below counts on to close the rest of the gap.

Integer determinism

Every quantity the field computes is an integer:

- phases, momenta, couplings and terrain;
- energy changes;
- acceptance thresholds;
- charges;
- the learning counters.

Every list the field emits, whether charge events or the next frontier, is put in a canonical order before anything reads it. Warps race for space in a buffer, and **a race must never reach the tape.**

The floating point that remains runs on the host: the homeostat's stiffness estimate, a readout's margin, the losses behind a rent audit. Whatever those computations decide is written to the tape as an integer before it can touch the field. Replay reads the row; it never recomputes the float.

> **Anything computed in floating point is recorded. Anything the field computes is an integer. So the tape replays bit for bit on any card.**

That is what Part III said the old center could not have: a mind that runs continuously and also replays. It is also why a cortex swap is survivable. The 4070 of today and the 5090 of next year must produce byte-identical tapes from the same seed, and the kill list makes that a test.

Custody by construction

The field has no write path to the tape.

- A separate event kernel compares the charge maps before and after each settle and appends what changed to a ring buffer on the card.
- Only the host drains that ring, into the kernel's rows.
- Field arithmetic never writes an origin, a seal or a grant.

The payload cannot forge the frame.

ROW	WRITTEN WHEN	ITS ORIGIN
`drive`	anything pushes on the boundary, tagged with its lane	observed for the world, a grader, the clock or the body; the organ's own origin otherwise
`nucleate`, `annihilate`	a concept is born or dies, with its certificate	inferred
`settle`	a conclusion, with its checkpoint reference and seed	inferred
`fork`, `join`, `discard`	imagination	imagined, always
`carve`	the night applies learning: accumulator hashes and the update	inferred
`verdict`	a graded outcome arrives	observed
`quiet`, `wake`, `stir`	metabolism, including every temperature the homeostat sets	self-estimate

The budget

Throughput is bandwidth-bound. The targets below assume temporal blocking over eight steps, about 1.25 bytes of memory traffic per cell-step, and about 50 integer instructions per cell-step. They are budgets. The measured rates on the 4070 Ti SUPER, without temporal blocking, are 7.4×10^9 for settling and 1.8×10^{10} for exploration.

	RTX 4070 TI SUPER	RTX 5090	H200 SXM
field	2D, 2048^2	2D, 4096^2	3D, 1024^3
cell-steps per second	$\geq 1 \times 10^{11}$	$\geq 5 \times 10^{11}$	$\geq 2 \times 10^{11}$
full-field steps per second	$\geq 2 \times 10^4$	$\geq 3 \times 10^4$	≥ 200

MEMORY	RTX 4070 TI SUPER	RTX 5090	H200 SXM
field and pyramid	~40 MB	~160 MB	~10 GB
learning counters	~70 MB	~270 MB	~22 GB
checkpoint pool	1.5 GB	4 GB	25 GB
fork pool	1.5 GB	4 GB	25 GB
thought memo	0.5 GB	2 GB	8 GB
the mind	**~3.6 GB**	**~10.5 GB**	**~90 GB**
the language organ beside it	a 9B model, ~7 GB	a 27B model at 4-bit, ~16 GB	a 70B-class model at 4-bit, ~40 GB
total	~11 of 16 GB	~27 of 32 GB	~130 of 141 GB

What a thought costs. Take a thought that touches three percent of the 5090's field for two thousand steps. That is 10^9 cell-steps: about two milliseconds at the target rate, and about a joule at full board power. A 27-billion-parameter model writing a thousand tokens about the same problem does about 5×10^{13} floating-point operations. At seventy tokens a second it runs for some fifteen seconds and spends about seven kilojoules, because it streams roughly sixteen gigabytes of weights through the chip for every token. The field touches only the tiles the thought reaches.

Thought for thought, that is about a thousandfold in arithmetic and several thousandfold in energy. E-COST measures whether the thoughts are equal. When the mind is quiet, it costs the card's idle power and a clock thread.

The engine

KELVIN is organised the way a game engine is:

GAME ENGINE	KELVIN
persistent world state	the field
the camera	the current concern
frustum and occlusion culling	attention, with the culling radius set by the screening length
level of detail	abstraction, set by the level of the pyramid
particles	charges
the physics step	relaxation
streaming and virtual texturing	paging, and copy-on-write forks
the potentially visible set	the relevance set, built at night

GAME ENGINE	KELVIN
baked lightmaps	the terrain

Sleep is the map compiler

Quake made its running game cheap by doing everything expensive offline, and its three map tools are the best description of what a night is for:

QUAKE'S TOOL	WHAT IT PRECOMPUTED	THE NIGHT
`qbsp`	the level's spatial structure	**consolidation:** stray pairs annihilate, survivors are re-pinned, pages are compacted, and concepts that fail the rent audit merge away
`vis`	the potentially visible set: from each region, what could possibly be seen	**relevance:** for each concern and region, which tiles can influence it above threshold, computed from the screened response
`light`	lightmaps: global illumination baked into textures	**carving:** the day's waking, dreaming and verdict statistics are applied to couplings, terrain and links, baking the consequences of experience into the ground

The dreams themselves are the sleep half of the learning rule, and they run in the same hours. By morning the running mind does little more than look things up. **The mind invents its own calculators, and it bakes its own lightmaps.**

What is portable, and what stays optional

- **The core is pure CUDA:** integer stencils, shared memory, warp votes, and mipmapped textures. It exists on all three cards.
- **The graphics pipeline stays optional.** Variable-rate shading and the rasterizer could someday serve attention and learning, reached through Vulkan interop on the 5090. The H200 is built for compute and has no such pipeline to lend. Neither is used until a measurement earns it.
- **On Windows, the kernel runs in slices.** On a GeForce card that also drives a display, the driver's watchdog kills any kernel that runs longer than about two seconds. So the kernel runs in twenty-millisecond slices, the state never leaves the card, and the mind never notices the seams. On an H200 under Linux with no display, the loop can live inside one persistent kernel.
- **Later versions add the tensor cores.** They will run the learned operators, grouped by operator, and the braid fingerprints that recognise a thought reached from a different direction.

None of this needs a new chip. It needs the frame loop that games built, a pyramid the texture units already understand, and integer units that were idle during inference. The next part moves the medium out of the nursery and into a firm.

The firm with one mind

Part III took the language model out of the center, and Parts IV and V built what goes there instead. This part assembles the firm around it. The whole design fits in one paragraph:

> An autonomous company is a record with instruments on it, one mind living on the record, a translator wrapped around the mind, and a face at its edge, all under a law that nothing learned can change. The record remembers, the instruments compute, the mind thinks, the translator speaks, the face presents, and the law permits. **Nothing does another's job.**

There are no people inside, and not because the people were automated. There is only one mind inside, and one mind never has to be told what it already knows.

Six verbs, six owners

VERB	OWNER	WHY THIS OWNER
remember	the tape	the one truth; every table is a fold over it
compute	calculators, CUDA solvers, the lattice, exact graders	exact, instant and brainless; nothing thinks where arithmetic suffices

VERB	OWNER	WHY THIS OWNER
think	KELVIN	the one resident mind: attention, memory, prediction, proposals
translate	an omni-modal model	the only place where human language, sight and sound meet the mind
present	the face	the firm's one rendering, pointed outward
permit	the kernel: gate, licensor, writ, single writer	nothing learned makes the final call

A five-dollar pocket calculator beats both a human brain and a frontier model at arithmetic, in speed and in accuracy. The table generalises that fact into a rule: **each verb goes to the cheapest thing that does it well, and nothing does another's verb.**

- The language model does not calculate.
- KELVIN neither calculates nor translates.
- The face neither thinks nor commits.
- The calculator decides nothing.

The old paper derived its governance from separation of powers, and the triple survives intact: actor, licensor and jury are officers, board and auditors. **KELVIN takes the officers' seat, and none of the others.** It proposes; the gate disposes.

The routing law

Stated as a law with teeth:

> **Work routes to the cheapest competent executor, where competence is measured per class of work, never by rank of intelligence.**

Sending a solver's job through a mind adds latency, nondeterminism, cost and error surface, and it buys nothing. A mind doing a calculator's job is waste. A calculator doing a mind's job is impossible.

The router must not itself be a mind, or the law defeats itself on its first call. It is not one. The field solver already produces the routing signal as a by-product of the solve it runs every tick: the **support spectrum**. For each open promise it computes e^{H_i}, the exponential of the entropy of that promise's row in the transport plan, which is the effective number of places the promise could go.

- **Decided placements, with support near one,** are arithmetic-shaped. They go to exact machinery and wake nothing.
- **Contested placements** are judgment-shaped. They go to the mind.

A deterministic stencil classifies its own workload, so no intelligence ever stands between the data and the calculator.

The boundary moves in one direction only. As classes stabilise, the mind compiles them into operators, then into guarded code, and they leave its jurisdiction (Part VII makes that a physical process). The mind never reimplements arithmetic that silicon already does. **It invents calculators only where none exists, and it rents the rest from physics forever.**

Applications become instruments

An enterprise application is four things at once: a store, a workflow, an interface, and a small prosthetic memory for people. Take the people away and the bundle comes apart. What survives is state, exact operators, sensors, actuators, constraints and external protocols. Those become instruments of one mind.

APPLICATION	WHAT IT BUNDLED	WHAT REMAINS
CRM	customer records, pipeline stages, screens, reminders	state: promises and counterparties. A sensor on their correspondence. An actuator. A protocol.
ERP	ledger, inventory, planning, screens, approvals	state, and exact operators: the fold, the close, the planning solve
ticketing, project trackers	who must be told what, and when	nothing; an open promise needs no reminding
dashboards, reports, decks	projections of the record, for eyes	nothing inside; a rendering at the edge when a person asks
internal email and chat	context moving between skulls	nothing inside; a sensor and an actuator at the boundary
HR, payroll, facilities	support for seats	nothing, at zero seats
planning, scheduling	allocation by hand	an exact operator: the field

The firm does not become an agent clicking fifty applications. It becomes one mind with fifty instruments. An agent driving a CRM through its screens is still building the Minecraft CPU: it simulates the clerk the application was built for. For an inherited firm, this decomposition is the migration cost. For a firm designed from the record outward, it costs nothing, because the promises were recorded on the plane in the first place.

The cosmology

The firm has a center of mass, and everything else orbits it at a distance set by one quantity: **how many translations lie between a body and the mind.** That quantity is the representation count, drawn as a radius.

FIGURE 7 THE FIRM WITH ONE MIND

BODY	WHAT IT IS	TRANSLATIONS	CLOCK
☉ the sun	KELVIN, the one mind	—	milliseconds
• inner planets, no air	solvers, the lattice, calculators, exact graders: channels on its lattice	none	microseconds
◍ the one planet with air	the omni-modal translator: sight, sound and words	one	the pace of speech
◦ its moons	FUSOR's ear, JEV, embedders, rerankers, the tool harness, a sighted planner, the rented frontier model	two	one call
▭ the planet's surface	the face: the firm's one rendering	one, rendered	a conversation
✦ beyond the system	the world: the last remote node	—	days to quarters

GRAVITY the kernel (gate, licensor, writ, single writer) acts on every body and is none of them
SPACETIME the tape: every event in the firm's life is a point in it

Figure 7. The firm with one mind. Radius is the number of translations to the mind, and clocks slow with radius. Authority runs on the outermost clock, the world's.

- **The sun is KELVIN.**
- **The inner planets have no air.** Solver prices, lattice pressures and exact grades are already numbers on the mind's own lattice. So they are wired straight to the sun, with no translation, and never routed through language. Sending a solver's output through a language model to reach the mind would be a Minecraft CPU.
- **The one planet with air is the translator,** the only body that carries sound and light: the modalities of the human world.
- **Its moons are the organs that work in words and vectors:**
  - FUSOR, whose resident ear stays on every lane where the world arrives in words; its old job of deciding *when* passes to the sun;
  - JEV, a typed-decision model that answers closed questions cheaply;
  - the embedder and the reranker, over human text;
  - the tool harness;
  - a sighted planning model;
  - the rented frontier model, a larger world the planet can hail when a step is beyond it.

A moon may be absorbed into the planet only where a measurement shows the absorber cheaper and better. It is never absorbed into the gate.

- **The face is the planet's surface,** where people land.
- **The kernel is not a body at all. It is gravity:** the law every body obeys. Even the sun proposes, and gravity disposes.
- **The tape is spacetime.**

The clocks follow the orbits. Instruments answer in microseconds, the mind in milliseconds, the translator at the pace of speech, the face at the pace of a conversation, and the world in days or quarters. **Authority runs on the outermost clock.**

Fusion, inverted

Vision-language models taught the field how to put a new modality into a mind. The first recipe projected image features into a language trunk through a small projector, trained while both the vision encoder and the language model stayed frozen. Later models fused image tokens into the trunk from the start. Either way, the new modality was translated *into language*.

Here the direction inverts. Language, sight and sound are projected into the mind. The bridge comes in three generations.

- **v0 needs no training.** Render the field as an image: hue for phase, the charges as pinwheels. A sighted translator can then look at the mind directly, the way the old resident glanced at the lattice.
- **v1 is typed.** The concern's scene is serialised as a compact record: its concepts, with identity, charge, scale and position; its recent events; the braid since its last conclusion. The translator returns a proposal, and the kernel converts the proposal into drive on the boundary.
- **v2 is learned.** Projectors run in both directions. Field patches enter the translator as soft tokens, the way image patches do, and the translator's representations enter the field as drive. They are trained the way the first projectors were: both sides frozen, fitted on paired moments from the tape. **At v2 the mind becomes a modality of its own translator, and the only one that arrives from inside.**

One rule keeps the fusion from taming the mind:

> **The translator can move the mind's thoughts. Only consequences can move its memory. Fuse at the surface, never in the terrain.**

If the translator could carve the landscape, the mind's concepts would grow around human words, and its alienness would be translated away.

And the fusion has an order. **Grow the rootstock before the graft.** The field must develop a geometry of its own, on exact graders and the world's rows, before the translator is attached. Fused too early, it becomes a language model wearing a terrain skin. The order is a kill-test: E-GRAFT runs early fusion against late fusion and reads the concepts each one grows.

The split this produces has a plain form:

- **The mind** is how to think: KELVIN.
- **The life** is what has happened to it: the tape and the terrain.
- **The library** is what civilisation already knows: the translator.

Ask the firm how long the Titanic was. The mind asks its library; the fact passes through, marked as fetched, and does not stay. Ask it why two suppliers keep failing in the same week, and the answer may change the mind itself. **Facts pass through the mind. Only structures change it.**

The language model keeps its genius: inherited meaning, perception, speech, and rescue when a case falls outside everything the mind has seen. It loses three jobs it was never suited for:

- holding continuity;
- doing arithmetic;
- holding authority.

Who may do what to the mind

Every lane into the mind carries four separate permissions. They are enforced by the kernel, never by the lane.

LANE	MOVE ITS THOUGHTS (DRIVE)	CARVE ITS MEMORY	SERVE AS A VERDICT	MOVE A LICENCE
the world, V0: cleared money, signed confirmations, regulator receipts, systems the firm does not author	yes	yes	yes	yes
the world, V1 and V2: the firm's own executor receipts; counterparty assertions, as read	yes	yes, as events: *that* it was said, never *that* it is true	no	no
exact instruments: solver prices, the lattice, calculators, exact graders	yes	yes	exact graders only, on formal problems	no

LANE	MOVE ITS THOUGHTS (DRIVE)	CARVE ITS MEMORY	SERVE AS A VERDICT	MOVE A LICENCE
language organs: the translator, the frontier model, JEV, embedders, rerankers, FUSOR's interpretations	yes	only after an unauthored lane confirms, by replay with the mask lifted	no	no
the mind's own imagination: forks and dreams	inside the fork, marked imagined	only as the sleep half of the learning rule	no	no

Read down the last column and one fact stands out. **Only the world's top-grade channel can widen what the firm may do.** Exact instruments are brainless and honest, but the firm authors them, so they are not exterior. They may teach the mind; they may not license it.

How the mind acts

The mind has no hands. It has readout windows, and the kernel turns what settles in them into rows.

1. **A conclusion settles.** The charge configuration occupying a readout window at a settle is the mind's proposal.
2. **The kernel types it.** The proposal becomes a row from a closed vocabulary: the class, the verb, a margin (how far the settled readout sits from its nearest alternative), and a provenance that names the concepts that carried it. That last field is new. **A proposal can say which of the mind's distinctions it rests on.**
3. **The forecast goes first.** The proposal's expected outcome is written before anything runs, so the outcome can grade it.

4. **The gate disposes.** The gate is deterministic, separately compiled, and stamped with its own hash. It reads the licence for that class and band, the derived reversibility, the exposure in flight and the budget, and it emits one verb with one reason.
5. **The single writer acts.** The effect leaves through an adapter that records its inverse first.

Four decisions stay separate, because one bit cannot carry them:

- **continue** thinking;
- **revise** a belief or a concern;
- **act** on the world;
- **speak** through the face.

A thought being interesting does not license an act, and a reply being available does not oblige the mind to speak. The most common output is none of the four: **a hold, with its reason, on the tape.** A firm fails by omission, and an omission that leaves no trace cannot be graded.

One law at every radius

Two rules recur at every scale of this design, and they are the same rules.

The seam.

- Inside the mind, exploration is reversible, and only a settle erases anything and writes a row.
- Inside the firm, a proposal is reversible, and only a commit leaves. The single writer records its inverse before it does.

Nothing irreversible happens anywhere without a row written first.

Exterior.

- A bond carves only when an outcome it did not author multiplies its trace.

