

The Physical Prices of Knowledge

To Learn Is to Forget

Alex Stremsky
alex@stremsky.zavinagi.org

July 2026

Abstract

Every act performed on knowledge — discovering, selecting, comprehending, storing, transferring, forgetting — is a physical process, and physics prices some of these acts exactly. We assemble, from established results in the thermodynamics of information, a price list for the elementary operations of knowledge, and show that its entries share one grammar: every floor is a conditional entropy priced at Landauer’s rate, evaluated in the executor’s frame, under four transactional laws. We then work out what the grammar implies for anything that maintains and grows a body of knowledge: a feed law for the subjects of knowledge (capability decays below a minimum intake), the economics of selection, chunking, transfer, and storage, the silent economics of decay and its detection, the statistics of reserves, and the pricing of boundaries and coercion. The framework converts several standing disputes — including the Kurzweil–Modis controversy — into defined measurement problems, and closes by naming the observations that would falsify it.

Contents

1	Introduction	3
1.1	The question	3
1.2	What this paper contributes	3
1.3	How this paper is written	4
1.4	A reader’s map	4
2	Objects: items, corpora, subjects, bodies	5
2.1	One ontology: items and corpora	5
2.2	Content space and diversity	6
2.3	Two senses of having: subject and body	6
2.4	Reach: the boundary as a surface	7
2.5	Complex corpora	8
2.6	Scope of this ontology	8
3	The two currencies and the transactional laws	8
3.1	What we inherit	8
3.2	What we add	9
3.3	The fare: what forced forgetting costs	10
3.4	The loan: what doing costs	10

3.5	The exchange rate is a clock	10
3.6	The conditional floor schema	11
3.7	The four transactional laws	12
3.8	Status	13
4	The subject and its feed	13
4.1	The feed condition	13
4.2	The criterion is contextual — and owns its circularity	14
4.3	Different domains, same law	14
4.4	Frozen subjects	15
5	The exchange asymmetry	16
5.1	What a discovery earns	16
5.2	How the knowledge stock changes	16
5.3	The asymmetry	17
5.4	Demonstration: pricing a famous dispute	17
6	The costs of knowing	18
6.1	Selection: the pan and the fee	18
6.2	Comprehension: the model and the wall	19
6.3	Chunking: the handle and its risks	20
6.4	Transfer: the recipe and the chef	21
6.5	Storage: the rule and the librarian’s mercy	22
7	Decay, its discovery, and its recovery	23
7.1	All rot is silent	23
7.2	Discovery: the remnant	24
7.3	Recovery: three answers to rot	25
7.4	One dial, two hands	26
8	Reserves: the stock and the catalog	26
8.1	Valleys are statistics	26
8.2	The pricing theorem	27
8.3	The release toll	27
9	Boundaries and stranding	29
9.1	Stranded correlations	29
9.2	The commandable set	29
9.3	Reach is not absorption	30
10	What would prove us wrong	30
11	The honesty ledger	31
11.1	Three tiers of claim	31
11.2	Two label distinctions the reader should keep	32
11.3	The open seams	32
11.4	This paper, priced by itself	32

12 The open programme	33
Coda	33
Acknowledgments	34
References	34

1 Introduction

1.1 The question

Knowledge is usually studied as if it were weightless. Epistemology asks what knowledge *is*; economics asks what it *does*; neither tradition asks, with any persistence, what it *costs* — not in money, which is a social convention, but in the conserved quantities of physics, which are not. Yet every act performed on knowledge — discovering it, selecting it, comprehending it, storing it, moving it, sharing it, and forgetting it — is a physical process, executed by physical systems, and physics has had, since 1961, an exact and experimentally confirmed answer to what the irreducible price of at least one such act is.

This paper asks **what does it physically cost to know, to keep knowing, and to know together?** We assemble, from established results in the thermodynamics of information, a *price list in joules and bits*, for the elementary operations of knowledge. We show that the entries on that list share a single mathematical grammar. And we work out what that grammar implies for anything that maintains and grows a body of knowledge — an organism, a laboratory, a firm, a science, a civilization, a machine.

In plain terms. Libraries have budgets; librarians have salaries; students burn glucose while reading. Everyone knows knowledge costs *something*. What is not widely known is that the laws of physics set exact minimum prices for some of these acts — and no technology, however advanced, can ever reduce them further — and that these minimum prices have a common shape. We show here why forgetting is expensive but decay is free, why understanding requires discarding, why sending an expert can beat sending a report, and why keeping an archive alive is a permanent bill rather than a one-time purchase.

1.2 What this paper contributes

The paper makes two linked contributions.

The first is a **vocabulary**: a small set of exactly defined objects — knowledge items and corpora (one ontology, §2.1), the *subject* and the *body* of a corpus as two distinct senses of "having" knowledge (§2.3), and *reach* as the multi-axis condition that draws a corpus's boundary (§2.4). These definitions are chosen for fitness: they are the minimal objects on which the paper's results can be stated without ambiguity.

The second is a **price list with a grammar**: the demonstration that the physical floors under the elementary operations of knowledge — erasing, selecting, reconstructing, recreating, detecting, storing, repairing — are all instances of *one conditional-entropy schema* (§3), interacting through *four transactional laws* that determine who pays, when, in whose frame, and what "paying" even means. The schema and the laws are assembled from established physics, with

the assembly’s own additions marked; their application to the economics of knowledge corpora is, to our knowledge, new.

1.3 How this paper is written

This paper is written in two registers. The **formal register** — the main text — states definitions, results, and their derivations at the precision a technical reader requires. The **plain register** — the passages marked "*In plain terms*" — restates each load-bearing result in ordinary language, with concrete analogies and worked numbers.

The formal register carries the content in full — definitions, derivations, exact statements. The plain register is a second, coarser corpus overlapping the first: incomplete by construction, but for many readers a bridge that lowers the cost of reaching the formal claims. A reader fluent in the formal register **may skip every plain insert without loss**; a reader who is not will find the inserts the cheaper road in.

This design is an application of the paper’s own results. A paper is a message between knowledge corpora, and the cost of receiving it is governed by the same reconstruction economics the paper derives (§6); writing for two different reader populations is building two linked corpora, each overlapping its readers’ holdings, with the overlap lowering comprehension costs: two smaller strides are easier than one big. The paper practices what it prices. The prose throughout is deliberately readable, in the tradition of physicists who wrote prose; wherever the formal tier uses a vivid image, the image accompanies a precise statement and never replaces it.

Every substantive claim belongs to one of three kinds, stated as: *verified* (proved in-model and checked by independent computation; the code is part of the paper’s record), *inherited* (established results of others, cited, with what we take stated exactly), or *candidate* (our proposals, marked as such, with what would falsify them).

In plain terms. This paper is a bilingual book, where both languages say the same, a Rosetta stone. If you read mathematics fluently, skip the plain-terms passages; you will miss nothing. If you do not, read the plain-terms passages; these will help you understand the math.

1.4 A reader’s map

The paper in twelve sentences, one per section.

1. Knowledge has physical costs; this paper assembles their price list and the vocabulary to state it.
2. The vocabulary: items and corpora — one substance at two resolutions — and two senses of having: the body that operates, the subject that thinks.
3. Everything a body does pays in two currencies — a fare for forgetting, a loan for doing — every fare one schema, under four transactional laws, in the executor’s frame.
4. Subjects must also eat: a minimum content intake, below which the make-judge-weave capability decays, in every domain the same.
5. The feed is bought: knowledge unlocks resources, resources buy knowledge — and since knowledge compounds while resources deplete, the loop’s gain is a race, whose most famous dispute the machinery prices.

6. The acts of knowing — selecting, comprehending, chunking, transferring, storing — are the schema's faces.
7. The stored silently decays; discovering the damage is paid; recovery takes three roads.
8. Between resource tranches lie valleys; crossing them takes a reserve stock sized by statistics and a catalog whose speed is the payout rate.
9. Where bodies end, correlations strand, commands are priced in the executor's frame, and reach never absorbs a subject.
10. How to prove all of it wrong.
11. What does, and does not, rest on it.
12. The fields it opens.