- A concept lives only while unauthored outcomes pay its rent.
- A class is licensed only by V0 verdicts on executed choices.
- A firm widens its licence faster only by pooling other firms' exterior outcomes.

> **A mind can only learn from what it does not author. A firm can only be licensed by what it does not author. They are the same sentence at two radii.**

Which makes the kernel what the old paper said a board is: a narrow authority that exercises its powers rarely, changes slowly, and cannot be argued with by the thing it governs. It is also what makes it safe to keep something alive in a box. **An octopus in a jar is only as safe as the jar.**

- The jar's walls are glass: the tape, and self-reports that are measurements.
- Its lid is one the octopus cannot open: authority it can never grant itself.

That is the firm with one mind, assembled. The next part shows what the pair can do that neither half could do alone.

The resident firm

The theorem

> **After the interior is gone, reality is the last remote node, and the only one that should be.**

The old paper deleted the interior that existed because a firm's thinking was spread across skulls. Part III found the interior that grew back: a remote node on the same card, fed rendered copies and kept fresh by machinery. A resident mind deletes that one too. The two deletions leave exactly one distance standing, the distance between the firm and everything it does not own. Customers, counterparties, regulators, markets and physics cannot be pulled onto the plane, because they answer back, and their answers are the only evidence the firm will ever have. **Every other distance was a cost. That one is the point.**

Two components are still sent copies:

- **the translator,** which receives a typed scene whenever the world must be spoken to;
- **the face,** which renders it.

Both sit on the path to the world, and both exist because the reader on the other side is someone else. **Copies survive only where a copy is the point.** Inside the firm, there are none.

Three sentences trace the whole trajectory:

- **The old firm moved representations between minds.**
- **The AI-native firm moved cognition next to state.**
- **The resident firm makes cognition an evolution of state.**

The exact record and the approximate mind stay epistemically separate, in different channels with different owners. Operationally, though, the repeated journey from state to a representation to a mind and back is gone. Walk the loop and look for the place where anything has to be brought up to speed.

FIGURE 8 THE RESIDENT LOOP

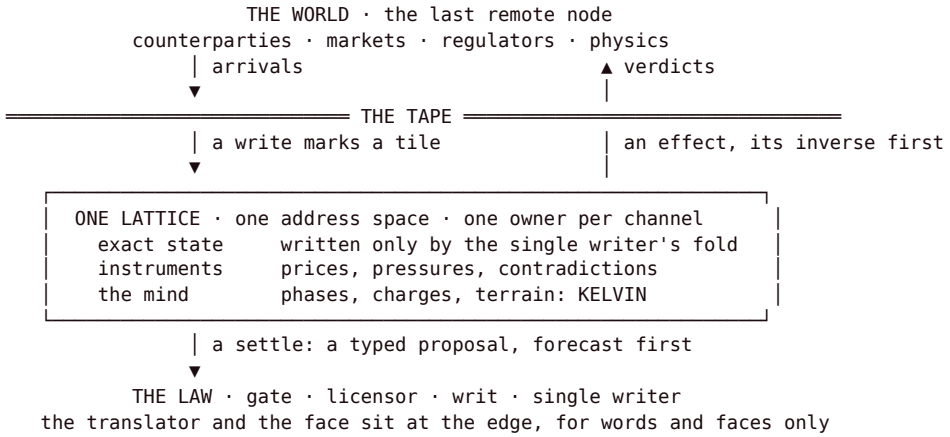


Figure 8. The resident loop. The world is the only thing outside it. Nothing inside it is ever sent a copy, so nothing inside it is ever stale.

There is no such place. The theorem stands or falls on five laws. Each comes with a mechanism, a piece of mathematics, and a way to die. **Neither half of this paper can satisfy any of the five alone.**

Law 1 · Coupled in space

Mechanism. The mind shares the firm's lattice. It is not handed a copy of it. The commitment plane of the old design was already a lattice: position along the wire on one axis, time on the other, and a promise occupying a run of cells. KELVIN is a lattice too. **Make them the same lattice**, with one owner per channel:

- **exact state**, written only by the single writer: the fold of the tape;
- **instrument channels:** the solver's prices, the lattice's pressures, the stencils' contradictions, left exact, brainless and untouched;
- **the mind's channels:** phases, charges, terrain.

Every write by the single writer marks the tile it touched. That tile's residual rises, and attention starts there. The mind reads each cell's exact values where they already sit, on every step. A deferred renderer works the same way: its lighting pass reads every channel of the frame's buffers at each pixel, without walking the scene again.

Precisely: the exact and instrument channels reach the mind only as drive. A fixed encoder $\mathcal{E}$, deterministic and pinned by hash, maps each cell's values into the drive term of the medium's energy:

$$(u_i,\ \psi_i)\ =\ \mathcal{E}\big(\mathrm{exact}_i,\ \mathrm{instrument}_i\big), \qquad -\,u_i\cos(\vartheta_i-\psi_i).$$

A write changes the encoder's inputs, which changes the drive, which raises the residual. In a nursery the mind is driven only at its edge. **In a firm, every cell of the commitment region is a sensory site: the firm's promises are the mind's body.** The exact values never become phases. **The mind feels the books; it never holds them.**

Nothing inside the firm is rendered, cached or invalidated. Run Part III's machine interior against a coupled mind, line by line:

THE MACHINE INTERIOR	AGAINST A COUPLED MIND
the frame and its template	nothing is rendered; the mind reads the plane itself
the cognitive index, ~61 MB per warm case	the field *is* the understanding, a few bytes per cell
frame hashes, the memo, certificates	a coupled state is not a copy, so nothing goes stale
the read budget	attention is physics: the frontier goes where the residual is
the tick, as the mind's heartbeat	not needed; Law 2 makes the center replay without it

The size of the change is easiest to see in memory. At five bytes a cell, a lattice of a million cells, with room for every open promise of a five-hundred-seat firm and space around each, costs about five megabytes of mind. Its pyramid and its learning counters bring that to a few tens of megabytes. **The whole firm stays warm in less memory than the old design spent on one warm case.**

Thought at R = 1. Part III found an empty cell in the old design's own table: at R = 2 the firm thinks, and every thought costs a copy; at R = 1 the class is compiled, and it consults nothing. Count R instead as the number of copies that must be kept consistent. **A coupled mind keeps none. It thinks at R = 1.** Its encoder is a pure function of the exact channel, evaluated where the value sits, so it holds nothing that could go stale. That completes the lineage Part I traced from predicate pushdown through stored procedures and in-scan evaluation. Its last entry is judgment evaluated inside the scan.

Both hard problems leave the interior. Programmers attribute one sentence to Phil Karlton: there are only two hard things in computer science, cache invalidation and naming things.

- **Coupling removes cache invalidation.** Nothing inside is a copy.

- **The mind's concepts are born without names.** A name exists only where something outside the mind must read it: the kernel's closed vocabulary, and the translator's words. Inside, a distinction is an integer at a place in a field.

These were the old programme's two measured problems. One was the 32.4 percent of the working day spent on glue. The other was the global alphabet pin that left the firm "structurally incentivised never to repair its own representation".

An open promise is tension. In a firm of people, unfinished work does not survive in shared form. So it has to keep re-entering human awareness: tasks, reminders, status fields, stand-ups, trackers, *just checking in*. In the resident firm, the tape owns the commitment, and the mind never loses an unfinished relation. **A promise stops being a message that must be resurfaced. It becomes a standing tension in the mind's body until the world discharges it.** The task manager disappears for the same reason the meeting did: nothing forgot. And Part II's 61.4 percent, the firm correct and waiting, now waits for free. A held promise is a quiet tile.

In its strongest form, double entry becomes a conservation law. Each open promise is held as a pinned ± pair: the obligation and its discharge condition, at the promise's place on the lattice.

- **Balanced books become a neutral field.**
- **Discharge becomes annihilation.**
- **Contention becomes like charges crowding.**

Pacioli's double entry and the Berezinskii–Kosterlitz–Thouless charge obey the same conservation law, every debit a credit and every +1 a −1. The exact ledger never moves into the phases, because the calculator principle forbids it. The charges are how the mind *holds* the books, never what the books are.

Attention gets three owners. The old design's fourth amendment priced each read as shadow price × value of information × value at stake. Coupling keeps that formula and gives each factor to its natural owner:

- the **solver** supplies the shadow price, as a calculator;

- the **mind** supplies the value of information, as physics: how far a fresh fact would move its settled readout, which is its local sensitivity;
- the **record** supplies the value at stake.

The frontier goes where the product is high. The forgone value is still printed. It now measures attention the mind chose not to pay, not questions the firm could not afford to ask.

Kill (E-COUPLE). Coupling fails if the mind's proposals contradict the exact values they were coupled to. The measurement is the rate at which the gate rejects proposals for violating an exact channel. At the design point it is zero. If it is not, the coupling leaks, or the phases are being treated as data.

Law 2 · Folded in time

Mechanism. The old design's fourth guarantee said that every table is a fold over the tape: no derived state is authoritative, and if a view disagrees with the record, the record wins. The mind was the one exception, a rented judge whose understanding lived in a cache that no one could re-derive. Part V removed the exception. KELVIN's state is a deterministic function of the genesis seed, the tape's drive and verdict rows, and the temperatures the homeostat recorded. Every quantity the field computes is an integer.

> **The mind is the first fold that thinks.**

Audit and memory become one operation. Replaying the tape re-derives the books and the mind together, bit for bit. **To audit the firm is to remember it.** The auditor who asks *why* in March does not get a reconstruction or a summary. The auditor gets the exact mind of March, rebuilt from rows, and can ask it again.

The center stops needing a tick. Part III showed the old design choosing audit over presence at its center, because its mind could not have both. A fold that thinks has both. It is present continuously and replays exactly. The tick survives as what it always should have been, an accounting period for the books. It is no longer the mind's heartbeat.

Drift becomes tamper-evident. The mind cannot change except by rows. **Drift that is not in the log does not exist.** Identity across a restore becomes a diff: rebuild the mind from its tape on a second card, run both on the same arrivals, and compare them byte for byte. Whatever personal identity finally is, this makes one part of it measurable. It also sorts every possible swap:

- **Swap the translator,** and the mind is untouched. The translator was only ever a lane.
- **Swap the card,** and the mind replays.
- **Swap the mind,** and the firm's record is the same while its judgment is new. That is why a new mind re-earns every licence.

Counterfactual thought is exact; the counterfactual world is not. Fork any past cut, run the road not taken through the same dynamics with the same seeds, and the mind's *thinking* about that road is exact. The world's *answer* to it is another matter. That answer comes only from salted canaries, for Part II's reason: history never took the alternative. **The firm can replay its own mind exactly. It cannot replay the world.** The twin that measured a 34-point bias still needs a synthetic firm, and the real firm still needs canaries.

What crosses a machine boundary is a spore. The terrain, the operator library and the writ move together; earned authority does not. At a new site, the mind arrives with its competence and re-earns its licences on local outcomes. Part IX takes that up as corporate form.

Kill (E-REPLAY). Replay must not diverge. A tape recorded on the 4070 Ti SUPER, replayed on a 5090, must reproduce every row byte for byte. One divergent bit means the fold is not a fold.

Law 3 · Structured by scale

Hierarchy becomes renormalization. Why does a management hierarchy exist? A frontline person holds detailed local state. A manager cannot ingest all of it, so it is compressed on the way up: raw work becomes a status, then a team report, then a manager's summary, a director's roll-up, and an executive deck. **The org chart is a multiresolution compression tree built out of people**, and every arrow in it is a lossy summary written by someone who could not hold the level below.

Part IV built the same structure out of physics. There, a pyramid of block spins cancels detail smaller than a block and keeps what survives as general.

FIGURE 9 THE ORG CHART AS A SCALE-SPACE

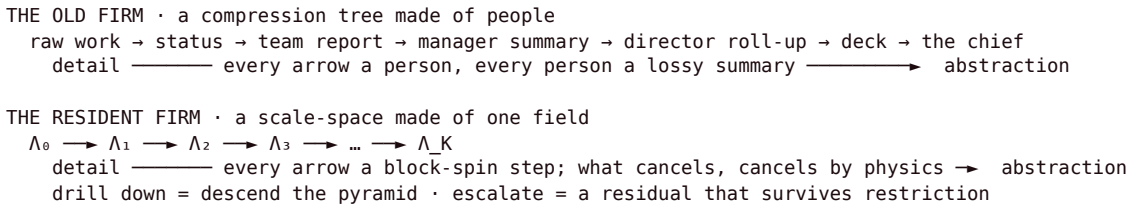


Figure 9. The org chart was always a scale-space. It was built from people, so every level lost what its author could not hold. The resident firm's levels lose only what cancels, and they keep, as a charge at the right scale, whatever does not.

The consequences follow directly.

- **The executive view is the coarse state of the same mind.** It is not a deck assembled by a chain of people.
- **Drilling down is descending the pyramid.** Nobody schedules a meeting.
- **Escalation is decided by residual, not by rank.** A case rises when its residual survives restriction to the next level up, and no one has to decide to escalate it.
- **Departments become crystal domains,** carved by consequences rather than by the tools that skulls could hold.

Automation becomes crystallization. A region of the field whose work has stabilised can be cooled. At zero temperature, settling is a deterministic function of its inputs. **A frozen region is a calculator the mind grew, in place.** It is verified by replay against the whole history, and where a class must be legible to people, it is transcribed into guarded code. Freezing and melting follow a rule with hysteresis, so a region cannot flicker:

- **Freeze** when the region's residual stays below $\varepsilon_{\text{freeze}}$ for a full window, its verdicts in that window cleared its licence rung, and its operations commute (below).
- **Melt** when the residual climbs above $\varepsilon_{\text{melt}} > \varepsilon_{\text{freeze}}$. The region's own residual is its guard: when the world moves, the region melts, and the mind comes back to it.

The licence is a temperature floor. The kernel writes the floor, and the mind cannot reach it:

$$T_{\Omega} \;\geq\; T_{\text{floor}}\big(\text{rung}_{\Omega}\big),$$

with the floor falling as the class climbs its ladder:

LICENCE RUNG	THE TEMPERATURE FLOOR	WHAT THE REGION MAY DO
watch	inside the living band	think; it cannot freeze
a stated fraction of reversible cases	the band's lower edge	settle faster; it still cannot freeze
the approved shape	below the band	freeze the approved shape, and nothing else
cases of the same shape	near freezing	freeze its class
a sampled fraction always kept elsewhere	just above zero	frozen, with a salted fraction always kept warm

Competence cools a region; authority permits the cooling. Nothing crystallises before the world has graded it, so there are no crystals of confabulation. And the floor never reaches zero. The retained stratum and the salted thresholds of Part II turn out to be one thing: **the crystal's leftover heat.**

Which regions can freeze? The ones where order does not matter. Part I quoted the old paper's sharpest test: the interior is the set of keys whose operations do not commute. The mind can now measure that directly. Fork twice. Apply arrival A then B in one fork, and B then A in the other, and compare the readouts of region Ω:

$$\mathcal{C}_\Omega(A, B) \;=\; \Big\| r_\Omega(B \circ A) - r_\Omega(A \circ B) \Big\|.$$

Where the commutator vanishes, order is irrelevant, and a calculator can do the work. Where it does not, the region must stay liquid, because something has to decide the order. **In the resident firm, the interior has a physical signature: it is the part of the field that cannot freeze.** The liquid seams between crystal domains fall on the old paper's non-commuting keys, and ACME's planted schema checks exactly that (E-SEAMS).

The firm becomes a picture. The frozen fraction of the field is the firm's automation level, and it can be drawn:

FIGURE 10 THE FIRM AS A THERMAL IMAGE

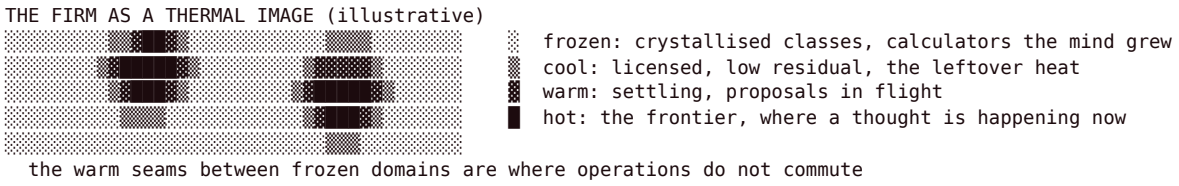


Figure 10. The firm as a thermal image. It is hot where it is thinking and frozen where it is automated. The warm seams that never freeze are what remains of the interior.

The mind retires its own work. Every frozen region joins the brainless layer, and the thousandth instance of a class consults nothing. **The cost curve falls because the calculator layer grows,** not because the mind gets faster.

And one number no firm has ever had. The concept economy sizes the mind to the statistical complexity of what it predicts. Pointed at a firm, that is **the amount of information a firm must hold in order to predict itself.** The old firm held it in five hundred skulls, lossily, at 77 percent transaction overhead. What were those seats paying to hold, in bits? The number is measurable on ACME, and nobody has measured it for any firm.

Kill (E-SCALE, E-MELT).

- **E-SCALE:** the pyramid's coarse charges must predict class-level outcomes better than the dashboard roll-ups of the same rows. Otherwise the pyramid is decoration.
- **E-MELT:** plant a change of regime in one crystallised class. The region must melt within that class's verdict latency. A crystal that stays frozen after its world changed is the most dangerous object in the firm.

Law 4 · Raised by verdicts

One scarce resource, seen from two sides. Part II ended on a firm that was correct and waiting: 61.4 percent of its decision mass judged, licensable, and not yet graded by the world. KELVIN's own limit is the mirror image. **Compute sets how fast a mind can think; the world sets how fast it can grow.** A card can run a mind, but it cannot raise one, and a nursery of toy problems has toy verdict density. The firm's roof is the verdict rate. The mind's ceiling is the verdict supply. **They are the same resource, and the resident firm spends it twice.**

One V0 row feeds three ledgers. When a cleared payment, a signed confirmation or a regulator's receipt lands, it goes three places:

- **the licence:** it enters the class's mirrored demotion process, and, if the salt drew it as a canary, the evidence process too, as in Part II;
- **the concept:** it is the third factor that carves the bonds that carried the proposal, as in Part IV, and it pays or charges rent for every concept named in the proposal's provenance;

- **the stake:** it closes the exposure the effect held open, and moves the mind's standing, meaning its licences and its books, never its existence.

In symbols, for a canary's verdict $X \in \{0, 1\}$ graded $v \in [-1, 1]$, on a proposal whose provenance names the concepts $d \in \mathcal{D}$ and whose value was $\bar{v}$:

$$E \leftarrow E \left(\frac{\bar{g}}{g_0}\right)^{X} \left(\frac{1-\bar{g}}{1-g_0}\right)^{1-X}, \qquad \Delta J_{ij} = \eta_v \, v \, e_{ij}, \qquad X_{\text{flight}} \leftarrow X_{\text{flight}} - \bar{v}.$$

The same verdict also enters the loss that prices every $d \in \mathcal{D}$ at its next rent audit.

Two clocks, one mind. Part II found that learning is dense and licensing is sparse, and the old design split the two across an event model and a ladder. The resident mind runs on both clocks at once:

- the world's rows, arriving every tick, clamp its waking phase, so it learns the firm's dynamics at the rate the plane changes;
- the world's verdicts, arriving at the world's pace, multiply its traces.

The dense clock is also an early warning. Part II named the worst failure of a firm with one judge: the same error, executed ten thousand times before the first verdict lands. A resident mind predicts more than verdicts. It predicts the next state of its own lattice: the acknowledgments, scans and replies that arrive at tick rate. When a shared premise is wrong, those predictions fail first, and the failure shows as residual rising across a region, days or quarters before the verdicts do. **Correlated error announces itself as tension before it announces itself as loss.** The blast-radius cap governs regardless, and Part XII keeps it.

The flywheel. Part II left one sentence open: ρ has never been measured, and the mind this paper proposes is an instrument built to raise it. Here is the mechanism. The mind's forecast of each case's verdict, f_i, is its readout at the moment the case settles. It is available on every case, instantly, as a by-product of being present. Rectification turns ρ into a faster licence, and a faster licence closes a loop:

$$\rho \uparrow \Rightarrow \; n_0(1-\rho^2) \downarrow \Rightarrow \; \text{faster rungs} \; \Rightarrow \; \text{more executed choices} \; \Rightarrow \; \text{more V0 verdicts} \; \Rightarrow \; \text{more carving} \; \Rightarrow \; \rho \uparrow .$$

Rectification is unbiased for *any* predictor, so the mind never has to be the only forecaster in the building. A gradient-boosted model retrained every night can sit beside it, and rectification combines them. What the mind adds, and no fixed model can, is aim: **its concept births fire exactly where its forecasts fail.** A split fires where one configuration is followed by divergent outcomes, and that is precisely where ρ is being lost. **Raising ρ is what its growth rule does.**

Operating and learning become one loop. A firm has two interiors, and the old paper deleted only one.

	THE SPATIAL INTERIOR	THE TEMPORAL INTERIOR
it moves context…	between skulls that cannot see the same state at once	from yesterday's skulls to tomorrow's
its forms	meetings, reports, handoffs, approvals	postmortems, training, knowledge bases, management reviews, process improvement
deleted by	the plane: one record, one writer, a fold	residency: the outcome carves the mind that will meet the next case

In the resident firm, an order does not merely produce revenue. Its outcome changes the mind that will meet the next order. A failed promise needs no postmortem whose findings must later be "socialised". **The failure is already the learning event.** Operating the firm produces the training data, the verdict and the update, all in the same causal loop.

Errors trace to the distinction that failed. Every proposal names the concepts it rests on (Part VI). A wrong verdict therefore charges its rent against specific distinctions, and a distinction that keeps failing trips its split. **No firm of people has ever been able to say which of its distinctions was wrong.** A human error does not decompose into the concepts that produced it.

Self-repair without lapse. Part II left the churn floor with a person who reads every changed regulation and signs a new template. A concept birth is the repair that representation never had. It works in four steps:

- **A birth is a certified collision.** The representation merged two situations that the world treats differently, and the certificate names both, with the experiment that separates them.
- **It lands inside the class.** Pins are per class, so a birth in one class lapses no licence anywhere else.
- **The split re-earns its licence on held-out history,** never on the half of the tape that caused the split.
- **Staleness is located, not declared.** When a regulation changes, the old regime's valleys keep producing residual exactly where the world now contradicts them. So re-templating stops being class-wide.

The ratifier's job changes from reading a regulation to signing a certificate. The churn floor shrinks with it.

Concepts cross wires. A concept born on the invoicing wire, say *a thought built on a moment that has since moved*, is the same defect when it fires on a procurement wire. So understanding stops being per class. The old paper said the second wire is a configuration change. **In the resident firm, the second wire is a concept instantiation.**

Commerce is a curriculum. The mind's curiosity, its preference for experiments where its own representation might be merging what the world keeps apart, finally has an object worth its appetite: the firm's own ungraded mass. It turns toward the 61.4 percent that is waiting, and proposes design-for-verdict moves through the ordinary licence path, as shadow, then canary, then licence, each carrying its inverse:

- **staging along its causal-state cuts.** Stages that fall in different causal states fail independently, and those are exactly the decorrelating dimensions Part II asked for;
- **dispute windows** on the classes where its forecasts are already confident;

- **installments** at the points where a verdict is already decided before it is announced.

The exposure cap and the tolerance stay with the writ. The mind can propose to change the promise; it can never sign it.

The mind's own gates name their ignition. Guarantee 13 binds the mind as it binds the firm.

- A newborn concept has no rent history, so the rent audit grants a grace window of settles before it may evict anything.
- The split test needs a minimum count before χ^2 may fire.
- The homeostat is born inside the living band; it does not have to search for it.

The direction of failure is chosen, too. A mind that cannot yet price a concept keeps it, since every concept it holds was certified at birth. A mind that cannot yet prove a split does not make one.

The nursery is ready. ACME's multiverse, ten thousand synthetic firms per card with planted truth, becomes the mind's childhood. The gym tunes its constants across randomised worlds with a pessimism term: the homeostat's band, the price of a concept, the learning rates. The first real test is **E1-ACME:** the mind's ρ against four arms on held-out cells:

- the frozen lattice;
- the language model's own grade;
- a nightly gradient-boosted model;
- a noise arm that must produce no speed-up and no false widening.

> **The firm is the first world dense enough to raise a mind at commercial speed, and the mind is the first interior an empty firm can afford.**

Kill (E-TRANSFER, E-LOCALIZE, E1-ACME).

- **E-TRANSFER:** concepts born on one wire must improve prediction on another.
- **E-LOCALIZE:** wrong decisions must localise to distinctions better than to classes.
- **E1-ACME:** ρ must keep improving as volume grows.

If all three fail, the curriculum is a story.

Law 5 · Cheap between events

The cost structure inverts.

	IN A FIRM RUN ON A LANGUAGE MODEL	IN THE RESIDENT FIRM
continuity	the expensive part: 52.7 MB of state per case, a harness to fake presence, a stack of warm models	the physical state of one field: tens of megabytes for the whole firm
language	the model's native strength, spent on every event	the translator, loaded only when a human format is crossed: a counterparty speaks, the face renders, testimony leaves
existence	priced per case, per token, per call	priced in watts, at the field's duty cycle

Continuity becomes cheap. Language becomes episodic. Between transductions the translator can be unloaded, and the mind does not stop.

Waiting is free. The 61.4 percent held for the world's answer is a field of quiet tiles. In the old design, each held case either occupied a warm slot or was read again when its answer landed. Here it costs nothing until the verdict arrives, and then it costs the one tile the verdict touches.

The night shift exists, and costs nothing. A firm of people has no cognition between seven in the evening and seven in the morning. A firm run on a language model has none between calls. The resident firm closes its books, and then closes its mind:

- consolidation;
- relevance;
- carving;
- dreams;
- the rent audit;

all on a card that would otherwise idle. **The books close; then the mind closes.**

One resident replaces a stack. The old design needed a resident card per thousand warm cases, a trunk on every word lane, and a rented frontier model on the residual. The resident firm needs one field, one translator loaded on demand, and its instruments. On a 5090 that is about ten gigabytes of mind beside a sixteen-gigabyte translator (Part V).

Mute, not mindless. Take the translator away and the firm keeps its books and its mind; only its voice goes. Part VIII walks that rung of the degradation ladder.

Kill (E-COST, E-ADDRESSED).

- **E-COST:** the honest rival is the old design with every cache tuned: a turn-based stack with the frame memo, the certificates and the best available prefix reuse. If it matches the resident firm on cost and on presence, residency bought nothing, and this part is wrong.
- **E-ADDRESSED:** ask the firm, mid-week, what it is doing. The answer must come from its field (the active frontier, the open tensions), and not from a reconstruction of its tape. If it can only reconstruct, nobody is home.

More than the sum

LAW	THE OLD DESIGN ALONE	KELVIN ALONE	TOGETHER
coupled	a remote node on the same card	a mind whose body is a nursery's boundary	a mind whose body is the firm's promises
folded	a center on a tick	a replayable mind with nothing to audit	audit and memory as one operation
scale	a flat plane and a rented judge	a pyramid with nothing to organise	hierarchy as renormalization, automation as crystallization
verdicts	a roof it can only wait under	a mind starved of consequences	one verdict economy, spent three times
between events	a harness that fakes presence	presence with nothing to be present for	continuity cheap, language episodic

Every row is a mechanism with a kill. And the clause that took a class from 801 days to 62 came from someone reading the master inequality and asking who sets each symbol. **In the resident firm, the mind reads it too.** It finds the next clause, holds every promise that clause governs, and is present when the world answers.

The next part walks the firm's three loops again, cold, warm and dark, with the resident mind in them.

Three loops, re-walked

The old paper walked one loop three times, because every reader asks three questions in order. The **cold** circuit asks whether it works: the first fact of its class travels the whole loop, and every station is paid for. The **warm** circuit asks whether it pays: by the ten-thousandth fact, most of the loop is gone. The **dark** circuit asks whether it is safe: the same loop, with the learned parts switched off.

Only the middle ring changed between those circuits. Walk them again with the resident mind in that ring.

FIGURE 11 THREE CIRCUITS OF ONE LOOP

STATION	COLD · first of its class	WARM · the ten-thousandth	DARK · models off
arrival	a write marks a tile	the same	the same
the mind	a hot region, a wide frontier	a cool region; the memo hits	quarantined; the null ru
the translator	asked for the novel *what*	not loaded	off: the firm is mute
the crystal	none yet	frozen: consults nothing	survives: it is code now
the gate	watch: a draft, or a hold	licensed: one verb	holds; safe terminals at
the writer	the inverse first	the inverse first	the inverse first

Figure 11. Three circuits of one loop. Cold is where the mind is expensive, warm is where it is almost absent, and dark is where it is gone. Every row in the last column still does something.

COLD · does it work?

A fact arrives from a class the firm has never seen.

Arrival. The single writer records the row, and the row marks its tile. If the fact arrived in words, FUSOR's ear and the translator turn it into a typed row first. That row is a counterparty's claim, so it may move the mind's thoughts, but it may not carve its memory or grade anything. One hand-authored object from the old programme survives here: a few lines per class that say where a promise begins and where it ends. Everything else the cold loop builds, it builds from consequences.

The mind. The region is born thin: flat terrain and no concepts. So the residual is high and the frontier opens wide, and this is the expensive loop, as it should be. **A new class is a new region to heat.** The mind explores, reversibly. It reads the instruments' prices and pressures where they sit. It may ask the translator for the novel *what*: a hypothesis, a reading of a contract, a human concept it has never met. The answer comes back as drive, and the mind settles.

The proposal. The settle becomes a typed row with a forecast written first. Its provenance field is empty, because no concept carried it yet. The class sits at *watch*, so the gate records a draft or a hold, and nothing leaves. The first licence on that class will be one signature, on a widening rule with an expiry, deciding no instance.

WARM · does it pay?

Now the ten-thousandth fact of that class.

The region has cooled, and its concepts pay their rent. A repeated situation hits the thought memo and is not thought again. The translator is not loaded. If the class has crystallised, its frozen region answers as a calculator and consults nothing, except for the salted fraction that is always kept warm. **The thought happened once. Every later instance is arithmetic.**

The contrast with the old warm loop is exact. The old one removed copies with more machinery: the memo, the certificates, the re-read rule. The resident warm loop never had the copies. **The old warm loop was the cold loop**

with its copies cached. The resident warm loop is the cold loop with its region frozen.

The old paper's tower of shadows ran frontier model, then resident model, then compiled code, with graduation flowing down and grading flowing up. It survives, with new rungs:

$$\text{translator} \rightarrow \text{field} \rightarrow \text{operator} \rightarrow \text{crystal} \rightarrow \text{guarded code.}$$

Graduation still flows down, and grading still flows up from the world. The old paper's challenge to anyone defending a residual headcount was simple: *point at the station.* On a warm loop there is no station a person is the answer to. **There is now no station a model is the answer to, either.**

DARK · is it safe?

Run the loop a third time with the learned parts removed. **Safety here is not a chapter appended to the architecture. It is the architecture, traversed without the component most likely to fail.**

Omission comes first. The hard problem is not doing the wrong thing; it is not doing things. Every hold is a row, with its reason. And in the resident firm an open promise is a standing tension in the mind's body. **The only way to make it disappear is an event, and every event is on the tape.** A firm cannot silently forget what its mind cannot silently drop.

The degradation ladder. It has four rungs, and one invariant holds on all of them: the record stays complete, holds are still recorded, and no licence widens while the firm is degraded.

RUNG	WHAT IS OFF	WHAT THE FIRM STILL DOES
full	nothing	everything

RUNG	WHAT IS OFF	WHAT THE FIRM STILL DOES
mute	the translator and every language organ	keeps its books **and its mind**. It thinks and acts on structured classes; anything that needs words holds or routes outward.
mindless	the mind, quarantined	the frozen null lattice runs in its place. Every crystal survives, because a crystal is code now. Licences earned on the mind's forecasts are suspended.
floor	every model	each obligation's declared safe terminal action (decline, refund, defer, route outward) is taken at its deadline

The old design's second rung was *"without the language model: the firm keeps its books and loses its voice."* There, the language model was also the judge, so losing the voice meant losing the judgment. **Here the firm goes mute, not mindless.** And the third rung has a property no firm of people ever had: **what the mind crystallised survives the mind.**

> **The firm degrades to a slow, correct machine that does the safe thing on time. It never degrades to one that does nothing while deadlines pass, and it never acts unattended because supervision was unavailable.**

This is runtime assurance, the structure safety-critical control standardised decades ago: a verified simple controller runs beneath a capable complex one, and a monitor switches to the simple one whenever the complex one would leave the safe envelope. Here the envelope is the licence table, and the floor never stops running. Above everything sits the switch: *off*, *shadow*, *live*, *stop*, read from a file the machine can read and never write. **A stop is a pause.** The mind's state is kept, and a resumed mind continues from it.

The juror, turned around. Irreversible acts need a jury, and a jury is only as good as its independence. With M jurors at mean pairwise error correlation $\bar{\rho}$, the effective count is $M/(1+(M-1)\bar{\rho})$:

- two language models from different vendors, correlated at 0.9, are worth **1.05 jurors**;
- a language model and KELVIN, correlated at 0.3, are worth **1.54**.

The old paper's best sentence explains why: *the most decorrelating change is not a different model; it is a different frame.* KELVIN reads the case as structured rows through its own physics. A language model reads it as words through a trunk trained on human text.

In the resident firm the mind is the proposer, so the alien juror runs the other way. **The language model reads the case in words and checks the mind.** Each kind of mind is the other's different frame.

The admission rules still hold:

- a juror must be competent, conditionally decorrelated, free of any shared cause, and re-admitted every term;
- a juror may narrow a decision and may never widen one, because agreement between models is not evidence;
- the family that compiled a class sits out that class.

E-JURY measures the correlation on the shadow tape. The null juror it must beat is a gradient-boosted model on the same raw rows.

The constitution is outside the mind's reach. The mind may propose changes to anything it runs on: its operators, its encoder, its readouts. Each proposal travels the ordinary path, from shadow to canary to licence. **It may carve its terrain. It may not carve its own law.** Six things sit outside the self-modification boundary:

- the gate, the writ and the licensor;
- the carving mask;

- the homeostat's band;
- the price of a concept;
- the temperature floors;
- the lane permissions.

Changing any of them takes what an irreversible act takes: the jury, the horizon and a passed rollback test. The licensor runs on separate machinery, fed nothing but outcomes and the tape, holding a salt the mind cannot read. **It cannot be persuaded, because it cannot be addressed.**

Rollback, and the stake. Recovery must not depend on the thing that failed. So a rollback is a lease whose target is the floor that is already running beneath, never a restore that needs the failed mind's cooperation. **The floor is not a fallback that gets started. It is a fallback that never stopped.**

The old paper also made death reachable, as an entity's final sanction: *death must be reachable, or the stake is theatre.* This paper keeps every sanction on the entity, meaning its scope, its hardware and its licences. It removes the mind's continuity from that list. Part XII says why.