2 Objects: items, corpora, subjects, bodies

2.1 One ontology: items and corpora

A **knowledge item** is any persistent, physically instantiated structure that carries information usable by some process: a remembered fact, a skill, a theorem, a recipe, a dataset, a trained AI model's weight configuration, a paragraph of this paper. A **knowledge corpus** is a collection of items together with the structure among them — the links, dependencies, and shared vocabulary that make the collection more than a heap.

We treat these as **one kind of object at two resolutions**. An item is a corpus whose internal structure we are not currently naming: zoom into any "atomic" fact and structure appears (terms, relations, presuppositions). A corpus is treated as an item whenever it is handled as a single unit: cited, transferred, priced, or combined whole. The distinction is a choice of resolution, not a difference in nature, and the paper moves freely between resolutions, always saying at which one a statement is made.

This identification is not a convenience but a load-bearing fact, because one of the paper's central operations — *chunking* (§6.3) — is precisely the act that converts a corpus into an item: packaging an internally structured collection behind a single handle so that it can participate in further combination as one unit. Concepts, theories, and words are corpora that have been successfully converted into items. That the conversion is possible, what it costs, and what it risks are results of this paper.

Chunking removes nothing: the parts do not stop existing as items — uniting algebra and geometry into mathematics stopped no one from using either separately — so no operation is needed to reach a chunk's contents. The one genuine reverse operation is *dissolution*: displacing the packaging itself, so the contents can be re-chunked differently — a paid operation, and §6.3 prices it.

In plain terms. Is a book one piece of knowledge or a bundle of many? Both, depending on how closely you look — it is one object on a shelf and thousands of sentences when opened. We refuse to decide the question globally, because the interesting action is in the *zooming*: the mind's ability to wrap a whole theory into a single mental token (so it can think "relativity" instead of re-deriving it) is one of the most consequential operations knowledge undergoes, and this paper prices it. From here on, "item" means "knowledge, viewed as one unit" and "corpus" means "knowledge, viewed as

a structured collection" — same substance, different magnification. And boxing things up destroys none of them — algebra and geometry lived on after mathematics united them; only melting a box down for recasting costs.

2.2 Content space and diversity

Items live in a **content space**: an abstract space in which distance means the amount of translation, prerequisite-building, or re-derivation that separates two items. We do not need coordinates for this space, only distances, and only to the precision our results use them. A corpus occupies a region of content space; its **diversity** D is the effective breadth of that region — how many kinds of content it holds that are *substantially different* — far apart in exactly this distance — counted so that near-duplicates count once. [Formal construction of D : the companion paper [Stremsky 2026]; this paper needs only its existence and monotonicity properties.]

In plain terms. Picture all possible knowledge as a library where books are grouped by similarity: geometry sits nearer to algebra than to ornithology. A corpus is a shape drawn on the space within this library. Its diversity is not how many books it holds but how much *ground* it covers: a million near-identical cookbooks cover an area only slightly bigger than a single book; a hundred very different books can cover a continent.

2.3 Two senses of having: subject and body

Ordinary language says a person "has" knowledge, a library "has" books, a firm "has" expertise — one verb concealing two different relations. This paper needs both, but separately, and much confusion dissolves once they are named.

A corpus is a **subject** to the degree that it hosts an ongoing capability to generate new knowledge, adjudicate it, and integrate it — in plainer words, to make new knowledge, judge it, and weave it in. Subjecthood, in this paper, is a matter of capability rather than activity. The precision behind that choice: a subject is not a separate thing behind these processes — it *is* these processes, taken as a standing pattern, dispositional when dormant and occurrent when running. That is why a sleeping scientist and a research institute during holidays remain subjects. And the choice has a precise feed-condition, derived in §4: sustaining the capability — that is, keeping its processes runnable — requires a minimum ongoing intake, below which no intention can substitute for fuel.

The criterion is *contextual*: it defines "subject" as this paper's results require — persistent knowledge-*processing* — and makes no claim about consciousness, moral status, or any other sense of subjecthood that other contexts rightly require.

Each of the three component capabilities — generating candidates, adjudicating their viability, integrating survivors — comes in degrees, so subjecthood is graded rather than binary: a database server that mechanically derives entailments, checks fixed constraints, and updates indexes is a subject, of some small degree. The grading calculus, and what follows from distributing the three roles across different bodies, lie beyond this paper.

A corpus's **body** is everything a subject can *jointly operate*: all it can read, write, compare, combine, and erase together, as one accounting unit. The body is defined by operational reach, not by legal ownership, physical containment, or storage location — and the two senses can and do come apart. A frozen archive is a body without a subject: operable, priceable, not generative. A scientific tradition sustained across scattered institutions can be a subject whose

body is patchwork. In this paper, static prices — of operations — are always relative to a body (whose reach determines what the operation can touch); dynamic prices — those of survival, growth, or starvation — are always relative to a subject.

In plain terms. "Having" knowledge means two different things that usually travel together, but sometimes don't. A working scientist both *makes new knowledge* (a subject) and *commands* a body of notebooks, files, and memories she can use together. An abandoned archive keeps the second without the first: everything in it can still be read and copied, but nothing new will ever come from it — a body outliving its subject. The paper keeps the two rigorously apart because they obey different laws: bodies obey the price list; subjects additionally obey a survival condition — they must eat.

2.4 Reach: the boundary as a surface

What exactly does a body include? The criterion is **reach**: a vector of conditions, each an independent axis — **physical access** (can signals travel there and back), **awareness** (does the corpus know the target exists and where), **capability** (can it execute the operations, given formats, languages, tools), and **permission** (is operating there allowed, by whatever enforcement is real). An element belongs to the body when all the axes it requires are satisfied; the boundary is therefore not a line but a **surface in condition-space**, and "is X part of this corpus?" decomposes into per-axis questions with possibly different answers.

Two of the axes are constituted by knowledge: awareness is declarative (the target's existence and address), capability procedural (how to operate) — which is why both can be held by things that are in no sense subjects: a map holds awareness; a calculator holds capability. Hence the definition's main consequence: **a body's boundary is endogenous to its knowledge**. Learning moves the boundary exactly where the epistemic axes bind — learning a cave's location brings the cave into the body; no amount of learning grows an arm. Bodies whose elements are knowledge-stores — this paper's concern — sit near the fully epistemic limit, where boundary and contents co-evolve.

Permission we treat with deliberate brevity, because it may not be a primitive axis at all: its *blocking* half reduces to tokens — keys and credentials, which are capability and access under a social history of distribution — and its *deterrent* half never gates an operation, entering the books only as prognoses and, when enforcement executes, as ordinary operations under Law 1. Likewise **ownership** — a standing claim that can hold at zero reach, functionally an option on *future* reach — contributes nothing to the physical cost of informational processes and lies outside this paper's field: the prices here attach to bodies, and the executor pays in the executor's frame, whatever the title deeds say.

The definition has one philosophical export [Clark & Chalmers 1998]: whether a constantly consulted notebook — and "notebook" may be replaced by "AI assistant" without loss — is *part of the mind* becomes measurable rather than debatable. The notebook is inside the body on the axes it satisfies, to the degree it satisfies them; we do not claim this settles the question — we claim it prices it. And the axes have an outermost shell physics itself supplies: within reach is only what lies within the body's light cone.

Bodies can **overlap**: a storage volume two corpora can both operate lies within both bodies, and "internal versus external" is a predicate of regions and axes, never a verdict on whole agents. What overlap does — to the price of erasure, the possibility of coercion, the stranding of correlations — occupies §9.

In plain terms. Your working body is not everything inside your skin: a paralyzed limb is inside the skin but barely within reach, while the car you drive skillfully, or the notebook you consult every hour, is outside the skin and fully inside it. Reach splits into separate questions: can you get at it, do you know about it, are you able to use it, may you. The middle two are made of knowledge — so learning literally grows your body.

Whose they are — the borrowed car, the torchlight forgotten in a locked cellar — is a claim about the future, not a fact about now: another story, outside this one. And two people can genuinely share a limb of this kind: a folder both can edit belongs to both bodies at once, which will matter enormously (§9) for what erasing it costs, and who pays.

2.5 Complex corpora

Corpora compose. A research group is a corpus whose elements — the members and their private archives — are corpora; a discipline is a corpus of groups; a civilization, of disciplines. The composition structure is a directed acyclic graph rather than a tree, because bodies overlap: a shared archive can be an element of multiple parent corpora at once. Throughout the paper, treating any corpus as one unit is a *declared simplification of resolution*, exactly as in §2.1: the results hold at the level at which the accounting is drawn, whatever internal structure the accounted units possess. Where internal structure changes a result — as it does for the boundary questions of §9 — the paper says so and descends a level.

In plain terms. Everything here is corpora made of corpora, like departments made of teams made of people made of — for our purposes — smaller knowledge-holders still. When we talk about "a corpus" doing something, we've simply chosen a floor of the building to stand on. The theorems don't care which floor; they care only that the bookkeeping is drawn consistently at one level at a time. When the interesting action is *between* floors, we'll say so and take the stairs.

2.6 Scope of this ontology

These definitions are chosen to be sufficient for this paper's results, not exhaustive of the objects a full theory will need. A more complete ontology — translator corpora, transferred objects, institutional roles such as reserve-holders and maintainers, the lifecycle of corpora from founding to freezing — lies ahead of this paper, and the present definitions are expected to be refined as the program grows. We regard this as the normal condition of foundational vocabulary: defined to current fitness, versioned openly, never final by decree.

3 The two currencies and the transactional laws

This section is the paper's physical core. It states what every change to a knowledge corpus is paid in (two currencies, §§3.3–3.4), exhibits the single schema that the paper's floors instantiate (§3.6), and fixes the bookkeeping rules that govern who pays, when, and in whose frame (§3.7). Everything after this section is application.

3.1 What we inherit

- **Erasing one bit of information dissipates at least $kT \ln 2$ of energy** at temperature T — Landauer's principle [Landauer 1961]. The bound has been approached in laboratory

experiments on single colloidal particles and nanomagnetic systems [Bérut et al. 2012; Hong et al. 2016].

- **Performing one elementary operation takes at least $\pi\hbar/2E$ of time** for a system of energy E — the Margolus–Levitin theorem [Margolus & Levitin 1998], one of the quantum speed limits; the quantity $E \cdot t$ it bounds is what its authors themselves call *action volume*.
- These two budgets — one on irreversibility, one on rate — are joined for any physical computer by Lloyd [Lloyd 2000], who computed both for a kilogram of matter and coupled them through the rate at which erased bits can be rejected as heat.
- The connection between speed limits and Landauer’s bound has since become a studied topic in its own right [Deffner & Campbell 2017].
- Del Rio and colleagues [del Rio et al. 2011] find that the work of erasure is governed by a **conditional** entropy — what the eraser must destroy *given everything it holds* — and is therefore **relative to the erasing observer**.
- That correlation itself is a fuel — one bit of mutual information convertible, reversibly and in principle, into $kT \ln 2$ of work — and that measurement and feedback have their own exact thermodynamics is the content of a literature we cite as a whole through its reviews [Sagawa & Ueda 2010; Parrondo, Horowitz & Sagawa 2015].
- That *semantic* information — the part of a system’s information that matters for its own persistence — can be defined without an external labeler, by counterfactual intervention, is described by Kolchinsky and Wolpert [Kolchinsky & Wolpert 2018].

3.2 What we add

What we add is an assembly and a reading:

- The treatment of the two budgets as **currencies of a knowledge economy** — a *fare* and a *loan*, with an exchange rate that is a frequency (§3.4).
- The observation that the floors under seven distinct operations of knowledge share **one conditional-entropy schema** (§3.5).
- The four **transactional laws** as the bookkeeping that makes corpus-level accounting consistent (§3.6).
- The application of all of it to the economics of knowledge corpora, which occupies the rest of the paper.

To our knowledge these are new.

Where any part turns out to be anticipated, priority belongs to the anticipator and the assembly loses nothing.

3.3 The fare: what forced forgetting costs

Definition (contraction). An *erasure* is a deliberate contraction of a corpus's occupied state space (which may occur even as the corpus's stored volume grows): an operation whose completion requires that some configurations the corpus could have been in are no longer available to it — forced forgetting. Three near-neighbors are explicitly *not* erasure, and the distinction carries theorems: **relocation** (moving content between stores, reversibly, is free in principle — no contraction occurs); **spontaneous decay** (rot destroys content but is not an operation the corpus executes and pays for — the universe does it, gratis, a point §7 develops into the direction law); and **overwriting with retention** (if the old value is recoverable from what is kept, nothing has contracted).

The floor. For an executor holding side information O , erasing content S costs at least

$$W \geq kT \ln 2 \cdot H(S | O)$$

joules [Landauer 1961; del Rio et al. 2011]: the fare is proportional not to the size of what is destroyed but to the *conditional* entropy of what is destroyed given everything the executor keeps and can reach. Destroying a copy whose twin you hold is free — $H(S|O) = 0$; destroying the last copy of something structureless pays full price. The fare is paid in **exergy**: ordered, work-capable energy, degraded to heat at the moment of contraction.

In plain terms. Physics runs a strange toll booth: it charges nothing for carrying a crate across the warehouse, nothing for the crate quietly rotting, but a fixed minimum fee for *deliberately destroying* the crate — and the fee depends on what else you own. Shredding one of two identical documents costs nothing at the toll booth (the information still exists in your other hand); shredding the last copy costs full fare. The booth does not care what the document *means*, who wrote it, or whether destroying it was wise — only how much of it, given everything else you hold, is now gone from the universe as far as you are concerned. "As far as you are concerned" is not a loose phrase; it is the law (§3.7, Law 4).

3.4 The loan: what doing costs

Operations — comparisons, transformations, transfers, any elementary step — are bounded in *rate*, not in irreversibility. A system with energy E above its ground state performs at most $2E/\pi\hbar$ elementary operations per second [Margolus & Levitin 1998]. The energy is not consumed by the operation; it must merely be *present*, orthogonalizing the system's states fast enough. It is a **loan**: capital that must be held for the duration, borrowed against time, returnable afterward. A perfectly reversible computation pays only loans; the moment it forgets anything, fares begin.

In plain terms. Thinking fast requires being rich, not spending. The energy that powers an operation is like the deposit a library demands for borrowing: you must have it, it is tied up while you work, and — for ideal, reversible work — you get it back, provided you return the book undamaged. Forgetting is the only thing the universe insists on being *paid* for, which is the fare of §3.3. Real machines, of course, are far from ideal and burn vastly more than either floor; the floors matter because they are the only part no engineering will ever remove.

3.5 The exchange rate is a clock

Two currencies invite an exchange rate, and here it is a *frequency*. A fare (energy, joules) and a loan (energy held over time, joule-seconds) compare only through a time window τ : one bit's

fare equals the loan of one elementary operation sustained for

$$\tau^* = \frac{\pi \hbar}{2 k T \ln 2},$$

about **58 femtoseconds at room temperature** — equivalently, the two floors trade at $\nu_{\text{T}} = 2kT \ln 2 / \pi \hbar \approx \mathbf{17 \text{ THz at 300 K}}$. The window-indexed *merger temperature* — at which one erasure’s fare equals one operation’s loan for a window τ — is $T(\tau) = \pi \hbar / (2 \ln 2 \cdot k \tau)$: about 17 picokelvin for a one-second window, and of the order of the Planck temperature ($2.27 T_{\text{P}}$) when τ is the Planck time.

The natural clock of the comparison is the **thermal tick** $\tau_{\text{T}} = \hbar / kT$, at which the two floors merge up to the dimensionless factor $\pi / (2 \ln 2) \approx 2.27$. The reading of the exchange rate as a temperature-set clock, and the two poles just quoted, are ours to our knowledge; the deep end this touches — thermal time, and the empirical "Planckian dissipation" bound observed in strongly correlated matter — has its own literature, which we cite and do not develop [Connes & Rovelli 1994; Zaanen 2004; Hartnoll & Mackenzie 2022].