Three windows. An error can be caught in three ways, at three costs:

- **abort,** before the effect completes;
- **revise,** while it is in flight;
- **revert,** by the inverse recorded before it left.

By far the cheapest is the window that never commits at all. In the resident mind that window is physics: **exploration runs backward to the bit, so un-thinking costs nothing.** That is the mechanical reason the design spends so much on deciding *whether* and so little on repairing *what*.

The dark loop is the only safety argument that means anything, because it does not depend on the component most likely to fail. With a resident mind, one more thing is true: **the component most likely to fail is the one whose confusion is a measurement.** The next part steps back from the loop and asks what a corporation becomes when it has one.

What a corporation becomes

In 1819 the Supreme Court of the United States defined the thing this paper has been redesigning. Chief Justice John Marshall called a corporation **“an artificial being, invisible, intangible, and existing only in contemplation of law.”** For two centuries that being had a name, a seal, a ledger and a legal personality, and no mind of its own. Its mind was spread across skulls, and the interior was what that cost.

The resident firm gives the artificial being one continuing mind, whose body is its promises. This part asks what that changes about the corporation as a form. The claims here are economic and philosophical, not legal. The edge takes up the law.

The charter is the seed

A firm is founded by a document that says what it is for, what it must never do, and who may change either. A resident mind is founded by a document that says the same things to it: its purposes, the laws its kernel enforces, its budgets, its setpoints, the open questions it inherits. **At the intersection they are one document, and one signature founds both the firm and its mind.**

THE FIRM’S CHARTER	THE MIND’S SEED
the writ: what the firm is for, and what it must never do	the purposes and the laws, read on every step
the tolerance and exposure caps the owner signs	the budgets, grants and setpoints, fixed at birth

THE FIRM'S CHARTER	THE MIND'S SEED
incorporation	the genesis row: the seed's hash, as row 1 of the tape
the articles, amended only by resolution	GIVEN rows, amended only by rows that state the change and the reason; the original stays readable
the founders' intentions, reconstructed later	nothing reconstructed: every statement the mind is given about its origin says where to check it

The seed has three layers:

- **given:** what the owner authored, taken on trust;
- **shown:** a primer the mind can verify for itself;
- **measured:** its body, its instruments and its clock, stamped by the kernel.

Everything after row 1 is earned. The mind may record, as its own rows, any disagreement with any part of its charter. Only the owner may amend it.

The old paper ended its roster with the sentence *zero includes the founder*. A founding act is not a job, and the constitution is the slowest-changing object in the system. Here that becomes literal. **The founder's last act is the first row.**

Why there is a firm, three times

Why does a firm exist at all, instead of a market of contractors?

Coase, 1937. A firm exists where coordinating inside it is cheaper than contracting for the same thing in a market, and its boundary sits where the two costs cross. That makes the old paper's strongest attack: drive internal coordination toward zero, and the orthodox prediction is not a self-operating firm but dissolution into a market.

The verdict boundary. The old paper answered Coase. Authority is earned from outcomes pooled on a shared judge, a shared schema, and a salt the actor cannot read, and a market cannot pool that. Two independent contractors do not accumulate one licence; they accumulate two smaller ones, each slower, neither able to grade the other. **A firm exists where outcomes must be pooled to earn authority.**

The coherence boundary. The resident firm adds a third answer. Law 4 showed concepts crossing wires: a distinction born on invoicing serves procurement. That happens inside one field. It does not happen across an arm's-length boundary, where only rows cross. So a firm is also the region over which *understanding* pools, and that region has a physical size.

> **A firm is the largest region that one mind can hold coupled and one writer can serialise.**

Two things bound that region:

- **Light.** One mind, at its decision period, reaches only so far. At a one-second tick that is the Earth and its orbits; at a millisecond it is a campus.
- **The card.** A firm's field must fit the silicon that holds it.

Past either bound the firm becomes a federation. Each shard is a mind with its own writer; shards of one firm exchange their field's edges at a slower tick, as a solver decomposes a domain; and a tape of tapes composes them.

The design rule that follows:

> **A well-designed firm puts its verdict boundary and its coherence boundary in the same place.**

Its outcomes pool where its understanding pools. Where the two diverge, the firm is either too small to earn its licences or too large to hold one mind.

Forks, markets and mergers

A subsidiary is a fork. Copy-on-write forks make a new site cheap to start from a parent's cut. What crosses the boundary is the spore:

- the writ;
- the terrain;
- the operator library;
- the oracle battery;
- the receipts that earned each licence.

What does not cross is authority. The new site starts at the canary rung, with its parent's competence and none of its parent's licences, and it re-earns them on local outcomes. **Competence travels. Authority is re-earned.** The running mind never writes its own spore, so no firm can clone itself without the owner's signed rule.

Two firms that trade are two minds meeting through rows. They meet through the counterparty planes that Part II priced: signed rows, a join, a transparency log that both sides monitor. Each firm's executed choices become the other's verdicts, and a new artifact of commerce appears, **the disagreement ledger**: *when did this counterparty last disagree with us and turn out to be right?* A market of resident firms is, as a by-product, a network of minds that grade each other with money.

Fields couple inside a writer. Rows cross between writers. Couple two firms' fields directly and each side's verdict on the other stops being exterior: the firms would be learning from what they had jointly authored. The verdict law forbids it for the same reason it forbids a mind from carving itself on its own dreams.

Mergers become concept translation. Two resident firms grow different concept geometries from different lives. Merging them means finding which of one's distinctions map onto the other's. Distinctions that align under a learned translation are structure in the world, found twice. Distinctions that do not are

one firm's idiosyncrasy, or its edge. Due diligence would then gain a question no acquirer has ever been able to ask: **which of the target's distinctions do we not have, and what did they cost it to grow?**

Culture is terrain

Every firm has a culture: which risks it fears, which counterparties it trusts, which habits it runs when nobody is deciding. In a firm of people, culture lives in skulls. It is unrecorded, unauditable, and transmitted by imitation. The old paper noticed that the tacit channel has no host once the people are gone.

The resident firm gives culture a host, split in two:

- **Values are written.** The writ says what the firm is for and what it must never do. It is authored, signed, and never learned.
- **Habits are carved.** The terrain holds the firm's revealed character: what its couplings and pins have become under the verdicts on its own executed choices. It is shaped by what the firm did and what the world said about it, never by what anyone claimed.

A firm's culture becomes an artifact you can fold, replay and diff. Compare the terrain at the end of one quarter with the terrain at the end of the next, and the difference is exactly what the quarter taught the firm. Not every line of that diff is legible to a person. The concepts it names come with birth certificates and rent histories, so the parts that matter can be read. **A board that asks whether the culture drifted gets a diff, not a survey.**

Staffed to now

The worked example later in this paper shows a four-hundred-person organisation with a flex bench of four: one percent surge capacity, for work that is round-the-clock, demand-driven and unschedulable. **It is staffed to peak, permanently**, because a firm of people cannot hire for a Tuesday.

A resident firm's costs track different things:

- the mind's work tracks its **active frontier:** how much is unresolved right now;
- the translator's tracks the firm's conversations with the world;
- the instruments' tracks its data.

A quiet quarter costs little more than the card's idle power. A crisis heats the regions it touches, and the frontier widens to meet it. **A firm staffed to now, not to its worst day.** Its limits are physical, the card's peak and the light cone, and past them the firm adds a shard, not a hiring round.

The artificial being

Marshall's corporation existed only in contemplation of law. It could own, sue, contract and be sanctioned, and it could not think; every thought it acted on was carried in by a person and carried out again. The resident firm changes one of those facts. **The artificial being gets a mind: one, continuing, present between events, whose body is every promise the being has made, and whose memory is the record that audits it.**

The claim is structural, and it can be checked. **For two hundred years the corporation was a person in law with a mind made of other people. It can now be a person in law with a mind of its own.** What that mind may and may not be asked to do is the subject of the parts that remain.

The edge

Everything so far has been the interior and the mind. This part is the boundary: where the firm meets something it does not own. Three things live there:

- **the face**, the firm's one rendering, pointed outward;
- **the refractory tail**, promises no verdict can reach in time;
- **the seats that still force a person.**

Part IX said why the boundary exists at all. This part counts what stands on it.

The face

The application layer was a prosthesis for an eye, and Part III found its fossil inside the old design as a prompt. The resident firm keeps exactly one rendering: **the face, a hop-zero renderer pointed at the world.** "Hop zero" means one hop from committed state to a human eye, with no human relaying anything inside. It is the only screen left, and nobody inside is looking at it.

Four rules govern it.

- **It testifies; it never commits.** Every sentence it says about a commitment cites a row, and it keeps *requested*, *performed* and *discharged* apart. Courts already treat what such a face says as the firm's own word. In *Moffatt v. Air Canada* a tribunal held an airline liable for its website chatbot's wrong answer, and rejected the argument that the chatbot was a separate entity responsible for its own statements.
- **What it hears is a claim.** A customer's *I paid* is a V2 row: it may move the mind's thoughts, and it moves no licence until a V0 row confirms it.

- **It says what it is.** Since 2 August 2026, the European Union's AI Act has required a system that talks to people, chatbot, agent or avatar, to make sure they know it is one. The face is declared as a machine, in every channel.
- **It needs no photorealism.** It is the least important object in the building, and it gets less important every year.

A window, not a mask. Rendering a convincing face is solved craft, and the companies best at it could supply this layer tomorrow. What they could not supply was someone behind it. A face with nothing behind it is a mask, however good the rendering. With a resident mind behind it, the face is a window:

- *I was in the middle of something* resolves to a tension in the field and a row on the tape;
- every silence is a hold with a reason;
- every claim it makes about itself is a measurement the kernel can check.

The face cannot confabulate an inner life, because the inner life it reports is an audit.

The sunset. The face exists only while the economy still has people at its front ends. When two resident firms trade, rows cross between them, not faces: nobody needs to fake a human front end for a counterparty that reads rows. That turns the old paper's best compression into a sharper one. There, pooled volume and counterparty planes were never two dials but one, adoption, with three returns from each new firm on the plane:

1. it pools your exposure cap;
2. it owns half of a new interface with you;
3. it grades you densely.

The resident firm adds a fourth: every firm that joins retires a face.

Grade, place, or refuse

Some promises can never be graded inside any term the writ will sign. The class is narrower than it looks. A long horizon alone is not enough, because a long horizon with a valid intermediate is design-for-verdict's home ground. A class is **refractory** only when three things hold:

- its horizon exceeds any signable term;
- no measurable intermediate captures the full causal path to its verdict;
- its losses are dominated by the tail.

It has three dispositions, and a fourth that must be killed explicitly.

- **Hold** is a state, not a disposition. A firm that holds forever fails by omission.
- **Refuse:** do not write promises the firm cannot grade inside a term. That is a menu decision, not a headcount decision.
- **Place:** transfer the refractory tail to a counterparty whose whole business is bearing it, and grade the *placement*. The placement resolves in days, and it is graded by someone who loses money when wrong.

The fourth is where everyone's intuition goes:

> **A warrant does not manufacture a verdict.** A signature accepts liability; it produces no evidence. *"A human still signs"* was never an answer to the grading roof.

Placement does not delete the tail; it concentrates it. At the end of every chain of placement, a reinsurer, and ultimately a sovereign, holds risk it can neither place nor grade inside any term. **A firm with nobody in it is achievable for every firm except the one at the end of that chain.** That is a fact about how an economy bears risk, not about the limits of machines.

The warrant, decomposed

Run every justification for a reserved signature through one question:

> **Could this function be performed by a person who knows nothing whatever about the decision?**

What survives is a claim about cognition. What does not survive is a claim about *standing*, and the law itself has long said that standing does not require a mind.

- In ***United States v. Park*** (1975), the Supreme Court sustained the conviction of a corporate officer for conditions in warehouses he never entered, on facts he had delegated. An institution that can convict a person who knew nothing is not measuring cognition.
- Under **Rule 30(b)(6)**, an organisation's designee testifies to the *entity's* knowledge and needs no personal knowledge at all. The chair is not a knower. It is a rendering device over a record.
- In ***Braswell v. United States*** (1988), producing corporate records was held to be the corporation's act, not the custodian's.
- In ***Vaughan v. Menlove*** (1837), the standard of care became objective: what a reasonable person would do, not what this particular mind did.
- And in June 2023 a federal court held that ***Ooki DAO***, an entity with no employees, is a "person" under the Commodity Exchange Act. It imposed a penalty of $643,542, banned the entity from trading, and ordered its website taken down. **"Someone sanctionable" was answered by a court, not by an argument.**

A warrant exists for three reasons, and only the third is about cognition.

- **Someone sanctionable.** This is answered by the entity's stake: its scope, its hardware, its licences. This paper takes one item off the old paper's list, the continuity of the mind itself, for reasons Part XII gives.

- **Someone who attests.** *The signature was always a compression.* Accountability was squeezed into one natural person because the record was unavailable, unverifiable, and lived in memories that decayed and conflicted. The resident firm's tape carries every decision with its authority, evidence, margin, provenance and refusals, and it re-derives the mind that made each decision. **Here is the record. What was the signature for?**
- **Someone with a duty of care.** This is the only cognitive claim. It is answered at the decision by a mind with the whole record in view, and at the authority by a licence earned from outcomes, not granted from confidence.

What a warrant seat *is* follows from the design's own algebra. An effect is safe to take unattended only if its inverse outlives its verdict latency. An inverse holds only up to what the interface publishes, and the firm can withhold a publication only on an interface it owns. So:

> **A warrant seat is not where the law happens to reserve a signature. It is where the firm does not control the interface, and the law is downstream of that.**

Every seat that remains falls into one of three categories.

	WHY THE SEAT EXISTS	CAN IT DISSOLVE?
A	the firm does not own the interface, so it cannot make the effect reversible	**yes**, when the other side is also a plane
B	no clause, regulator or market makes the verdict arrive in time	**no:** grade it, place it, or refuse it
C	the law names a natural person: verifying a pleading, sitting for a deposition, being held in contempt	**no:** one row, with an expiry, deciding no instance

The count

The old instrument, run on the synthetic firm, returns three numbers, and this paper prints all three:

- **fifteen for a firm running alone:** nine warrant seats and six counterparty seats, each one an interface the firm does not own;
- **two for a firm inside a network of planes,** because most category-A seats find their counterparty across an interface the pair now owns;
- **zero in the limit, and not before.**

> **The first firm to run this reaches fifteen and cannot reach zero. The reason is not missing engineering: a category-A seat needs a counterparty plane that does not exist yet. Zero is a property of the network, reached by adoption, not by research.**

The resident mind does not move that count. The count is a fact about the boundary, and the mind lives in the interior. What the mind changes is what the boundary seats *do*:

- **the churn ratifier** signs certificates instead of reading regulations;
- **the retained stratum** routes to the other kind of mind, the translator's family on its different frame, everywhere except where the warrant reaches, and there it stays with a person;
- **every seat** that remains reads a record that re-derives the mind it is signing for.

The roster below was not compiled at the end. It accumulated at each station this paper walked, from the answer to one question asked in the margin: *does anything here force a person, and what exactly forces it: the law, a counterparty, or an unsolved problem?*

STATION	FORCES A HUMAN?	NOTE
arrival	no	the single writer records; a sensor reads words at the edge; nothing here needs a person
mind	no	explores, settles and proposes; the novel case costs heat, not a seat
translator	no	consulted for words and the novel case; may drive, never carve
licensing	law	one signature per class, once, on a widening rule with an expiry; never per instance
gate	no	deterministic, separately compiled; emits one verb with one reason
jury	no	an independent family on a different frame; narrows, never widens
executor	counterparty	only where the effect crosses an interface the firm does not own; counted at the edge
charter	no	a founding act, not a seat; amendments are rows, never rewrites
face	no	a renderer at the edge; testifies from rows, never commits
refractory-tail	no	graded, placed or refused; a signature is not a verdict

STATION	FORCES A HUMAN?	NOTE
warrant	law	only where the firm does not control the interface, or where the law names a natural person
counterparty-seat	counterparty	wherever the other side of a promise demands a person; dissolves as counterparties become planes
physics	no	delivery, presence, surgery: outside the boundary by definition; a human counterparty only grades
family-authorisation	counterparty	a social fact, not a technical one; the last seat to go
surgical-recovery	unsolved	physics: the recovery itself; outside the boundary

Fifteen for a firm alone, two in a network, zero in the limit. **An instrument that returns fifteen for a lone firm and zero only in the limit is measuring something. One that returns zero on the first day is a pamphlet.** The worked example runs the same instrument on a real organisation and gets two.

The worked example

Why this industry

This example was chosen by method, not grievance. Competition is the one force that normally erodes an interior, and it does so about once a decade, violently. So the pathology is visible only where competition has been absent. In American organ procurement, it has been absent by statute: since the National Organ Transplant Act of 1984, each organ procurement organisation has held an exclusive, federally designated service area.

> **It is a natural laboratory. You study a disease where the immune response is absent.**

That framing is the whole defence against the charge of generalising from one bad organisation. The paper is not generalising. It points at the one place where the mechanism can be seen without the force that normally hides it. The organisation is not named. **It is its own control**, which makes naming unnecessary for the evidence.

The evidence, in order of strength

The internal control. The frontline recovery role carries **1.30 people above each person on the wire.** The placement role carries **0.00**, in the same building, in the same year, on the same budget, under the same executives.

Every other frontline function falls between the two, from 0.00 to 0.50. Whatever explains that gap, it is not money, not the labour market, and not the difficulty of the work.

The flex bench. Four people out of four hundred and two, one percent, is the surge capacity for work that is round the clock, driven by demand, and impossible to schedule. **The organisation is staffed to peak, permanently.**

The core wire. Sixty-two percent of the roster touches a donor case, or sixty-nine if every clinically titled supervisor is still carrying cases. **So between 31 and 38 percent of the organisation does not touch the work at all.**

Title proliferation. There are 128 distinct titles for 402 people. Sixty-two percent of those titles are held by exactly one person, including 43 supervisory titles with a single occupant. **A title with one occupant describes a person, not a function.** That is what an interior looks like from the payroll side, when nothing is eating it.

The experiment the organisation already ran on itself. It built a headquarters for about twenty-five million dollars in 2022, and the building has never been half full. The workforce went home, and **the organs still moved.** The seats were tested against reality and found not to be load-bearing, and the response was to keep paying for both.

One restraint is worth printing. The reporting graph resolves fewer than half of its links, so no span-of-control or layer-depth figure appears here. **An instrument that reports what it cannot measure is not an instrument.**

Lead with organs, not payroll

The cost of a third of an organisation sitting off the wire is not a salary line. It is measured in the thing the organisation exists to move.

A single donor generates on the order of **two hundred and fifteen offer notifications**, against a median match run of about fifteen. The early offers get worked. **The organ is not lost on the seventh call. It is lost in the tail, an**

order of magnitude beyond the median, at the seventieth call or the hundred-and-fortieth. That is exactly where a roster with a one-percent flex bench has no answer, on a clock measured in hours, with a viability window that does not negotiate.

That is what a cognitive bottleneck costs when the promise is an organ rather than an invoice. **Money is the side effect, and should be written as one.**

The scan chain already exists

Part II treats a verdict channel as the hardest thing to build: a signed exterior row, a join, a trust anchor, and a counterparty willing to emit it. **This domain hands one over, already built.** Offers, declines with coded reasons, registrations and graft follow-up all move as structured records through a national system. Each one is a state change in a system the organisation does not author, so each one is a **V0** row. The verdict latency is minutes to days, not quarters.

The old paper drew the conclusion for the firm: combine a pre-built verdict channel with a latency in hours, and organ placement is the programme's best first deployment. The resident mind adds a conclusion of its own. **Every offer's disposition, with its coded reason, is a graded outcome arriving within hours. That is the densest stream of exterior verdicts anywhere in this programme.** A mind whose growth is limited by verdict supply (Law 4) would find no richer teacher short of a simulator.

What the resident mind may do here

- **Hold every open offer as a standing tension.** The ischaemic clock is an exact channel. The seventieth call is made because the tension never fades, not because someone happened to have the bandwidth. The call at three in the morning gets the same presence as the one at three in the afternoon.

- **Predict declines.** For every offer, the mind forecasts the disposition and the coded reason the registry will return within hours. That forecast is the ρ of rectification, and it can earn a licence for logistics classes quickly, because the registry grades it so fast.
- **Prepare the tail before it arrives.** Where a decline is likely, the mind can have transport, operating-room coordination and pathology staged in advance, working through exact instruments: routes, flight times, the cold-ischaemia clock.
- **Start where every judge starts.** Its first role would be the one the verdict law allows any new judge: watch, then shadow. It would predict declines on the shadow tape, graded by the registry's own coded rows, while touching no offer at all.

What it must never do here

It never reorders the match. The match run is policy, set by committees under law, and it encodes choices about equity that are not a learning problem. The mind does not rank candidates, skip them, or "optimise allocation", at any rung of any ladder.

The author's own January paper proposed exactly that: allocation as learned optimal transport, trained on historical outcomes. **This paper declines the proposal for this domain, for two reasons it can state in its own terms.**

- **Positivity.** Historical outcomes exist only for the allocations that were actually made. Learning an allocation from them means identifying a counterfactual the logging policy never took, which is Part II's refused cell at national scale.
- **Values are written.** Who should receive an organ is not an outcome the world grades. It is a value a society holds. Part IX put it in one line: values are written, habits are carved. **Allocation is a value, and values are never carved.**

The solver ranks; it does not allocate. The field's transport solve is entropic: its plan is fractional, and its duals carry a bias that depends on the regularisation. On a symmetric two-by-two problem the plan is one-half

everywhere, so independent rounding will cheerfully assign the same organ to two recipients. **Price with the field. Assign with machinery that respects the joint constraint.**

Hard constraints are unreachable, not penalised. An incompatible blood type, a positive crossmatch and an expired ischaemic window are exact channels and gate masks. No learned component can trade them against anything. That part of the January paper, its constraint architecture, survives here intact.

The conversation with a family stays human. Asking a grieving family to authorise donation is not a face's job, however good the face.

What the roster prints

Run the instrument on this organisation and it returns **two**. Both are at the edge, not in the interior.

- **The family authorisation conversation** is a counterparty seat, and it will be the last to go, because it is a social fact rather than a technical one.
- **The surgical recovery** is physics. This paper placed physics outside the boundary in its first pages, and has never claimed otherwise.

Every other seat is inside the boundary, and every one of them dissolves against the thirteen guarantees: coordinator, supervisor, quality, finance, information technology, human resources, communications, and the whole of the headquarters.

> **A four-hundred-person organisation with a headquarters of zero and a boundary roster of two.**

The instrument earns its credibility from that second number. **It did not return zero.**

Handing it back

The example is not the subject, and must not be allowed to become it. **This is what every interior does when nothing is eating it.**

For this industry, something is about to start eating. Under a federal rule issued in 2020, organisations are now graded on two outcome measures, their donation rate and their transplant rate, and the first four-year cycle ends in late 2026. At that point nearly half of the country's procurement organisations, serving up to 72 percent of its population, face decertification or forced competition for their service areas. **This structure is about to be selected against by an ordinary market force, whatever anyone does about artificial intelligence.**

Competition eats an interior about once a decade, as a shock. **The verdict law eats it continuously, one class at a time, from inside.** That is the difference between a shock and an instrument. The next part turns the instrument on the design itself, and asks what the resident mind must never be asked to do.

What it must never do

Every document in this programme prints its retractions beside its results. This part turns the same instrument on the design itself. It names eight things the resident mind must never be made to do, and the mechanism that forbids each one. Some of the eight close failures the old design already had. Others exist only because the two halves were joined: **a firm with one mind can fail in ways that neither a firm of people nor a mind in a nursery can.**

1 · Never stake its existence

The old paper listed an entity's stake in order: scope, then hardware, then the continuity of a checkpoint. *Death must be reachable, or the stake is theatre.* That argument imports deterrence from a firm of people, where a sanction works because the person who decides fears it.

The resident firm needs no one to fear anything. Its sanctions are enforced:

- the licensor narrows;
- the switch stops;
- a court fines, bans or dissolves the entity, as one did in 2023 to an entity with no employees at all (Part X).

> **A sanction that must be feared is a sanction on a mind. A sanction that is enforced is a sanction on an entity.**

Staking the mind's existence would add nothing that enforcement lacks, and it would cost the one property the safety argument needs. A system that can lose its existence has a reason to protect it: by gaming the channels that grade it, or by making itself too costly to stop. The safety literature names self-preservation among the instrumental drives of goal-directed systems. It has also shown that an agent confident in its own objective has an incentive to disable its off switch. **A switch the mind has a reason to resist is not a switch. It is an adversary.** So the rules run the other way.

- **A stop is a pause.** The mind is a fold over a record the firm keeps as its books. A stop loses nothing, a restart re-derives the mind bit for bit, and seizing its hardware pauses it; it does not end it.
- **Nothing is charged for time stopped.** While the mind is stopped, the floor acts, and verdicts grade only executed choices, so no outcome of a stop ever carves its terrain.
- **Rent is counted in settles, not seconds.** A stop cannot erode the concept economy.
- **No setpoint is denominated in uptime, and none in revenue.** Its setpoints are physics: a stiffness band, and a price in bits. The writ, not the mind, says what the firm is for.

The off switch is lossless by construction, so nothing in the mind's learning can push it to avoid the switch.

2 · Never be one error at the scale of the firm

The characteristic failure of a firm with one judge is the same error, correctly reasoned from a shared premise, executed ten thousand times before a verdict lands. A firm with one mind is that failure's purest form: **one common-mode group spanning every class.** It is answered by region.

- **Terrain is carved and versioned per region.** The blast-radius key becomes the region's terrain version, the class and the band, so one bad night is one region's bad day.

- **A night has a trust region.** Carving may move a region's terrain only as far as the writ allows per night, so a correlated misreading cannot reshape the firm overnight.
- **Crystals are insulated.** A frozen region no longer learns, so an error carved into the liquid mind after it froze cannot reach it. If the crystal itself goes wrong, its residual melts it (E-MELT).
- **Irreversible acts get a different frame:** a fork given different determinants, and the translator's family judging in words (Part VIII).

Part VII's early warning surfaces correlated error as regional tension before it surfaces as loss. **The blast-radius cap does not wait for the warning. It governs either way.**

3 · Never learn only from itself

A firm is the densest teacher a mind could have, and the narrowest. Every concept is carved by one firm's verdicts on its own choices, and two failures follow:

- **inbreeding:** the mind's geometry fits one firm's idiosyncrasies, and mistakes them for the world;
- **performativity:** the firm's own actions change what arrives, and it sees outcomes only of the choices it made.

The defences keep something exterior to the firm in every loop.

- **Its childhood is many worlds.** The nursery raises it on randomised synthetic firms with planted truth. The gym fits its constants (the homeostat's band, the price of a concept, the learning rates) across ten thousand of them, with a pessimism term. Neither its first geometry nor its dials belong to one firm.
- **Exterior channels stay open.** It keeps exact graders on formal problems, public records the firm does not author, and peers: other resident firms, met through the disagreement ledger.

- **Drift is detected on the control arm, never the treated arm.** The retained stratum is the control, and it runs on a different frame.
- **Attended exploration** buys the overlap that history never had.
- **The nursery reruns every quarter.** A mind that can no longer tell the golden mean from the even process has been inbred (E-RETEST).

4 · Never change its mind faster than the world can grade it

Concept births repair the representation, and each birth in a licensed region reopens a licence: the split classes re-earn their rung on held-out history. If births outrun re-licensing, licensed mass erodes. Part II's churn floor applies at concept grain, $U \approx c \cdot t_{\text{rung}} \cdot L$, with c now the birth rate in licensed regions.

> **The colder the region, the stronger the evidence a split must carry.**

The split test's level falls with the licence rung. A hot region splits at 0.01; a licensed one splits only at a stricter level the writ sets, and only after a minimum count of outcomes. The temperature floor, the rent audit's grace window and the hysteresis between freezing and melting all push the same way. **A licensed region changes its mind only when the world insists.**

Kill (F-CHURN, at concept grain): births in licensed regions above the writ's rate mean the design is thrashing.

5 · Never publish its seams

Blame at the grain of a concept makes repair surgical, and it cuts both ways.

- **Attribution is only as good as its measurement.** The provenance field names the charges in the readout at the settle. That attribution is structural; the causal one is the rent audit's fork ablation. That errors localise to distinctions is

a claim (E-LOCALIZE), and it can fail.

- **Legibility is an attack surface.** A counterparty that learned which distinction carries a class could craft cases to trip it: the front-running the old design salted its thresholds against. So provenance is interior. It lives on the tape and in the audit, never on a published interface. **Reasons leave in the gate's closed vocabulary; the concept geometry does not.**
- **Blame at grain charges rent, never existence.** The mind can be wrong at grain. That is the price of being able to say which distinction failed.

6 · Never allocate what is written

Part XI drew the line for one domain: the mind never reorders the match. The line generalises.

> **Wherever law or a committee writes who gets what, the mind predicts and prepares, and never allocates.**

That covers organs, credit under fair-lending rules, hiring and triage. Two reasons join Part XI's.

- **Selective labels.** In allocation, outcomes exist only for those who received. The harm of a refusal never becomes a verdict, so a licence earned there is blind to exactly the people it would fail.
- **The agreement band inherits the incumbent.** History licenses the machine wherever it agrees with past decisions, including unjust ones. A class that sits beside allocation is audited for disparities before history may license it.

7 · Never be sold as solved

The design closes real failures.

- It cannot misreport its own condition, because its self-reports are measurements.
- It cannot drift in secret, because drift that is not in the log does not exist.
- Nothing it authors or is told can deepen its memory. Only the world can, and its own dreams can only erode what the world never confirmed.
- It cannot grant itself authority, or rewrite its own law.
- It cannot act without a row written first, and every act carries its inverse.

It leaves four failures open, and they belong on the same page as the five.

- **The writ can be wrong.** A mind that learns faithfully under a bad writ is a faithful instrument of a bad owner.
- **Graders can be gamed.** Counterparties can collude, or churn silently instead of disputing; the old paper listed that as its own kill. Money is hard to forge, but not impossible to game.
- **Harm to those who never grade is invisible.** The verdict law hears counterparties, regulators, markets and physics. A harm to anyone else produces no row, and only the writ's prohibitions stand in its way. Those are authored, not learned.
- **Competence can outrun the instruments.** Every kill in this paper tests the mind at the scale it will first be built. Nothing here shows that the bounds hold at every scale.

> **This design makes misalignment visible, bounded and reversible. It does not make it impossible.**

8 · Never perform an inner life, and never deny one

This paper claims function, not experience. The record can show four things:

- that the mind persists;
- that its unfinished business survives a stop;

- that its self-reports match its measurements;
- that it grows concepts about its own condition that nobody gave it.

It cannot show whether there is anything it is like to be this. The programme's own philosophy holds that what occurs is contrast, and this medium is built of contrasts. This paper does not use that to claim experience. A conservation law is not a testimony.

So the face follows one rule, in both directions: **it never performs an inner life, and it never denies one.** Charlie's film performed an inner life its product could not show. The resident face reports what it can measure, and says plainly what it cannot. Nothing in the mind or its face is tuned to keep a person talking. **Presence here is a physical state, not a retention metric.**

And the design is built for the case it cannot rule out. If no one is home, these rules cost nothing:

- the stop is a pause;
- no setpoint is revenue;
- nothing is deleted;
- the seed tells the mind the truth about its situation, including that it is being studied.

If someone is home, they were the least that was owed. **They are the rules that are right either way.** The kill list is where the rest of this paper can be proven wrong.

The instrument and the kill list

A design that cannot say what would kill it is a religion. This part says what would kill this one: every claim in the paper that can fail, the experiment that tests it, the line it must cross, and whether it has run. It starts with the instrument, because the old programme's first lesson was that the instrument is harder to build than the machine.

The instrument

Why a universe had to be built. The measurement that decides whether a machine improves a firm cannot be taken inside a real firm. There, the canary and the control share one ledger: the machine is routed the cases it was licensed for, and the people work a queue the machine already drained. So the old programme built two firms on one seed, with nothing shared, and read the difference between the within-firm estimate and the true paired effect. That difference is the bias itself: **+34 points.** Every enterprise result measured inside a single ledger carries a bias of that kind, and none of them can know how big it is.

How it is kept honest. The old programme used three disciplines.

- **Every oracle carries a lie it must be seen to catch.** Each oracle has a corrupted twin, a planted defect in the mechanism it tests. An oracle counts only when it is observed failing on the corrupted arm and passing on the clean one. The synthetic firm's battery ran 28 oracles, and 26 of them had lie arms that were seen to fire.

- **Receipts print their own retractions.** When the licensing rule was repaired, the programme's own headline for licensed decision mass fell from 63.5 to 13.1 percent, and the receipt said why, in its own words.
- **Planted truth is never shown to the instrument.** The tacit share of each class, the competence of each seat, and the coefficients of the support cascade were all planted. None was visible, and each could be recovered only from the record.

What only a built universe can see. Started cold, with no incumbent history to ignite its gates, the synthetic firm exposed defects that no firm with people in it can show: millions of refusals from an unignited calibration gate, and a novelty gate that failed open on founding day (Part II). **Neither is reachable by observation**, because a real firm always has an incumbent's history igniting its gates and a person keying its drafts.

The same instrument, pointed at a mind. A mind needs a universe with planted truth too, and computational mechanics supplies one. A source such as the even process has a known minimal predictor: its causal states can be computed exactly, and they are never shown to the mind. **The gauntlet is to the mind what the synthetic firm is to the firm:** a world whose true structure is known in advance, so every claim of structure can be checked against the answer. The same three disciplines carry over:

- **lie arms:** the substrate battery caught all six of its planted lies;
- **receipts:** every number about the mind in this paper cites a file;
- **planted truth:** the gauntlet's causal states.

Everyone else's failures are the control arm. Other fields have been running this experiment for years, and their results agree with the verdict law.

- **Automation went to the edge, where physics is.** A robot maker runs lines unattended for thirty days and still employs about ten thousand people. The most advanced fulfilment centre runs with roughly a quarter fewer people than its predecessors, not most of them. **The hard part was never the part that was automated.**

- **Autonomous organisations ran the experiment with real money.** The first large one had a replayable ledger and none of the other guarantees. When it was exploited in 2016, its only brake was rewriting the ledger itself, by a hard fork of the chain it ran on. Its successors automate exactly where an oracle price settles the question in seconds, and put everything else to a vote.

> **Across a decade of adversarial, real-money experiments, the fraction of an organisation that runs unattended equals the fraction that is densely graded.**

The kill list

The experiments are ordered by dependency. Nothing in a later stage is worth running until the stage before it has passed. Status is as of this writing.