This subsection is also where the framework is most directly *exposed to experiment*: the fare’s floor has been probed in the laboratory to within a small factor [Bérut et al. 2012], the approach of real computation’s energy cost toward the Landauer wall is a documented decades-long trend (§10’s Koomey anchor), and any demonstration of below-floor erasure, at any scale, falsifies the fare outright and this paper with it.

In plain terms. How do you compare a toll (paid once) with a deposit (held over time)? Only by fixing *how long*. At room temperature the crossover is absurdly short — tens of femtoseconds — which is physics’ way of saying that for every process slower than molecular vibration, **forgetting is the expensive part**. Cool the world toward absolute zero and fares become cheap relative to loans; heat it toward the Planck furnace and the two merge into a single price. Between those poles — where all libraries, brains, and data centers live — the ledger has two distinct columns, and the fare column is the one that never forgives.

3.6 The conditional floor schema

The paper’s floors now begin to repeat a shape, and we promote the shape to the central object.

Schema. *The irreducible cost of an information act is*

$$\text{cost} \geq kT \ln 2 \times H(\text{what is destroyed or missing} \mid \text{everything held and reachable})$$

— *a conditional entropy, priced at Landauer’s rate, evaluated in the executor’s frame.*

Seven operations of knowledge instantiate it; each face is derived where its section is cited, and the *assembly* — that these seven are one schema — is this paper’s claim:

Face	Conditioning	Where
Erasure	$H(\text{destroyed} \mid \text{kept})$	§3.3 [inherited]
Selection	$H(\text{rejected} \mid \text{retained})$	§6.1 — the Kuhn fee
Reconstruction	$H(\text{tacit} \mid \text{transmitted})$	§6.4 — transfer
Recreation	$H(\text{item} \mid \text{corpus})$	§6.5 — the storage rule
Detection	$H(\text{failure signal} \mid \text{watched variables})$	§7 — the remnant
Storage-worthiness	$H(\text{item} \mid \text{corpus})$, again	§6.5 — the same term, other side
Repair	$H(\text{true value} \mid \text{redundancy})$	§7 — the ECC split

The schema explains, in one line each, several facts the coming sections earn in full. Understanding is cheap for the prepared: a large O shrinks every $H(\cdot|O)$. The last copy is sacred: nothing conditions it. A well-coded archive repairs for almost nothing: redundancy drives $H(\text{true}|\text{parity})$ to zero. Re-deriving what a corpus already implies burns nothing at the floor: $H(\text{item} | \text{corpus})$ is then zero, and the re-derivation is paid in loans and search time alone. And the price of any act of knowledge is always, at bottom, *the price of what you must give up, discounted by what you already have*.

Three conventions make the schema exact, and all three were implicit in the laws below. *Which H* : the entropy is Shannon entropy under the **executor's own predictive distribution** — the same frame Law 4 fixes; operationally, the length of the executor's best available compression of the destroyed-or-missing part, given its holdings. The gap between the executor's distribution and the world's is not ignored but separately billed — it is Law 2's delusion surcharge.

Whose T : the temperature is that of the **bath that absorbs the discarded entropy** — the executor chooses where to dump, and cold dumps are cheaper, which is the fact §7.1's design law builds on. *When*: every conditioning is evaluated **immediately before the act** — afterwards the corpus has changed, and yesterday's price is no quote for today's.

The schema even survives the quantum extreme: with entangled side information the conditional entropy can run negative, and erasure becomes a source of work rather than a cost [del Rio et al. 2011] — the formula holds, sign and all.

In plain terms. Every bill in this paper is the same bill wearing different clothes: you pay for the part that is *gone or missing*, measured against *everything you still have*. A chef who loses a recipe she half-remembers pays half a recipe. A student who reads a textbook he mostly already knows pays only for the new part. An archive with parity disks repairs a damaged file for pennies, because "what is missing given what is held" is nearly nothing. One formula, seven costumes; the rest of the paper is fittings.

3.7 The four transactional laws

The floors say how much; the laws say *who, when, and in whose frame*. We state them once, boxed, and use them everywhere.

Law 1 (executed process). Costs attach to executed physical processes, and to nothing else. No process, no cost; no retroactive charges; an operation not performed is not owed. A late copy of content shared earlier is a *new acquisition*, paid at today's conditional entropies, not a completion of the old transaction.

Law 2 (correlation and belief). Correlation with the world is the asset; belief is the protocol it is used through. A protocol optimized for a distribution the world does not have pays a surcharge proportional to the mismatch (the relative entropy between believed and true) [cf. Kolchinsky & Wolpert 2018]: *delusion has a thermodynamic price*, and verification is the insurance premium against paying it.

Law 3 (direction). Entropy flows in for free and out only for payment. Spontaneous decay — rot — collects; deliberate contraction — erasure — pays; there is no symmetric refund. Landauer prices the filing, never the ink (§7 builds this into the economics of decay). Rot is

thus no exception to Law 1: decay executes no process of the corpus's and charges it nothing; its bill arrives only through the repair processes it makes necessary.

Law 4 (frame). Every conditional entropy in this paper is evaluated in a frame: the holdings and reach of the executor [del Rio et al. 2011]. Costs are observer-relative — the same erasure can be free for one executor and maximal for another — and there is **no view from nowhere** in the economics of knowledge. Where this paper appears to quote a frame-free price, a frame is implicit and stated nearby.

In plain terms. Four rules of the road. *One*: you are billed for what is actually done, when it is done — the universe keeps no tabs and honors no backdating. *Two*: what you know is your capital, but what you *believe* is the machine that spends it — and a wrong belief runs the machine against friction you pay for. *Three*: decay is free and cleanup costs; there are no refunds in the other direction. *Four*: every price on this list is a price *for someone* — the person doing the deed, with their knowledge and their reach — and two doers of the same deed can face wildly different bills. The frame rule (four) will do more work in this paper than any other single idea.

3.8 Status

Of the claims in this section: the two floors and their conditional refinement are **inherited** physics, experimentally supported where experiment has reached; the currencies reading, the exchange-rate clock with its two poles, the schema, and the four laws as stated are **ours** — the schema's faces are derived individually in their sections, and the laws are bookkeeping made explicit rather than new physics. What is *not* claimed anywhere in this paper: that corpus dynamics (growth, competition, collapse) follow from the floors alone. The floors bound; they do not drive. The paper's honesty ledger (§11) grades every downstream claim on exactly this line.

4 The subject and its feed

Section 2 separated the two senses of having knowledge; §3 priced what bodies do. This section states the one law that applies to subjects and not to bodies: **remaining a subject requires a minimum intake**.

4.1 The feed condition

A subject, per §2.3, is a corpus hosting live capability to generate, adjudicate, and integrate new knowledge. That capability is not a possession but a process — populations of skills, active vocabularies, calibrations against a moving world, working relationships among the corpus's parts — and a process must run to persist. Running spends resources, energy above all, and resources are bought with knowledge (§5 prices that exchange); but having currency alone keeps no subject alive — buying enough food with it does. This loop also needs **content**: new material to calibrate against, problems that keep the adjudication criteria honest, enough throughput to keep the machinery in working shape. It is this epistemic intake whose minimum the law states.

For a subject in a context there exists a **subsistence intake** X — a flow, in content per unit time — such that at sustained intake $Y < X$ the capability can no longer be maintained:

it decays, whatever the subject intends. Intention cannot substitute, and by Law 1 it could not: maintaining capability means executing processes, and executed processes are paid in the currencies of §3.

One inequality falls out here for later use: sustained operation requires an integration bandwidth $B \geq X$ — a mouth too small for the subsistence intake starves its owner at a full table — and §8 will lean on it.

In plain terms. A hearth-fire is not a thing but a process, and it must eat. A brilliant mind, a great laboratory, a living language — each keeps its ability to make, judge, and weave in new knowledge only while fed: problems in, conversations in, corrections in. Starve any of them — cut the students, the funding, the speakers — and that ability quietly dies, long before any book is lost. Wanting to stay brilliant feeds nothing. And the mouth matters as much as the pantry: whoever cannot swallow fast enough starves in front of food — remember this when we price rainy-day reserves.