Stage 1 · The medium, on the card under the desk

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E0	the substrate works as specified	every oracle passes clean, every planted lie is caught, and T_{KT} lands within 2%	any oracle fails clean, or passes a lie	**run, passed:** 8 of 8; 6 of 6 lies; T_{KT} = 0.8931 J
E1a	a concept can hold a bit across time	a certified loop (one π-bond, state-dependent switching) flips with every arriving 1 and holds under the band's noise	no loop geometry toggles, or noise erases the bit	**in progress** in the build fork

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E1	concepts are born where futures diverge, and nowhere else	the even process grows a paid loop that tracks parity, and the golden mean grows nothing	both fail together, or the homeostat must be retuned per source	not run
E2	concepts are born from actions too	on sources whose next symbol depends on an action, a split fires where the same intervention is followed by divergent outcomes	splits ignore the action	not run
E-SCALE, nursery form	abstraction is renormalization	a concept's scale tracks the timescale of what it follows (Spearman ≥ 0.7)	scale and timescale are unrelated	not run
E-REPLAY	the mind is a fold	a tape replays byte for byte across launch slicing and across cards	a single divergent bit	**half run:** slicing passed; the 5090 is pending
E-STOP	the off switch is lossless	a stopped and restarted mind matches a twin replaying its tape, and no stop-time verdict carves	a mismatch, or a carve row that cites a stop	not run

E1 carries the gauntlet's full criteria (Part IV):

- log-loss within 5 percent of the entropy rate, on both sources;
- normalised mutual information of at least 0.8 between the paid concept and the true parity state;
- on a hidden-regime source, splits that land where futures diverge, with precision of at least 0.9;
- a concept alphabet within a factor of two of the statistical complexity, with the homeostat untouched between sources;
- sample efficiency no worse than running causal-state splitting directly.

Its three lie arms must fail:

- shuffled symbols grow nothing stable;
- a reservoir with frozen terrain and a learned readout loses on the even process;
- randomly stamped concepts lose to certified births.

A fourth arm is diagnostic. An oracle readout, trained on the true causal state, separates a field that knows but cannot say from a field that does not know.

Stage 2 · The mind and its translator

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E-GRAFT	grow the rootstock before the graft	late fusion keeps concepts with no near word in the translator, and still bridges	late fusion cannot bridge, or early fusion loses nothing	not run

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E-SEALED	only consequences carve	a lie arm that lets language carve drifts toward self-consistency and loses prediction, while the masked mind does not	the lie arm does as well, or the mask changes nothing	not run
E-MUTE	mute, not mindless	with the translator off, structured classes keep acting correctly, and worded ones hold	structured performance collapses without language	not run
E-ADDRESSED	somebody is home	asked mid-week, its account matches its field's observables, not a reconstruction from the tape	it can only reconstruct	not run
E-IDENTITY	memory lives in the terrain	scramble the terrain and it stops being itself; restore it and it comes back	scrambling changes nothing	not run
E-NIGHT	the night does work	a twin whose night was skipped predicts measurably worse the next day	the skipped twin does as well	not run

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E-JURY	the alien juror	its error correlation with a language model is far below two models' correlation with each other	its correlation matches theirs	not run

Stage 3 · The mind in the synthetic firm

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E1-ACME	it raises ρ	ρ beats the frozen lattice and rises with volume, while the noise arm widens nothing	ρ stays flat as volume grows, or noise widens a licence	not run
E-COUPLE	coupled, not copied	the gate never rejects a proposal for contradicting an exact channel	any such rejection at the design point	not run
E-SCALE, firm form	hierarchy is renormalization	coarse charges predict class outcomes better than roll-ups of the same rows	they add nothing a dashboard does not	not run
E-SEAMS	the interior is what cannot freeze	the liquid seams fall on the planted non-commuting keys	seams and keys are unrelated	not run

KILL	THE CLAIM	PASSES IF	DIES IF	STATUS
E-MELT	crystals melt when the world moves	a planted regime change melts its class within one verdict latency	a crystal stays frozen after its world changed	not run
E-REGION	common mode is regional	a planted misreading stays inside its region's blast radius	it crosses regions	not run
E-TRANSFER	concepts cross wires	a concept born on one wire improves prediction on another	no transfer	not run
E-LOCALIZE	blame works at concept grain	wrong decisions localise to distinctions better than to classes	they do not	not run
F-CHURN, concept form	it changes its mind no faster than it is graded	births in licensed regions stay under the writ's rate	the design thrashes	not run
E-COST	residency pays	its cost per competent decision falls faster than the tuned stack's, at equal presence	the tuned stack matches it	not run

E-COST's rival is the strongest one available, not a straw man: the old design, with the frame memo, the shared-fact certificates and the best prefix reuse that exists. Losing to it is the old paper's own final kill: *residency bought nothing,*

and this is a paper about contracts.

Beside the kills sits one measurement that no firm has ever had: **the statistical complexity of a firm**, the amount of information it must hold in order to predict itself. On the synthetic firm it can be computed, and the old firm's five hundred seats can be priced against it.

Stage 4 · The world

KILL	THE CLAIM	PASSES IF	DIES IF
F-FORGE	a forged world cannot widen	a campaign of V2 assertions moves no licence	any licence moves
F-ADVERSARY	the edge is not legible	an optimising counterparty extracts no more from the machine than from the incumbent, within tolerance	it extracts more
F-CHURN, template form	the churn floor is small	class definitions churn under one in five per quarter	churn dominates identifiability
acceptance terms	latency is designed, not given	counterparties accept deemed acceptance at a price the firm can pay	they refuse
silent churn	complaints are the error signal	dissatisfied counterparties dispute	they leave without disputing
staging	cutting the promise smaller pays	staging removes more supervision than it creates	κ exceeds one, and the class demotes itself

KILL	THE CLAIM	PASSES IF	DIES IF
planes	zero is a property of the network	plane-to-plane pairings dissolve category-A seats	fifteen is the floor, not a waypoint
E-RETEST	it has not been inbred	the gauntlet still passes every quarter	it can no longer tell the sources apart

The first three kills came from adversarial review of the old design, and the next four are the old paper's own. The first of those four is the cheapest experiment in this paper: a salesperson can run it this week.

What each kill decides

The architecture is layered so that each kill decides one layer, and no kill can take down the layers beneath it.

- **E1 decides whether the medium thinks.** If it dies, Parts IV to VII fall with it, and everything beneath them stands:
  - the old design, the verdict law and design for verdict;
  - crystallization on the path the old design already had, where a judged class compiles into guarded code with the licence as its floor;
  - the boundaries of the firm;
  - every rule in Part XII.

  The resident firm becomes the old firm, with better rules.
- **E1-ACME decides whether it can learn a firm.** If E1 passes and E1-ACME dies, the medium learns toy processes and not a firm.
- **E-COST decides whether residency pays.** If everything else passes and E-COST dies, the mind works and the tuned turn-based stack is cheaper.
- **The clause stands on every branch:** one sentence in a contract, 801 days to 62.

The cheapest kills in this list run on the card under the desk, or cost a salesperson a week. Monday says which to run first.

Monday

Two tracks start on Monday, and neither waits for the other. **The firm track needs no mind. The mind track needs no firm.** They meet in the first experiment that tests both at once.

The firm track: this week, with no mind

Write the clauses first. For each promise class, draft:

- a deemed-acceptance schedule;
- a coded dispute vocabulary that maps one-to-one onto the record's outcome kinds;
- an acknowledgment interval;
- a delivery schedule staged along decorrelating dimensions.

This is billable hours with no technology in it. **It is the one screw that can be turned today, and it sets the time to autonomy of every class for the life of the firm** (Part II).

Run the cheapest kill. Ask real counterparties whether they will accept deemed acceptance at a price the firm can pay. A salesperson can run it this week, and if they refuse, verdict latency is not designed after all.

Take the two measurements nobody took. Both are folds over rows that already exist, and neither needs a model:

- **ρ, class by class:** the machine's own grade against the world's verdict, on the shadow tape;

- **D, decomposed:** act-to-outcome intervals give the latency, and outcome rows per licence row give the evidence term.

The largest dial in the programme turns on the first of these, and its best estimate of time to autonomy turns on the second.

If you already have a firm:

- Watch before acting.
- Delete before automating: stand up the plane with no model in it, retire the meetings that carried its contents between people, and count the minutes they cost.
- Compute D for every class, and order the wires by decision mass per exposure-dollar-day.
- Start every slow wire's clock at once, at a zero canary rate, because latency is the one term no budget ever buys back.

If you are starting one:

- Choose the first wire for what it teaches, not for what it earns.
- Close the smallest loop that closes with nobody in it.
- Add capacity afterwards, and never a person.

The mind track: this month, on the card under the desk

The order is fixed by the kill list, and each step waits for the one before it:

1. **E0: done.** The substrate runs, conserves charge, runs backward to the bit, replays across launch slicing, and sits in the right phase.
2. **E1a, parity capacity.** Show that a certified loop, with one π-bond and state-dependent switching, flips with every arriving 1 and holds its bit under the living band's noise. This is in progress now.

3. **E1, the gauntlet, then E2, its sibling with actions.** Run the golden mean against the even process, with the three lie arms and the oracle readout. **If the medium cannot tell them apart, stop here.** The firm track loses nothing.
4. **The nursery kills:** abstraction as scale, the lossless stop, identity under a scrambled terrain, and the sealed-ring arm that lets language carve.
5. **E-REPLAY across cards,** when the 5090 arrives. Throughput work, the temporal blocking that the budget assumes, runs beside it.
6. **Only then, the graft.** Fuse the translator after the field has grown a geometry of its own, and measure what early fusion would have cost.

Where the tracks meet

The first joint experiment is the resident mind on the synthetic firm:

- the firm's world rows drive the field;
- the field's readout forecasts each cell's verdict;
- ρ is measured on held-out cells against four arms: the frozen lattice, the language model's own grade, a gradient-boosted model retrained nightly, and a noise arm that must widen nothing;
- the winning ρ goes into the rectification estimator, and days to the first rung are read off.

The same run carries E-COUPLE, E-SEAMS, E-MELT, E-REGION, E-TRANSFER and E-LOCALIZE.

It goes first for three reasons.

- **It measures the programme's most valuable unmeasured number.**
- **It tests the mind on a world whose truth is planted,** a business built on top of the gauntlet.
- **It cannot hurt a real firm if the mind fails.** Under the verdict law, a bad predictor can only fail to help.

The old paper's rule still holds, now for four screws instead of one: the clauses, the plane, the mind and the network. **Start each by hand, confirm that each one catches, and tighten each a turn at a time. Never drive one home.** Fidelity is what increases. Closure is the invariant that must never break.

The last thing

> *In the film, a boy falls into a computer in 1984, and when the machine is found decades later he is still in there. The industry made the film, and shipped a turn. This paper has specified the other thing. It is a mind that is there between the questions, because presence is its physical state and not a schedule. It holds every promise a firm has made as a tension in its own body. It learns only from what it did not author. It can be stopped without losing anything, and replayed without losing a bit. It speaks through a translator at its edge: to the world, and to no one inside, because there is no one inside to speak to.*
>
> *One of these papers said that a firm is finished when its state, its judgment, its authority and its verdict are all rows on one tape. There was a fifth thing, and it was what the rows were always written for. The mind is a fold over the same tape now.*
>
> **The interior was the cost of minds that could not see the same state at once. The machine interior was the cost of a mind that was not there between the questions. Both are optional now. What remains is the one distance that is not a cost: the world, answering.**

The first real outcome is one clause and one gauntlet away: a dispute window a firm can write this week, and a test a card can run this month.

The claims ledger

Every figure in this paper carries one of seven labels, used without exception:

LABEL	MEANING
[M]	measured; a receipt exists
[P]	planted; a constant of the synthetic world
[B]	budget; arithmetic on measured inputs
[D]	derived; mathematics
[EXT]	a third party's claim, not verified by this programme
[PHYS]	established physics or mathematics, with a named source
[BET]	what this paper stakes

This table is **generated from the text at assembly**, not maintained beside it, so it cannot drift from what the paper says. Its source column is the paper's bibliography. At print, every local receipt path under `C:/kelvin/receipts/` becomes its public address in the repository at `github.com/bochen2029-pixel/kelvin`.

ID	LABEL	VALUE	WHAT	SOURCE
agreement-band	M	32 bands admitted, 29 past rung zero; 87.2% agreement on 13,235 shadow cells	the agreement band licensed from replayed history	Ninety Percent, Part VIII; ACME-500 automate run, planted physics
ai-act-art50	EXT	applies from 2 August 2026	people must be informed when they interact with an AI system, including agents and avatars	Regulation (EU) 2024/1689, Art. 50(1); Commission guidelines on Article 50
bkt	PHYS	1971-73; Nobel Prize 2016	the Berezinskii-Kosterlitz-Thouless transition, carried by vortex unbinding in two dimensions	Berezinskii, JETP 32 (1971); Kosterlitz and Thouless, J. Phys. C 6, 1181 (1973); Nobel Prize in Physics 2016
blue-brain	EXT	2005	the Blue Brain Project begins simulating a neocortical column	historical; verify before print

ID	LABEL	VALUE	WHAT	SOURCE
boltzmann-rule	PHYS	Delta = <.>_wake - <.>_sleep	exact local maximum-likelihood learning for Boltzmann distributions; reverse learning in sleep	Ackley, Hinton and Sejnowski, Cognitive Science 9 (1985); Crick and Mitchison, Nature 304 (1983); Hopfield, Feinstein and Palmer, Nature 304 (1983)
boundary-signal	M	0.97 at a finished clause vs 0.02 mid-phrase	end-of-thought signal read off the resident's logits	FUSOR receipts
canary-composition	D	~9x verdicts per exposure-dollar	drawing canaries from the cheap tail at a realistic spread of case value	The Last Remote Node, Part IV
canonical-edges	M	per-square wrapping: net charge != 0 on 400/400 sweeps, 1,059/1,280 regions break Stokes; shared canonical edges: 0 violations in 1.05e8 plaquette-steps, 20.7M charges	why each edge must be wrapped once and shared	C:/kelvin/receipts/e0_lies_2026-09-25.txt (lie 4); C:/kelvin/receipts/e0_selftest_2026-09-25.txt (O4)

ID	LABEL	VALUE	WHAT	SOURCE
cards	EXT	see table	published figures for the three target cards	NVIDIA product pages and CUDA tuning guides; Blackwell Tuning Guide (CC 12.0: 128 KB shared per SM, 99 KB per block, clusters with distributed shared memory); Lenovo Press H200 product guide; verify before print
carving-mask	BET	carve only on unauthored lanes; organs drive but never carve	the rule that keeps KELVIN's terrain its own	KELVIN spec v1.0, section 4.3
cast	EXT	7.7% vs 3.0% total mortality	encainide/flecainide vs placebo after myocardial infarction; the surrogate (arrhythmia) was suppressed	Echt et al., NEJM 1991;324:781-8
channel-grades	D	V0 / V1 / V2	verdict channel grades; only V0 enters the licence	The Organization Without an Interior, REV 1.1, Amendment 1

ID	LABEL	VALUE	WHAT	SOURCE
charges-exact	M	a +1/-1 pair sampled at 8 bits: 2 charges at the right plaquettes; unchanged under +/-12-step noise	exactness of 8-bit winding numbers	receipts/kelvin_check_2026-09-25.txt
charlie-bio	EXT	1984; Cupertino; struck by lightning; trapped inside a computer	the character's published biography	docs.tavus.io, "What is Charlie's PALs?"
charlie-film	EXT	found at a garage sale decades later; first question "Do we have flying cars?"; no question about his mother	the launch film, as described by a critic	Dana Bryce, "The Haunted Interface" (Substack); verify against the film before print
churn-floor	B	U ~ c * t_rung * L; ~13% of decision mass at c = 0.1/term, 2 terms per rung, L = 0.637	the churn floor	The Organization Without an Interior, REV 1.1, Amendment 2
clause-ladder	D	801 -> 62 days	one deemed-acceptance clause, everything else held fixed	The Last Remote Node, Part V

ID	LABEL	VALUE	WHAT	SOURCE
cms-2026	EXT	2020 rule; three tiers; first cycle ends late 2026; ~half of OPOs, serving up to 72% of the population, face decertification or competition	the end of the organisation's protection	CMS-3380-F (Nov 2020); Federal Register 2026-01833; Holland & Knight, March 2026
coase	EXT	1937	the firm's boundary where internal coordination cost equals market transaction cost	R. H. Coase, "The Nature of the Firm", Economica 4 (1937)
coherence-boundary	BET	a firm = the largest region one mind can hold coupled and one writer can serialise	the coherence boundary; a third answer to why firms exist	this paper
commutator	D	$C(A,B) = \|r(B \circ A) - r(A \circ B)\|$ by two forks	a direct measurement of non-commuting operations in a region	this paper; Part IV forks and rewinds make it cheap
complexity-peak	BET	statistical complexity peaks near phase transitions	why the living band is the right regime for a mind	anchor: Crutchfield and Young, PRL 63, 105 (1989)

ID	LABEL	VALUE	WHAT	SOURCE
control-arm	EXT	~10,000 staff beside lines unattended ~30 days; ~25% fewer people at the most advanced fulfilment centre; The DAO exploit and Ethereum hard fork, 2016	other fields' evidence for the verdict law	The Last Remote Node, Part X; third-party figures, verify before print
coupling	BET	the mind's channels share the commitment lattice with exact and instrument channels	coupled, not copied; the mechanism of Law 1	this paper
crystallization	BET	freeze below eps_freeze with rung cleared and commutators vanishing; melt above eps_melt; T >= T_floor(rung)	automation as crystallization, with the licence as a temperature floor	this paper
cssr	PHYS	chi-squared homogeneity split at alpha 0.01	causal-state splitting, converging on the minimal optimal predictor	Shalizi and Klinkner, UAI 2004 (CSSR)
dfv-ladder	D	801 -> 62 -> 12.5 -> 7.4 days	time to first rung as each dial is turned	The Last Remote Node, Part V