4.2 The criterion is contextual — and owns its circularity

What counts as "the capability being sustained"? This paper's criterion is deliberately contextual: the subject persists — remains a subject — if its knowledge processing stays *fit for its context*: it continues to generate, adjudicate, and integrate content that its niche accepts as knowledge. We write this $\dot{K} \geq 0$ informally: the corpus's viable output is not shrinking.

The definition is circular in the way all viability definitions are — life is what maintains itself as life — and we own the circularity rather than hide it: the anchor is Kolchinsky and Wolpert's [2018] definition of semantic information, where the information that *matters* is exactly the information causally necessary for the system's own persistence, identified by counterfactual intervention rather than by any external labeler.

A subject's feed, in these terms, is the influx of content causally necessary for the persistence of its capability as a subject.

4.3 Different domains, same law

The feed condition is exhibited not only by organisms, but by everything that hosts a subject. We state its shape once.

The shape. In each domain there is a **stock** — the stored knowledge, which is a body and obeys §3's prices — and a **subject** — the living generate–adjudicate–integrate process of §2.3, which obeys this section's law. The subject has a concrete **feed**: some flow of problems, corrections, exchanges, newcomers.

Cut the feed, and one specific signature follows: **the stock persists while the subject decays**. The holder keeps what it knew and loses the capability to know more. Which of the three capabilities fails first varies by domain and by what exactly is starved — an ordering this paper leaves open. One example of what it may contain: since deliberate forgetting pays fares (§3), a starved corpus may first lose the ability to *afford* forgetting — ossification, the inability to restructure, as a poverty symptom.

Here are examples of this shape:

Domain	The stock	The subject	The feed	The decay, by its household name
An organism	memories, skills	the learning brain	percepts, problems, practice	"rusty," "out of practice"
A firm	patents, documentation	the R&D capability	hiring, licensing, frontier contact	lost absorptive capacity
A science	the archive, the journals	the research community	students, fresh problems	scholasticism
A language	grammars, dictionaries, texts	the speaker community	children acquiring it, daily use	a <i>dead</i> language

The last column shows that we already have a name for this decay in every domain — proof the thing is known, just not as the same thing everywhere. One of those names is also a technical term with a literature behind it: economists studying firms found that the ability to *understand and use* outside results is a separate asset from the patents a firm owns — it decays when hiring and frontier contact stop, even while the patent stock stands untouched — and they named it *absorptive capacity* [Cohen & Levinthal 1990]. Different vocabulary, same law: the stock stayed; the subject starved.

If this subsection is correct, its law must read correctly in all four domains without stretching; the reader is invited to verify that against cases they know. The shape can be falsified — for example, by a subject that thrives unfed, or by a stock whose decay improves its subject.

In plain terms. Everyone knows this shape — under different names, in different areas; it has simply never been united into one law. A violin does not forget how to play music — hands do. Latin lost not the books written in it, but the children who spoke it. The body is an object; the subject is a process. A subject cannot persist without its body. A body can persist without its subject — but it is dead. Keeping an object is cheap, and §7 prices it. Keeping a process requires feeding it — that is this section’s law.

4.4 Frozen subjects

The feed condition admits a strategy: **freezing**. A subject may deliberately suspend its knowledge-processing — the mothballed laboratory, the seed bank, the codebase in maintenance mode, the language preserved in grammars awaiting revival — trading the feeding bill for body storage costs (§7 prices these), plus a **reactivation price** payable later.

Reactivation is not mere thawing: the frozen subject’s calibrations have staled against a world that moved, so revival pays a re-calibration bill that grows with how far the world has drifted meanwhile — the same conditional-entropy logic as everything else: what is *missing given what is held* has grown while nothing was fed. The revival of Hebrew as a spoken language is the exhibit that freezing can be survived and reversed; its cost and rarity are the exhibits that the reactivation price is real. The full pricing of freeze-and-revive belongs to the reserves and decay machinery (§§7–8); here we record only that the option exists, that it is a genuine option (bodies keep; subjects sleep), and that its price has the schema’s shape.

In plain terms. You can put a subject to sleep and pay only the rent, saving on groceries. The seed bank does it; so did Hebrew, for centuries, waiting in books. But waking is not free: the world walks on while the sleeper lies still, and the bill for catching up grows with every year of sleep. Sleep is a loan against the future’s grocery bill — sometimes the right loan, but never a free one.

5 The exchange asymmetry

Knowledge earns. A corpus's content unlocks resources — techniques, energy, materials, attention, funds. Those resources in turn feed the corpus: they enable it to obtain new knowledge. This section prices the exchange in both directions, states the one asymmetry on which everything compounds, and then demonstrates this machinery on a famous dispute.

5.1 What a discovery earns

The value a new item brings to a corpus travels through two channels, and they scale differently. The **quantity channel** is the item's direct contribution — one more usable result, priced at its use. The **structure channel** is the item's contribution to the corpus's *variety*: what it makes reachable, combinable, and comprehensible that was not before — the value of breadth at fixed size, which we denote m and which §6 will show doing independent work. The two channels come apart cleanly in practice: a thousand near-duplicate results pay much quantity but almost no structure; a first bridge between two previously unconnected regions pays little quantity but much structure.

Exchange under uncertainty carries a premium with an exact form. When an item's realized value is uncertain by a lognormal factor with log-variance σ^2 , the expected shortfall between naive and realized pricing carries the factor

$$\chi = e^{\sigma^2} - 1$$

— small when uncertainty is small, and growing *exponentially* in the variance [derived here and verified by independent simulation, per the verification protocol of §1.3]. The premium is why frontier knowledge trades so badly. At the frontier, σ^2 is largest. The spread between what a discovery is worth and what anyone will pay for it sight-unseen is therefore widest exactly where the content is newest. This is a floor under the awkwardness of markets for ideas — not a defect of any particular market.

In plain terms. A new fact earns twice: once for what it *is*, and once for what it *connects*. And selling new knowledge is hard for a reason with a formula in it: the newer the knowledge, the less both the seller and the buyer know what it is really worth, and the mistrust premium grows explosively with that ignorance. The used-car problem is mild; the unused-idea problem is the same disease at fever pitch.

5.2 How the knowledge stock changes

A corpus's knowledge stock K changes as the sum of three flows:

$$\dot{K} = G + T - R$$

Generation G : new content produced in-house — exploration spend times the corpus's conversion efficiency, the pan that turns effort into finds (§6.1 prices the panning).

Transfer intake T : content acquired from other corpora, each transfer paying the tolls of §6.4. (Outbound transfer, note, subtracts nothing: knowledge is nonrival — §5.3 — and sharing it does not shrink the sharer's stock. What a corpus can lose, it loses through rot or deliberate erasure, never through giving.)

Rot R : the silent losses of §7, always on, never announced.

The equation is bookkeeping, not theory. It says only this: every change in K belongs to one of the three flows. The theory is elsewhere — in what each flow costs, and the next two sections deliver it flow by flow: G is priced in §6.1, T in §6.4, and R , the silent one, owns §7 entire. How the flows feed back on each other is a dynamics this paper does not enter.

5.3 The asymmetry

There is an important difference between knowledge economics and resource economics; the entire program of which this paper is a part stands on it. **Knowledge is nonrival and compounds; resources are rival and deplete.** Using a fact does not use it up; sharing it multiplies its holders rather than dividing a stock; and new knowledge *unlocks* resource tranches that were physically present but operationally invisible — ore that was not metal until metallurgy, sunlight that was not electricity until photovoltaics. Resources, meanwhile, deplete with use, tranche by tranche. So the loop — knowledge unlocks resources, resources feed knowledge — has one compounding leg and one depleting leg, and whether the loop's gain g exceeds unity is not settled by either leg alone but by their race. That $g > 1$ is *possible* is the compounding leg's gift; that it is never *guaranteed* is the depleting leg's discipline.

In plain terms. Bread eaten is gone; a recipe used is still there, and teaches variations. But recipes alone bake nothing — flour runs out, tranche by tranche, and each new pantry must be *unlocked* by knowing something new about the world. Civilization is the race between the recipe book fattening and the pantry emptying. Neither side of the race wins by definition; that is why it is a race.

5.4 Demonstration: pricing a famous dispute

Two celebrated positions describe the long-run shape of technological development. One, associated with Kurzweil, holds that development compounds — accelerating, exponential or faster, because each capability speeds the acquisition of the next [Kurzweil 2005]. The other, associated with Modis, holds that all growth processes are logistic — that every exponential is the rising flank of an S-curve whose ceiling is not yet visible, and that extrapolating the flank is the oldest error in forecasting [Modis 2002]. Both positions are argued by serious people from real data, and the dispute has run for decades. We claim the machinery above prices it — and *dissolves it as posed*.

Let us decompose the context in which "the curve" is observed along two measurable axes: **which element currently binds** (the compounding one — knowledge and its unlocking flow — or the depleting one — the current resource tranche), and **the horizon of observation** (within one tranche's exploitation, or across tranche transitions — the latter always evaluated over the current situation, since no bounded observation decides an unbounded race).

Four sub-contexts result, and the framework assigns each its curve. *Within a tranche, ceiling distant*: exponential flank — both positions fit, and the data cannot distinguish them. *Within a tranche, ceiling near*: logistic bend — the Modis reading is correct, and locally always eventually correct, because every tranche is finite. *Across tranches, unlocking outpacing depletion*: the staircase of S-curves acquires an exponential envelope — the Kurzweil reading is correct, and the acceleration is real, because the compounding leg is winning the race. *Across tranches, depletion outpacing unlocking*: the envelope itself bends — logistic-of-logistics — and the Modis reading is correct again, at the scale he always intended.

The score, honestly kept: the Modis reading is correct in two cells, the Kurzweil reading in

one, and in the remaining cell — the exponential flank far from any ceiling — the data can crown no one. Each disputant's error consists of universalizing the cells where they are right.

We give the two regimes the names they have earned — development in **Kurzweil mode** (unlocking outpaces depletion; the envelope of stacked S-curves compounds) and in **Modis mode** (the current tranche's ceiling binds; the S-curve bends) — coinages of ours so far as we know, offered without claim of priority. The dispute, as posed — "which curve describes development" — dissolves: *development has no single curve; it has a regime map, and the regimes are diagnosable.*

The cell criteria are measurable in principle — is the binding constraint rival and depleting; is the unlock flow outpacing tranche depletion — without waiting to observe the curve's shape, so the decomposition is not circular. We do not claim the measurements are easy; we claim they are defined, which the dispute's own vocabulary never was. And the decomposition yields what neither position offers: a prediction of **alternation** — trajectories should switch between Kurzweil mode on fresh tranches and Modis mode as tranches exhaust, a signature visible in the historical record; deriving the regimes, the envelope mechanism, and their historical fingerprint from this structure lies beyond this paper.

One residue survives the dissolution, and we convert it rather than judge it. Each side keeps an asymptotic comeback: the finitist notes that the universe is finite, so depletion must win *eventually*; the accelerationist notes that knowledge keeps unlocking resource classes no one had counted, so "eventually" keeps receding. This paper referees neither.

The asymptotic question is exactly **the race of §5.3** — whether unlocking flow sustainably outpaces depletion — an open, defined, in-principle-measurable dynamical question, not a matter of philosophical temperament. A sharper form of the criterion is derivable — does an average tranche's yield, net of the subject's existence costs over the tranche's productive span, buy enough knowledge to unlock more than the tranche delivered? — and lies beyond this paper; we state it here so the question is seen to be well-posed. Converting a decades-old dispute of temperaments into a measurement problem is what a price list is for.

In plain terms. Two prophets have argued in the town's market for decades. One says: every curve bends — nothing grows forever. The other says: this curve escapes — each power buys a greater next. Our answer: look at which wall is closest. Climbing inside one pantry, the shelf always eventually empties — the first prophet is right. Leaping from pantry to pantry, the leaps have gotten faster for two centuries — the second prophet is right. Whether the leaping can outrun the emptying *forever* is not a thing either prophet knows; it is a thing the world will measure. We did not crown a winner; we drew the map on which their war was fought, and the map — unlike the war — has coordinates.

6 The costs of knowing

This section prices the first two flows of §5.2 — generation and transfer — by pricing the acts they are made of: selecting, comprehending, chunking, transferring, storing. Each act is a face of §3.6's schema: the same conditional entropy, evaluated in the executor's frame, wearing five costumes.

6.1 Selection: the pan and the fee

Selection is the act of choosing what to retain. Exploration produces candidates — observations, conjectures, variants — in quantities that dwarf what a corpus can keep, and selection is where naive accounting first goes wrong: its cost is not the keeping but the *judging and the discarding*.

A corpus with conversion efficiency η turns a flow of exploration effort into a flow of retained finds; the efficiency is bought with adjudication work — every candidate examined pays loans — and the discarding itself has a fare with the schema’s shape.

When retained content displaces an older organization of the corpus — a new framework replacing an old one, a better taxonomy overwriting its predecessor — the displacement is a genuine contraction — of possibility, not of size — priced at

$$kT \ln 2 \times H(\text{rejected} \mid \text{retained}):$$

the conditional entropy of what is given up, given what replaces it.

We give this fare a name we will use throughout the program: **the Kuhn fee**. The fee explains, in one term, a familiar asymmetry of intellectual life. Where the old organization is nearly *expressible in* the new one — a special case, a limiting approximation — the conditional entropy is small and the transition cheap: Newtonian mechanics survives inside relativity as a limit, and the fee for adopting relativity was, in this precise sense, modest. Where the old organization is poorly expressible in the new — where its distinctions have no images and its questions no translations — the conditional entropy is large, the fee is steep, and the transition meets resistance that is not mere stubbornness but *priced* reluctance. What the philosophy of science calls **incommensurability** is, in this vocabulary, exactly the high-conditional-entropy case of the Kuhn fee [Kuhn 1962].

The naming honors the observer of the phenomenon; the price tag, to our knowledge, is new here.

One consequence is testable and listed in §10: selectors under budget pressure harvest low-uncertainty candidates first — cream-skimming is *rational* under the fee structure — so fields should show a systematic late-surprise pattern: the high-variance candidates, deferred, deliver their occasional enormous values late.

In plain terms. The prospector’s pan is the whole economics of research in one tool: mostly it discards, and the discarding is the work. And when a better map of the goldfield arrives, adopting it costs — not printing costs, but the price of everything the old map organized that the new one has no place for. If the old map folds neatly into a corner of the new one, the switch is cheap and quick. If the old map’s marks mean nothing in the new one’s language, the switch is expensive, and the old prospectors’ famous stubbornness is, at least in part, an honest reading of a real bill.

6.2 Comprehension: the model and the wall

Comprehension is the act of building, within the corpus, a working model of an item: links to what is already held, expectations about its behavior, the ability to answer questions the item’s raw text never answers explicitly. The construction pays loans throughout, and its floor has the schema’s discount structure: the more the corpus already holds, the less remains to build — understanding is cheap for the prepared, at a rate the conditioning makes exact.

The discount has a dark inversion: the *unprepared* pay more for the same new content — a nearly empty corpus pays more to walk the same distance than a rich one — so preparation gaps self-amplify with no suppressor anywhere in sight, a passive divergence engine whose consequences for knowledge hierarchies reach beyond this paper.

The engine has a rival, and the rivalry is worth stating: the same growth that cheapens *accumulation* for the established corpus raises its price of *reorganization* — the Kuhn fee of §6.1 scales with how much standing structure a shift displaces, and the fresh corpus carries almost

none. So the two forces pull opposite ways: incumbents learn cheaper but turn slower; newcomers pay more per step but nothing at the revolution. When a shift is large enough that displacement dominates the discount, the advantage flips to the young — Planck’s principle, that science advances one funeral at a time, is this crossover wearing its bitterest clothes. A parametrized form of the crossover is derivable, and lies beyond this paper.

But *complete* understanding is not the standard, and cannot be. Every model omits; the question is what omissions matter, and that question has no answer inside the item itself. The answer lives at what we call the **purposes wall**: what counts as understood is relative to the purposes the understanding must serve, and purposes are supplied by the subject’s viability context, not by the content [Kolchinsky & Wolpert 2018 gives this a physical form: the information that matters is the information causally necessary for persistence]. The frame principle again — Law 4 wearing epistemological clothes: there is no purpose-free standard of comprehension, only comprehension *for* something.