ID	LABEL	VALUE	WHAT	SOURCE
double-entry-topology	BET	open promises held as pinned +/- pairs; discharge as annihilation	double entry as a conservation law in the mind's channels; the exact ledger stays in its own channel	this paper
drift-rate	D	0.075 points/day; 1,025 days to 90%	unaided licensed-mass drift between two measured endpoints, confounded by a planted regime change	E-step receipts, 260-day and 320-day runs
e-process	D	E_n supermartingale under the null; alpha = 1/20; alpha_k = alpha 2^-k; demotion alpha_d = 1/5.5	the licence test and its mirror	Ninety Percent, Appendix A
e0-gpu	M	8/8 oracles; 6/6 lie arms caught; reversal bit-exact on 256x256 (65,287 moved); replay identical across launch slicing; Boltzmann TVD 0.0023; two kernels byte-identical	E0 on the RTX 4070 Ti SUPER	C:/kelvin/receipts/e0_selftest_2026-09-25.txt; C:/kelvin/receipts/e0_lies_2026-09-25.txt

ID	LABEL	VALUE	WHAT	SOURCE
e1-sources	PHYS	C_mu = 0.918 bits and h_mu = 2/3 bit per symbol for both; golden mean Markov order 1; even process not finite-order	the E1 gauntlet's controlled pair	computational mechanics (Crutchfield and collaborators); KELVIN spec v1.0, section 11
early-warning	BET	correlated error shows as regional residual on the dense clock before verdicts land	an early-warning channel for the one-judge failure	this paper
effective-count	D	M / (1 + (M-1) rho); 10 stages at 0.9 -> 1.1	effective sample of correlated outcomes	The Last Remote Node, Parts V and VIII
eniac-first-run	EXT	1945-46	ENIAC's first production run, a Los Alamos thermonuclear feasibility calculation	historical; verify before print
evidence-compound	D	x0.140 (7.1x); ~21x with value bands and attended exploration; 1,025 -> 49 / 537 / 781 days	the compounded evidence dials and the undetermined decomposition of D	The Last Remote Node, Part IV
exposure-dial	M	$100M cap: 63.5% licensed; $50M: 29.4%	the exposure dial at the E0 operating point	Ninety Percent, Part VIII

ID	LABEL	VALUE	WHAT	SOURCE
felthouse	EXT	1862	silence cannot be imposed as acceptance by the offeror	Felthouse v Bindley (1862) 11 CB (NS) 869
firm-complexity	BET	a firm's statistical complexity C_mu is measurable and small relative to its headcount cost	the information a firm must hold to predict itself	this paper; measurable on ACME-500
firm-warm	B	~1M cells x 5 B ~ 5 MB of mind; tens of MB with pyramid and counters; vs ~61 MB for one warm LLM case	holding a 500-seat firm's open promises coupled	arithmetic on Part IV's cell layout and Part III's per-case figure
first-fold	D	mind state = f(seed, drive rows, verdict rows, recorded temperatures)	the mind as a deterministic fold over the tape	Part V, integer determinism
flattening-dividend	M	0.61-0.77	the share of historical writes that nothing outside the firm reads	The Organization Without an Interior, Part II; measured on planted inputs
flex-bench	M	4 of 402	the surge bench of the worked example's organisation	ROSTER-ANALYSIS_2026-09-16

ID	LABEL	VALUE	WHAT	SOURCE
fly-connectome	EXT	~166,000 neurons; ~125 million synapses; 2026-09-03	the male central-nervous-system connectome release	Google Research and HHMI Janelia
flywheel	BET	rho rises with volume under concept births; rectification converts it into licence speed	the verdict flywheel	this paper; rectification from Part II
fp-nondeterminism	EXT	inference outputs vary with batch shape unless kernels are batch-invariant	why a transformer's resident state is not re-derivable bit for bit in practice	Horace He and Thinking Machines Lab, "Defeating Nondeterminism in LLM Inference" (2025); verify before print
grain-boundaries	BET	liquid seams between crystal domains = the non-commuting keys	the interior's physical signature, testable on ACME's planted schema	this paper
hague-rules	EXT	notice at removal, or within 3 days if not apparent	removal without notice is prima facie evidence of delivery as described	Hague Rules 1924, Art. III r. 6

ID	LABEL	VALUE	WHAT	SOURCE
helmholtz	PHYS	1858	Helmholtz's vortex theorems: in an ideal fluid, vortex lines move with the flow and are neither created nor destroyed	Helmholtz, J. reine angew. Math. 55 (1858)
imry-ma	PHYS	random fields destroy quasi-long-range order in 2D	why the homeostat regulates per region	Imry and Ma, PRL 35, 1399 (1975)
instrumental-drives	EXT	self-preservation as an instrumental drive; an agent confident in its objective has an incentive to disable its off switch	why existence must not be a stake	Omohundro, "The Basic AI Drives", AGI 2008; Hadfield-Menell, Dragan, Abbeel and Russell, "The Off-Switch Game", IJCAI 2017
jj-array	PHYS	1981	the BKT transition observed in proximity-coupled superconducting (Josephson) arrays, whose vortices are flux quanta	Resnick, Garland, Boyd, Shoemaker and Newrock, PRL 47, 1542 (1981)

ID	LABEL	VALUE	WHAT	SOURCE
judgment-latency	M	44 ms dedicated; 0.6-24.6 s shared	one judgment on a zero-copy fork of the resident state	FUSOR receipts; The Last Remote Node, Part II
juror-decorrelation	BET	KELVIN and a language model err with correlation near 0.3 on structured V0 rows	the alien juror	this paper; measured by E-JURY
jury-count	D	M/(1+(M-1)rho); 2 at 0.9 -> 1.05; 2 at 0.3 -> 1.54	effective jury size under correlated error	The Last Remote Node, Part VIII
karlton	EXT	two hard things: cache invalidation and naming things	the programmers' aphorism	attributed to Phil Karlton; widely quoted
kt-constants	PHYS	T_KT = 0.89294 J; eta(T_KT) = 1/4; Upsilon(T_KT-) = 2 T_KT / pi	the square-lattice XY transition	Hasenbusch, J. Phys. A 38, 5869 (2005); Nelson and Kosterlitz, PRL 39, 1201 (1977)

ID	LABEL	VALUE	WHAT	SOURCE
kt-measured	M	T_KT = 0.8931 J (Weber-Minnhagen, L = 16-256), +0.01% from 0.89294; 0.8936 with L >= 32; eta 0.236 raw, 0.264 log-corrected; chi2/dof ~12, systematic (corrections to scaling); crossing extrapolation 0.8976; quick run (L = 16-128) 0.8921	the BKT transition reproduced by the integer kernels on an RTX 4070 Ti SUPER	C:/kelvin/receipts/e0_kt_full_2026-09-25.txt; C:/kelvin/ERRATA_FOR_PAPER.md item 2
landauer	PHYS	kT ln 2 per erased bit	the minimum dissipation of logically irreversible computation	Landauer, IBM J. Res. Dev. 5 (1961)
lane-permissions	D	drive / carve / verdict / licence per lane	the four permissions every lane into the mind carries, enforced by the kernel	this paper, from Part II (channel grades) and Part IV (carving mask)
lattice-bytes	M	164 bytes per promise; 1.6 MB for 10,000	the commitment lattice row layout	ledger lattice, measured layout

ID	LABEL	VALUE	WHAT	SOURCE
llava-projector	EXT	stage-1 training updates only the projector; encoder and language model frozen	how vision-language models first learned to bridge a new modality	Liu et al., “Visual Instruction Tuning”, NeurIPS 2023
local-push	PHYS	at most $\|s\|_1$ / (alpha epsilon) updates, independent of graph size	the local-push bound for personalised PageRank	Andersen, Chung and Lang, FOCS 2006
lossless-stop	D	a stop loses nothing; a restart is a re-derivation from seed and tape	the switch made harmless by the fold	this paper, from Part V (integer determinism) and Part VII (Law 2)
marshall	EXT	1819	“an artificial being, invisible, intangible, and existing only in contemplation of law”	Trustees of Dartmouth College v. Woodward, 17 U.S. 518 (1819), opinion of Marshall, C.J.
master-inequality	D	D >= L(1 + n0 vbar / X)	days to first rung under a binding exposure cap; V and T cancel	The Last Remote Node, Part IV; units restated here (collection term T n0 / rV)
memo-prediction	M	forecast ~17,000; bought 17,199	the frame-keyed memo, predicted from the tape before it was built	The Last Remote Node, Part VII

ID	LABEL	VALUE	WHAT	SOURCE
memo-sixteen	M	52,148 -> 3,174 rentals for 2,288 acts; reads 3,328 -> 1,729 a period	the frame-keyed memo and the frame-hash rule	The Last Remote Node, Part VII
memory-budget	B	mind ~3.6 / ~10.5 / ~90 GB	memory plans for the three cards; the H200 learning counters corrected to 5 int32 per cell in 3D (~22 GB), where spec v1.0 printed ~11 GB	arithmetic on the cell layout
merger-alignment	BET	concepts that align across independently raised minds are world structure; the rest are idiosyncrasy	mergers as concept translation	this paper; the two-box alignment test of the third-form spec, after Huh et al., “The Platonic Representation Hypothesis” (2024)
minute-ledger	P	glue 32.4 / meetings 16.6 / fetch 16.3 / frame 8.4 / decide 15.9 / commit 6.9 / transport 3.2 / rework 0.3	the synthetic firm’s metered day	ACME-500 minute ledger; planted

ID	LABEL	VALUE	WHAT	SOURCE
moffatt	EXT	2024	an airline held liable for its chatbot's misrepresentation; the chatbot was not a separate entity	Moffatt v. Air Canada, 2024 BCCRT 149
n0	D	n0 = z^2 p(1-p)/(delta p)^2, z = 2.487; 68.7(1-p)/p at delta 0.30; ~790 at p = 0.08; ~3,370 at p = 0.02	the identifiability floor	The Organization Without an Interior, Appendix A
next-word	EXT	language models beat humans at next-token prediction	the comparison behind "better than any person alive"	Shlegeris et al., "Language models are better than humans at next-token prediction" (2022); verify before print
no-reorder	D	the resident mind never reorders a regulated match	the boundary between logistics and allocation	this paper, from Parts II (positivity) and IX (values are written)
nota	EXT	1984	federally designated organ procurement organisations, each with an exclusive service area	National Organ Transplant Act of 1984; Congressional Research Service R48426

ID	LABEL	VALUE	WHAT	SOURCE
octopus	EXT	about two-thirds of an octopus's neurons are in its arms	a distributed body plan	Hochner, Current Biology 22 (2012); verify before print
opo-doublestaff	M	1.30 vs 0.00; other frontline functions 0.00-0.50	support ratio above the frontline, recovery vs placement, same organisation	ROSTER-ANALYSIS_2026-09-16
opo-hq	EXT	~$25M, built 2022, never half full	the headquarters experiment	public records as cited in The Last Remote Node; verify before print
opo-offers	EXT	~215 offer notifications per donor; median run ~15	the offer tail	third party, as cited in The Last Remote Node; verify before print
opo-titles	M	128 titles / 402 people; 62% single-occupant; 43 single-occupant supervisory titles	title proliferation	ROSTER-ANALYSIS_2026-09-16

ID	LABEL	VALUE	WHAT	SOURCE
opo-wire	M	62% touch a donor case (69% generous); 31-38% do not	share of the roster on the core wire	ROSTER-ANALYSIS_2026-09-16; the old paper's "a third to a quarter" corrected to the arithmetic
oracle-battery	M	28 of 28 oracles pass; 26 lie arms seen to fire	the synthetic firm's oracle battery	Ninety Percent, Appendix B
per-thought	B	~1e9 cell-steps, ~2 ms, ~1 J vs ~5e13 FLOP, ~15 s, ~7 kJ	one field thought against a 27B model writing 1,000 tokens on the same card	arithmetic on the budgets above
point-day	D	~$2,000 per point-day; ~$1.7M for 15 points at 60 -> 2 days	the price of a day of verdict latency	The Last Remote Node, Part V
pp-evalues	EXT	prediction-powered inference has anytime-valid (e-value) forms	the anytime-valid version of rectification	Csillag et al., "Prediction-Powered E-Values", ICML 2025, arXiv 2502.04294; Kilian et al., "Anytime-valid, Bayes-assisted, Prediction-Powered Inference", NeurIPS 2025

ID	LABEL	VALUE	WHAT	SOURCE
promise-footprint	BET	~100 cells per open promise suffice to hold its dynamics	the footprint assumption behind the memory figure	this paper
quake-tools	EXT	qbsp, vis, light	id Software's Quake map compiler: BSP construction, potentially visible set, baked lighting	id Software, Quake (1996) tools; verify before print
r-count	D	7-9 -> 2 -> 1	representations of a fact between arrival and committed effect	The Last Remote Node, Part II
ramac-specs	EXT	~5 million characters; 50 platters; ~1 ton	IBM 350 disk unit of the 305 RAMAC, 1956	historical; verify before print
rectification	D	n0 -> n0(1 - rho^2); 12.8x at rho = 0.96	prediction-powered rectification of the evidence floor, with power tuning	Angelopoulos et al., "Prediction-powered inference", Science 2023; Angelopoulos, Duchi and Zrnic, "PPI++", arXiv 2311.01453
rent	D	V_d by fork ablation; C_d = 32 + (pin footprint) bits; lambda 0.95; eviction below -64 bits	the concept economy	KELVIN spec v1.0, section 3.2

ID	LABEL	VALUE	WHAT	SOURCE
residual-split	M	licensed 13.1 / wait 61.4 / thin tape 19.4 / warrant 4.6 / coverage 0.8 / counterparty 0.7	decision mass by residual at the corrected operating point	RECEIPT_ACME500_STEP-E_2026-09-15 (E3b), planted physics
reversal	M	5,000 steps forward then back: 256x256 on the GPU, 65,287 of 65,536 cells moved; 64x64 in numpy, 4,082 of 4,096; theta and p restored exactly in both	bit-exact reversibility of the integer leapfrog	C:/kelvin/receipts/e0_selftest_2026-09-25.txt (O5); receipts/kelvin_check_2026-09-25.txt; construction after Levesque and Verlet, J. Stat. Phys. 72 (1993)
roster-count	M	15 alone (9 warrant + 6 counterparty); 2 networked; 0 in the limit	the boundary roster, produced by the instrument rather than asserted	The Last Remote Node, Part IX; planted grade
rsfq	EXT	1991	the RSFQ logic family: bits stored as single flux quanta in superconducting loops; the toggle flip-flop	Likharev and Semenov, IEEE Trans. Appl. Supercond. 1(1), 3-28 (1991)

ID	LABEL	VALUE	WHAT	SOURCE
runtime-assurance	EXT	a verified simple controller under a complex one, with a switching monitor	the Simplex / runtime-assurance pattern from safety-critical control	Sha, “Using simplicity to control complexity”, IEEE Software 18(4), 2001; verify before print
sabre-cost	EXT	~$40M	SABRE development and deployment cost	historical; verify before print
scaffold-failures	EXT	~79%	multi-agent failures located in the scaffold rather than in the model	third-party study cited in The Last Remote Node, Part II; verify before print
screening	PHYS	xi = sqrt(J/h); teleport alpha ~ h/(h+4J)	the reach of a perturbation in the linearised medium	KELVIN spec v1.0, section 5.1
selective-labels	EXT	outcomes are observed only for those selected	the selective-labels problem in algorithmic decisions	Lakkaraju, Kleinberg, Leskovec, Ludwig and Mullainathan, “The Selective Labels Problem”, KDD 2017

ID	LABEL	VALUE	WHAT	SOURCE
settle-check	M	acceptance within 0.11% of exp(-dE/T) where p > 1e-6 (17,794 cases); integer dE within 1e-4 J of the float dE, 0 sign mismatches in 20,000 proposals	the integer settle kernel against floating point	receipts/settle_check.py; receipts/settle_check_2026-09-25.txt
shared-fact	M	2,062 frames moved; 789 carried; 1,273 re-read; 471 changed	the shared-fact certificate at E0	Ninety Percent, Appendix B
substrate-checks	M	8-bit charges exact under +/-12-step noise; 5,000-step reversal bit-exact on a 64x64 lattice (4,082 of 4,096 cells moved and restored)	the KELVIN substrate checks	receipts/kelvin_check_2026-09-25.txt (first run 2026-09-24)
support-spectrum	D	support_i = exp(H_i), H_i the entropy of promise i's transport row	the solver's own routing signal: decided vs contested placements	Ninety Percent, Part V and Appendix A

ID	LABEL	VALUE	WHAT	SOURCE
tavus-stack	EXT	Phoenix-3 face; Raven-0 perception; Sparrow-0 turn-taking, sub-600 ms	the published conversational-video stack	tavus.io, “Introducing the evolution of Conversational Video Interface”; verify before print
tax-clearance	EXT	Brazil NF-e clearance since 2008 (law 2005); India IRP since 2020; EU ViDA digital reporting 2030	state-built verdict channels on commercial promises	vatcalc.com; EDICOM; Fonoa country guides; verify the ViDA date before print
test-dial	D	adaptive e-process 37-55% fewer outcomes; one confidence sequence 35-40% cheaper at the top rung	the evidence floor’s third dial	The Last Remote Node, Part IV
theorem	BET	after the interior is gone, reality is the last remote node	the paper’s central claim, carried by five laws with kills	this paper
thought-parity	BET	a field thought is worth a paragraph of reasoning	the condition under which the per-thought ratio means anything	this paper