This yields the paper’s first **corpus-maturity gauge**: the density of working links between an item and the rest of the corpus measures how far comprehension has proceeded — operationally, how many kinds of questions about the item the corpus can answer without going back to the original text.

Two corollaries of the purposes wall, still candidate-labeled: a corpus serving several purposes builds *parallel families of links* for the same item — one geometry of relevance per purpose — and, invertedly, comparing the families an item acquires under equal investment measures the *distance between the purposes themselves*: goals made measurable through the geometries of understanding they induce.

The gauge is therefore item-, corpus-, and *purpose*-dependent — a reading of the frame principle, not a defect. And the links themselves raise a structural question we defer with one preview: many-ended links do not, in general, reduce to pairs (three items can carry a joint regularity invisible in every pair), and the one ontology of §2.1 already whispers the rest — a link is a small item, and pays like one: created at a cost, stored at a cost, destroyed at a fare.

In plain terms. You understand a thing when you can use it, predict it, and connect it — and "enough" understanding depends entirely on what for. A commuter understands the train schedule; an engineer understands the train machinery. Neither understanding is fake and neither is complete; each is *paid up to its purpose*. The map is never the territory in full verbosity — but it is not trying to be; it is trying to be verbose enough to get you home. And the price of a map is set by the trip, not by the terrain.

6.3 Chunking: the handle and its risks

Chunking is the operation that converts a corpus into an item — and here §2.1’s promise, that the corpus–item identity carries theorems, comes due. The operation compresses an internally structured collection behind a stable interface — a name, a notation, a callable skill — so that the collection participates in further thought as *one unit*. The chunk’s value is combinatorial: what costs the working register several slots before chunking costs one slot after, and every structure built on top inherits the saving. Chunking is how a finite mind holds what carries more details than its working register can; concepts, theorems, words, and trained reflexes are its products.

The operation’s prices follow the schema: building the chunk pays comprehension (§6.2) plus the compression work, and **mis-chunking** carries the operation’s characteristic risk — an error sealed inside a chunk is inherited by *every reuse*, so the expected cost of a bad chunk is its error

times its reuse count; small mistakes in foundations compound precisely because foundations are what chunks are for. Hence the paper’s second **corpus-maturity gauge**: the depth of reliable chunking — chunks built of chunks built of chunks — measures how much structure a corpus can wield per unit of working register.

Un-chunking. Chunking needs no undo for access: the parts do not stop existing as items, and where a chunk’s internals are stored only in compressed form, recovering them is the storage rule’s decode (§6.5, §8.3), not a new operation. The one genuine reverse operation — un-chunking proper — is **dissolution**: displacing the packaging so the contents can be re-chunked differently. It is a contraction of possibility, and pays twice: the Kuhn fee on the interface (the old handle’s distinctions, given the new ones), plus a **migration debt** that scales with the old chunk’s reuse count — every structure, habit, and dependent chunk built on the old interface must be re-fitted or breaks silently. The same multiplier that makes mis-chunking dangerous makes dissolution expensive; reuse is leverage in both directions.

Dissolution is nonetheless sometimes mandatory, and the condition is a clean inequality — dissolve when the accumulated mis-fit of the old chunk (its error times its reuse) exceeds fee plus migration — which the software industry rediscovers continually under the name *refactoring*, and pays for exactly as written. The characteristic risks are dissolution’s own: orphaned dependents that fail without announcement (§7 prices unannounced failure), and *partial* dissolution, which leaves two packagings competing for the same contents and charges the register double until one wins.

A remark connects this to the machine-learning present: training a model *is* industrial-scale chunking, and the phenomenon called catastrophic forgetting is a mis-managed Kuhn fee — new content displacing old packaging without the displacement being priced, scheduled, or even noticed.

In plain terms. A chess master does not see thirty-two pieces; she sees four or five *situations* — chunks — and thinks with those. That is the whole trick of expertise: wrap a thousand details into one handle, then think in handles. The dangers are the obvious ones for anything sealed: a crack inside the handle travels into everything built with it, and taking a handle apart to re-wrap it costs you every habit that used the old grip.

6.4 Transfer: the recipe and the chef

Transfer is the act of moving knowledge between corpora — the program’s oldest question, and the schema prices it exactly. A transfer is **complete** when the receiver can do with the item whatever the sender could — when, with respect to this content, the receiver has become the sender. The floor is a correlation-creation cost at the *system* level: the receiver’s corpus must come to hold, reliably, structure it did not hold; who pays which share — sender’s encoding effort, channel, receiver’s reconstruction — depends on protocol, but the total is bounded below by what the receiver lacks:

$$\text{cost} \geq kT \ln 2 \times H(\text{item} \mid \text{receiver's holdings}),$$

the schema’s reconstruction face. Two established results give the floor its economics. **Side information discounts the message**: the receiver’s prior holdings reduce the required transmission exactly as coding theory says they can [Slepian & Wolf 1973] — shared context is not a nicety of communication but a quantitative discount on every message, which is why specialists converse in shorthand and why this paper spent §2 building shared context with its reader. And

the discount explains **apprenticeship** as an investment: years of co-location are the purchase of side information that makes the master’s subsequent messages cheap; the co-location premium across fields should track how much of each field resists codification — a §10 test.

The resistant part has a name and an anatomy. Every corpus has a **tacit core**: content whose reconstruction floor from any feasible message is highest — the calibrations, muscle-fitted procedures, and judgment textures that no codification exhausts [Polanyi 1966].

The definition is economic, not metaphysical: nothing physical is untransferable in principle — the floor is always finite — so the threshold where the tacit core begins is not absolute; it sits wherever the transfer cost becomes prohibitive, and prohibitive is priced by the payer and the stakes. Past that threshold, practice becomes the only affordable channel.

The core presents two interfaces to the corpus that contains it: **qualia**, the atoms of its sensor side, and **actia**, the atoms of its effector side. The tacit core’s economics produce the section’s most practical theorem: when $H(\text{tacit} \mid \text{codifiable})$ is large, **moving the expert is cheaper than sending the report** — sender mobility substitutes for message completeness — which is why laboratories poach scientists, not papers, and why the recipe has never yet replaced the chef.

Transfers also scale asymmetrically: a message composed once serves any number of receivers (the sender’s cost amortizes; each receiver still pays reconstruction), which is the economics of teaching, publishing, and broadcast in one clause.

In plain terms. Trying to learn what someone else knows is never free, and the bill lands on what you *don’t* already have: learning an explanation of chess costs a grandmaster an effortless eyeblink, and a beginner a semester of hard study — same content, different bills. That is also why apprentices live in the workshop for years: they are buying the background that makes the master’s every later sentence cheap. And some of what the master has never fits into sentences at all — the hands know it, the eyes know it — which is why, when a rival wants a bakery’s secret, they do not steal the recipe. They hire the baker.

6.5 Storage: the rule and the librarian’s mercy

Storage is the act of keeping; what to keep is a standing decision, and the store’s content is that decision’s product over time — not everything the corpus could recreate, and not nothing either. The decision has a rule with the schema on both sides. Keeping a derivable item costs its maintenance: $\text{rot rate} \times \text{size} \times \text{horizon} \times \text{repair price}$ (§7 supplies each factor).

Not keeping it costs an expected recreation: the probability it is needed times the price of recreating it — and recreation takes whichever of two paths is cheaper in the case at hand: *recall-with-repair*, where a degraded copy or a neighboring corpus still holds enough to restore from, or *full re-derivation*, priced by the schema’s recreation face, $kT \ln 2 \times H(\text{item} \mid \text{corpus})$, plus the search that finding a derivation entails. The rule:

store a derivable item iff $\rho \cdot B \cdot c_{\text{repair}} \cdot \tau < p \cdot c_{\text{recreate}}$

maintenance costs less than expected recreation.