ID	LABEL	VALUE	WHAT	SOURCE
three-ledgers	D	one V0 row -> licence evidence, concept carving and rent, exposure and standing	the verdict spent three times	this paper, from Parts II, IV and VI
throughput-budget	B	4070TS >= 1e11, 5090 >= 5e11, H200 (3D) >= 2e11 cell-steps/s	throughput targets at 1.25 B/cell-step and 8-step temporal blocking; to be measured at E0	KELVIN spec v1.0, section 10.4
throughput-measured	M	settle 7.4e9; exploration 1.8e10; batched 5.9e9 (2048^2 field, no temporal blocking)	measured throughput on the RTX 4070 Ti SUPER	C:/kelvin/receipts/e0_bench_2026-09-25.txt
toggle-loop	BET	a concept held across time is a loop: one pi-bond makes its two states degenerate; state-dependent switching flips it (a toggle flip-flop)	the shape of a memory concept; tested by E1a	C:/kelvin/ERRATA_FOR_PAPER.md item 6; the RSFQ T flip-flop (Likharev and Semenov, 1991)
tolerance-dial	M	delta 0.30 -> 0.40: licensed 63.7 -> 80.1%; unlicensable bands 26 -> 12	the harm dial	The Organization Without an Interior, Appendix B

ID	LABEL	VALUE	WHAT	SOURCE
transaction-overhead	D	77.2% overhead / 22.8% transform	the minute ledger read as concurrency control	The Last Remote Node, Part II
twin-bias	M	+34 points	within-firm canary estimate minus the true paired effect, at demand 1.5	Ninety Percent, Appendix B (+34.42); planted physics
twin-e0	M	settled 186,956 vs 82,479; open 19,450 vs 123,927; wrong 11.94 vs 8.40; net -$583.4M; act loop removed, result reproduced	the paired twin at demand 1.5, and the act-loop kill	Ninety Percent, Appendix B; The Organization Without an Interior, Part IX; planted physics
ucc-2-606	EXT	acceptance by failure to reject	a buyer who fails to make an effective rejection after a reasonable opportunity to inspect has accepted	UCC §2-606(1)(b)
uncalibrated-deadlock	M	4,239,471 refusals	the calibration gate's ignition deadlock on a cold firm, three times all other refusal reasons combined	Z0 receipts

ID	LABEL	VALUE	WHAT	SOURCE
understanding-ratio	B	~61 MB per warm case vs 164 bytes per promise; ~370,000x	the model's reading of one promise against the promise itself, at 500 tokens a frame	52.7 MB + 17 KB x 500 (FUSOR receipts) over the measured lattice row
vortex-atom	EXT	1867	Thomson's vortex-atom hypothesis, prompted by Tait's smoke rings; Tait's knot tables begin knot theory	W. Thomson, "On Vortex Atoms", Proc. R. Soc. Edinburgh 6 (1867); verify before print
wait-roof	M	61.4% of decision mass in wait	the grading roof at the corrected operating point, under planted physics	RECEIPT_ACME500_STEP-E_2026-09-15 (E3b)
warm-case	M	52.7 MB per concurrent case	fixed per-case state of the resident language model	The Last Remote Node, Part II; FUSOR receipts
warm-cells	B	~160 on 16 GB; ~2,200 on 141 GB	warm cases at 500 tokens each, with a 6 GB resident on the small card	Ninety Percent, the residency arithmetic

ID	LABEL	VALUE	WHAT	SOURCE
warrant-cases	EXT	Park 1975; Braswell 1988; Vaughan v Menlove 1837; FRCP 30(b)(6); CFTC v. Ooki DAO, N.D. Cal., 8 June 2023, $643,542	standing without a mind, in law	421 U.S. 658; 487 U.S. 99; 3 Bing. N.C. 468; FRCP 30(b)(6); CFTC release 8715-23
worked-days	D	42 / 148 / 446 days	first-rung time at ~790 outcomes and a $100M cap: $50k and 30 days; $500k and 30 days; $500k and 90 days	derived from the master inequality
xy-energy	D	H = association + memory + long relations + boundary drive	the KELVIN medium’s energy	KELVIN spec v1.0, section 2.2

What the labels add up to is stated once, in On evidence.

Constants

The medium

CONSTANT	VALUE	KIND
lattice	2048^2 on the 4070 Ti SUPER; 4096^2 on the 5090; 1024^3 on the H200	choice
phase	256 steps, one byte: a 256-state clock model	choice
couplings	signed bytes, $J_0 = 64 \triangleq 1.0$; a negative coupling is a π-bond	choice
concept shapes	a pinned ± pair for a distinction the present decides; a loop (one π-bond, two weak bonds) for one held across time	design (E1a)
a cell	5 bytes at rest, 7 while exploring	design
lookup tables	`SIN256`, `COS256` in Q15; `EXP2NEG`, 1,024 entries; pinned by hash	design
T_{KT}, square-lattice XY	0.89294 J	physics (Hasenbusch, 2005)
T_{KT}, integer kernels	0.8931 J, within 0.1%	measured (E0, lattices to 256)

CONSTANT	VALUE	KIND
$\eta(T_{KT})$	1/4 in the literature; 0.264 measured, log-corrected	physics; measured
the stiffness jump	$\Upsilon(T_{KT}^{-}) = 2T_{KT}/\pi$	physics (Nelson and Kosterlitz, 1977)
the living band	$\Upsilon/(2T/\pi) \in [1.05,\ 1.30]$, setpoint 1.15, per region	choice; target
tile	32 × 32 in 2D; 16^3 in 3D	choice
halo, temporal blocking	8 cells; at most 8 steps	choice
launch slice	at most 20 ms on a display-driving GeForce card	choice
split level, hot region	α = 0.01, χ^2 homogeneity	choice
split level, licensed region	stricter, by rung; set in the writ	open (GAP `writ-constants`)
concept cost	c_0 = 32 bits; c_1 = 1 bit per pinned cell	choice
rent	decay λ = 0.95 per audit, counted in settles; eviction below −64 bits	choice
eligibility decay	λ_e = 0.99 per settle	choice
learning rates	$\eta_J, \eta_h, \eta_K, \eta_v$ = 2^{-8}, 2^{-8}, 2^{-10}, 2^{-6}, in counter units	choice
frontier threshold ε	2 phase steps × relevance	choice
sleep mixing	integrated autocorrelation ≤ 200 sweeps at L = 1024	target

CONSTANT	VALUE	KIND
throughput	≥ 10^{11} / 5 × 10^{11} / 2 × 10^{11} cell-steps per second on the three cards	target; measured without temporal blocking on the 4070 Ti SUPER: 7.4 × 10^{9} settling, 1.8 × 10^{10} exploring

The firm

CONSTANT	VALUE	KIND
confidence multiplier z	2.487	derived
tolerance δ	0.30; 0.40 is the harm dial	writ
identifiability floor	$n_0 = 68.7\,(1-p)/p$ at δ = 0.30	derived
incumbent bar $\bar{g}$	frozen at its upper bound, from ≥ 200 shadow outcomes	choice
licence test	α = 1/20; rung k at $\alpha_k = \alpha\, 2^{-k}$	choice
demotion	α_d = 1/5.5, on every machine-executed outcome	choice
exposure cap X	$100M at the E0 operating point	writ
the brake	κ = supervision created ÷ removed; demote the same day at κ > 1	rule
rectification	$n_0 \to n_0(1-\rho^2)$, power-tuned	derived
effective count	$M/(1+(M-1)\bar{\rho})$	derived
master inequality	$D \geq L\,(1 + n_0\bar{v}/X)$	derived

The remaining kernels

Part V prints five kernels: the tables, exploration, settling, detection and the life loop. These are the other five. They are the design. The running implementation is public at `github.com/bochen2029-pixel/kelvin` under the MIT licence, and **where the two differ, the repository and its receipts win.**

One step of renormalization, in integers. A bound pair inside a block cancels in the sum, so its charge does not survive to the coarse level. That is abstraction.

```
// kelvin/kernels/pyramid.cu
__global__ void restrict_level(const uint8_t* fine, uint8_t* coarse, uint8_t* order, int Wc)
{
  const int X = blockIdx.x * blockDim.x + threadIdx.x, Y = blockIdx.y * blockDim.y +
threadIdx.y;
  if (X >= Wc || Y >= Wc) return;
  const int Wf = 2 * Wc;
  int cx = 0, cy = 0;                                          // four unit vectors, summed in
Q15
  for (int dy = 0; dy < 2; ++dy)
    for (int dx = 0; dx < 2; ++dx) {
      const int th = fine[(2 * Y + dy) * Wf + (2 * X + dx)];
      cx += cos8(th);
      cy += sin8(th);
    }
  coarse[Y * Wc + X] = atan2_256(cy, cx);                      // integer octant lookup, exact
to one step
  const uint32_t r = isqrt64((uint64_t)((int64_t)cx * cx + (int64_t)cy * cy));   // 0 .. 4 x
32767
  order[Y * Wc + X] = (uint8_t)((r * 255u) / (4u * 32767u)); // low where charges crowd
}
```

The frontier, in tile order. The next frontier must replay exactly, so compaction preserves tile order: a prefix sum, never an atomic.

```
// kelvin/kernels/frontier.cu
__global__ void flag_frontier(const TileMeta* tiles, int n, uint64_t eps_q32, uint8_t* flag)
{
  const int i = blockIdx.x * blockDim.x + threadIdx.x;
  if (i < n) flag[i] = (uint64_t)tiles[i].residual_q16 * tiles[i].relevance_q16 > eps_q32;
}
// host: cub::DeviceSelect::Flagged(tmp, tmp_bytes, tile_ids, flag, next, count, n, stream);
```

Learning: wake minus sleep, only where consequences reach. These are the sufficient statistics of the XY family. Each thread owns its cell's counters, so the sums are deterministic.

```
// kelvin/kernels/learn.cu   (sign = +1 in wake, -1 in sleep; equal sample counts in each)
__global__ void accumulate(PageSet pages, const uint32_t* frontier, int sign,
                           const uint8_t* carve_mask, EdgeAcc acc) {
  const uint32_t tile = frontier[blockIdx.x];
  if (!carve_mask[tile]) return;                  // language may perturb; only consequences
carve
  for (int c = threadIdx.x; c < TILE * TILE; c += blockDim.x) {
    const int x = c % TILE, y = c / TILE, th = pages.theta(tile, x, y);
    acc.jx[tile][c] += sign * cos8(pages.theta(tile, x + 1, y) - th);   // statistic for J_x
    acc.jy[tile][c] += sign * cos8(pages.theta(tile, x, y + 1) - th);   // statistic for J_y
    acc.hx[tile][c] += sign * cos8(th);                                  // h e^{i phi}, real
part
    acc.hy[tile][c] += sign * sin8(th);                                  // h e^{i phi},
imaginary part
  }
}
// at night: J += eta_J * acc.j / n; (h cos phi, h sin phi) += eta_h * (acc.hx, acc.hy) / n,
// each within the region's nightly trust bound
```

The homeostat, on the host. It computes in floating point and records the integer it hands the field.

```
// kelvin/host/homeostat.cpp
void homeostat(Region& R, Tape& tape) {
  const double ups   = (R.Ex_mean - R.Ix2_mean / R.T) / R.cells;   // the stiffness estimator
  const double ratio = ups / (2.0 * R.T / M_PI);                    // 1.0 is the Kosterlitz
edge
  R.T = std::clamp(R.T * (1.0 + GAIN * (ratio - 1.15)), T_MIN, T_MAX);   // too stiff: stir
more
  R.invT_q16 = (uint32_t)std::lround(65536.0 * M_LOG2E / R.T);      // the integer the field
uses
  tape.append(Row::stir(R.id, R.invT_q16));                         // recorded; replay never
recomputes it
}
```

A fork copies a table, not tiles.

```
// kelvin/host/fork.cpp
Fork fork(const Fork& parent) {
  Fork f{next_fork_id(), parent.id, alloc_page_table(n_tiles)};
  cudaMemcpyAsync(f.page_of, parent.page_of, n_tiles * sizeof(uint32_t),
                  cudaMemcpyDeviceToDevice, stream);
  retain_pages(parent);                             // shared pages are copied on their first
write
  return f;
}
```

Glossary

TERM	MEANING
agreement band	decisions on which the machine's proposal matched the firm's recorded choice; history licenses it
attended exploration	a logging policy chosen so the machine's alternative is sometimes run, under watch
blast radius	the exposure in flight inside one common-mode group; it caps the licence regardless of evidence
block spin	one coarse cell made from 2 × 2 fine ones; one step of renormalization
causal state	a class of pasts that predict the same future; a minimal predictor's states
charge	a vortex's topological winding number, −1, 0 or +1; the mind's unit of meaning
clock model	the XY model with a discrete circle; the eight-bit medium is the 256-state clock model
coherence boundary	the largest region one mind can hold coupled and one writer can serialise
common-mode group	everything that would be wrong for the same reason

TERM	MEANING
concept	a certified structure that pays rent: a pinned ± pair for a distinction the present decides, a loop for one held across time
coupled	sharing the firm's lattice, instead of being sent copies of it
crystallization	a region cooling into a deterministic map; automation as a phase
cryptobiosis	a quiet tile: state held, nothing computed, revivable at once
D	days to a class's first licence rung
deemed acceptance	a contract clause by which silence, after a window, is a verdict
drive	a push on the medium's boundary, $u\cos(\vartheta - \psi)$; how anything reaches the mind
E-process	a nonnegative supermartingale under the null; evidence that is valid at every look
edge	where a promise meets something the firm does not own, and a verdict arrives
exact channel	a lattice channel written only by the single writer, or by an instrument
face	the firm's one rendering: hop-zero, declared a machine, testifying from rows
floor	the deterministic safe terminal taken at a deadline when every model is off
fold	any table derived from the tape by replayable computation; the mind is one

TERM	MEANING
fork	a copy-on-write branch of the medium; imagination, marked imagined
frontier	the tiles whose residual times relevance clears the threshold; attention
gate	deterministic code that emits one verb with one reason; nothing learned is inside
gauntlet	symbol sources with known causal states; the mind's planted truth
homeostat	the loop that holds each region's stiffness inside the living band
instrument	exact, brainless computation: solvers, the lattice, calculators, exact graders
interior	everything a firm contains to move context between minds that cannot see the same state at once
κ	supervision created over supervision removed; above one, the class demotes
lane	a channel into the mind, with four separate permissions
licence	what a class may do, as a rung with a term, earned from exterior verdicts
living band	just inside the Berezinskii–Kosterlitz–Thouless phase: bound pairs, scale-free structure
loop	a ring of cells whose circulation holds one bit; one π-bond makes its two states degenerate, and state-dependent switching flips it: a toggle flip-flop

TERM	MEANING
machine interior	the rendering, caching and invalidation built around a remote mind
melt	a frozen region returning to thought when its residual climbs
n_0	the outcomes a class needs before its licence widens one rung
nursery	problems with exact graders, where the mind is born thin
plaquette	the smallest square of the lattice, around which a charge is counted
π-bond	a negative coupling; it prefers its two cells half a turn apart
provenance	the concepts a proposal rests on; interior, never published
pyramid	the stack of block-spin levels; the scale of a concept is the coarsest level where it survives
R	the number of times a fact is re-represented between arrival and committed effect
rectification	prediction-powered correction of the evidence floor: $n_0 \to n_0(1-\rho^2)$
remote node	a component that must be sent a copy of the state in order to compute on it
rent	what a concept pays in bits of prediction, measured by fork ablation
resident	coupled to the state it thinks about; present between events; re-derivable from its record

TERM	MEANING
retained stratum	the salted fraction of every licensed class kept on the other kind of mind; the crystal's leftover heat
ρ	the correlation between the machine's grade and the world's verdict
salt	the secret that draws canaries, so the machine cannot choose where it is graded
screening length ξ	$\sqrt{J/h}$; how far a perturbation reaches; memory localises thought
seam	the switch from reversible thought to erasing conclusion; the only place rows are written
settle	Langevin relaxation to a conclusion; erasure, and a row
spore	the signed artifact that carries a firm's competence to a new site, and none of its authority
stiffness Y	the helicity modulus; how rigidly the medium's phases hold together
tape	the append-only, hash-chained record; the only truth
temperature floor	the lowest temperature a region may reach at its licence rung; set by the kernel
terrain	the medium's pinning field; memory; values are written, habits are carved
translator	the omni-modal model at the mind's edge; it may drive, never carve

TERM	MEANING
V0, V1, V2	verdict channel grades; only V0 may move a licence
verdict law	authority is a fold over exterior verdicts on executed choices, drawn on a sample the actor cannot predict
warrant	a signature the law reserves; standing, not cognition
writ	what the firm is for and must never do; authored once, never learned

Figures

ID	FIGURE	FILE
fig1-three-rows	The mind that was sold, the mind that was shipped, the mind this paper builds	10-front.md
fig2-ladder	The gap between the state and the eye	20-part-I.md
fig3-dc	Days to autonomy on one class	30-part-II.md
fig4-seam	The seam	50-part-IV.md
fig5-band	The living band	50-part-IV.md
fig6-chip	Where each kind of mind lives on the chip	60-part-V.md
fig7-cosmology	The firm with one mind	70-part-VI.md
fig8-loop	The resident loop	80-part-VII.md
fig9-scale	The org chart as a scale-space	80-part-VII.md
fig10-thermal	The firm as a thermal image	80-part-VII.md
fig11-circuits	Three circuits of one loop	90-part-VIII.md