Its limits organize the whole storage economy. Both sides move with circumstance — need probability and recreation price on one, rot, size, and repair price on the other — but the horizon τ multiplies only the maintenance side. As τ grows, maintenance therefore grows without bound while expected recreation stays bounded, and the rule tightens toward *keep only the incompressible core and re-derive the rest*.

Where recreation is Perelman-priced, the rule says the opposite: keep it, armor it, and never let it lapse. Perelman-priced content is content whose re-derivation carries enormous cost or requires capabilities the corpus might never again possess. Re-deriving the proof of the Poincaré conjecture from nothing would take years of work by a mathematician of Grigori Perelman's caliber — the conjecture stood unproven for a century before his work. The same conditional entropy $H(\text{item} \mid \text{corpus})$ thus serves twice, as the schema's table promised: read one way it is the recreation bill; read the other it is the storage-worthiness of the item — what the item adds that the corpus does not already imply.

The rule's deepest consequence gives this paper its subtitle, and the chain deserves to be walked in full: thinking requires generalizing; generalizing requires discarding differences — forgetting; and forgetting, by §3, is never free. To think is to forget — that is Jorge Luis Borges's theorem; to learn is to forget at a price the schema sets — that is this paper's variation, and its subtitle.

A corpus that kept everything — every observation at full resolution, every intermediate, every near-duplicate — would not be a better knower but a worse one: its retrieval tolls would grow monstrous, its maintenance would eat its feed, its generalizations would drown in their own instances — it would not be able to think at all.

Borges wrote the theorem as a story: Funes the Memorious, cursed with perfect memory, could not think because *to think is to forget a difference, to generalize, to abstract* [Borges 1942]. Forgetting — priced, deliberate, schema-governed forgetting — is not the failure mode of a knowledge corpus. It is a load-bearing operation of knowing.

In plain terms. A library is usable only through its catalog, and a catalog exists only by discarding — of each book it keeps the title, the author, the shelf, and nothing else. A catalog that discarded nothing would be the library all over again, and no help to anyone. Keep what you cannot rebuild; keep what you will surely need; for the rest, keep the seed and trust the garden. Or keep the catalog alone — names, authors, and where each book lives — and let everything it merely points to live in the cellar, or in other libraries that hold these books.

The man who forgot nothing, in Borges's story, could not think at all — every leaf stayed with him forever, complete in every detail, and because of that he could never see the forest. Thinking runs on forgetting the way walking runs on falling: controlled, priced, and pointed somewhere.

7 Decay, its discovery, and its recovery

The third flow of §5.2 remains: R, the losses. Everything stored is under attack — silent, ceaseless — by time and the decay it brings. This section prices the attack (§7.1), the finding of its damage (§7.2), and the three ways of undoing it (§7.3) — and ends by exporting one instrument whose reach extends far beyond archives (§7.4).

7.1 All rot is silent

All decay is silent. No stored structure announces its own corruption; the appearance of "self-announcing" damage is always, on inspection, a paid watcher already in place (§7.2). And decay is *free* — for the universe: by Law 3, entropy flows into a corpus without any process of the corpus's executing or paying, while pushing it back out is always billed. The direction is the whole asymmetry of archival life: rot collects; repair pays.

The decay has an anatomy of three terms, distinguished by spectrum and therefore by countermeasure:

$$\text{rot rate} = A e^{-E_a/kT} + R_{\text{tunnel}} + R_{\text{env}}.$$

The first term is **thermal**: degradation is thermally activated, rising exponentially with temperature along the Arrhenius law — the basis of every accelerated-aging test in engineering. The second is **quantum**: tunneling processes that persist as $T \rightarrow 0$, measured as the flattening of Arrhenius plots at low temperature — cooling cannot touch them. The third is **environmental**: non-thermal bombardment — cosmic rays and kin — energy the environment delivers outside its thermal spectrum, attackable not by refrigeration but by shielding.

Two of the channels are environment-delivered — no energy arrives, no damage is done — and can be starved by isolation. The third is subtler: tunneling spends the stored configuration’s own metastability, and needs the environment only as a place to shed the released energy; isolation cannot touch it.

What fights it is **size**: tunneling rates fall exponentially with barrier width and mass, so heavier, wider barriers per bit suppress it — bought at the price of storage density. One subtler isolation can also reach it: the transition must shed its released energy somewhere, and an environment engineered to refuse that energy — offering no modes to decay into — inhibits the decay itself. But this is isolation from *sinks*, not from attackers, and shares only the word.

The countermeasures thus sort the anatomy: refrigeration fights the first term, size the second, shielding the third — each in a different currency: power, density, mass. Whether the quantum residue is ultimately connected to temperature-like properties of the vacuum — horizons in expanding spacetime carry a genuine temperature [Gibbons & Hawking 1977], and the thermal-time literature [Connes & Rovelli 1994] runs adjacent — is beyond this paper; if such a connection exists, the three-term anatomy would collapse toward a single environmental parameter.

The anatomy’s asymmetry against §3’s floors yields a design law. Cooling suppresses the thermal rot channel *exponentially* while cheapening the erasure fare only *linearly* — and computation is rate-bound by the energy present (§3.4’s loan), so speed wants energy concentrated where the operations run, while storage wants only stable wells. The optimum segregates: **store cold, compute hot** — visible in every data center’s tiering, every archive’s basement, and every genome’s chemistry, and derived here, to our knowledge for the first time, from the mismatch of two temperature dependencies; the practice, of course, is ubiquitous.

In plain terms. Nothing on any shelf is safe, and nothing on any shelf will ever tell you that it was hurt. Heat is the main vandal — cool the vault and his visits fall off a cliff — but two accomplices remain: a quantum pickpocket whom no cold deters, and cosmic buckshot from the sky, stopped by thick walls but never by thermostats. Hence the shape of every serious archive since the pharaohs: deep, cold, and quiet — while the workshops upstairs run hot. The cellar and the forge were always thermodynamics.

7.2 Discovery: the remnant

Damage found is damage paid-to-find. Detection is a process (Law 1), and its floor has the schema’s shape: distinguishing *this failure’s signal* from everything the corpus already observes costs at least

$$kT \ln 2 \times H(\text{failure signal} \mid \text{watched variables})$$

per check, per failure mode. (The fare localizes at the meter's reset: measurement itself can in principle run reversibly, but each check's record must be cleared for the next, and clearing pays Landauer on what the record held, conditioned on what is kept.)

The floor can be driven low by design, along two routes. One routes failures into variables already watched for other reasons: a server room's temperature sensor serves the air-conditioning loop and doubles as the fire alarm — one watched variable, two masters; crash-only software converts subtle corruption into a process exit the operating system watches anyway. The other buys a cheap thing engineered to fail *first*: the miner's canary is a purchased detector whose whole design is early failure — a small standing cost that converts an expensive rare check into a glance.

The floor can never be driven to zero: run the "I needed that watcher anyway" regress to its end and a remnant always survives — the marginal price of telling this damage apart from all innocence — and the regress is finite and constructive, so the remnant is computable, made by mode, not merely asserted.

In plain terms. Fires do not ring bells; bought bells ring. Every alarm in your life is a salary someone pays — and clever engineering mostly consists of making failures trip over wires that were already there for other reasons. The cleverness has a floor: some sliver of pure watching can never be economized away, because *telling bad from normal* is itself information, and information is never free. The smoke detector's battery is the schema's seventh face, hanging from your ceiling.

7.3 Recovery: three answers to rot

A corpus keeps content alive by one of three strategies, and the storage rule of §6.5 chooses among them.

Maintenance — watch and re-copy as things spoil — pays a wage forever: detection (§7.2) plus repair, per unit content, per unit time. It wins where rot is fast and the content is in active use, which is why living research fields do their maintenance under the name *review* without noticing it is archival work.

Parity armor — store error-correcting redundancy beside the content — converts repair into decoding: with adequate codes the true value is implied by the survivors, so the restorative erasure of corruption is conditioned on redundancy that pins it, $H(\text{true} \mid \text{redundancy}) \approx 0$, and the repair is **Landauer-free** — pure processing, no fare. The armor must locate damage as well as fix it, so detection hides inside the code overhead too; the engineering trade is armor mass against wage, and it favors armor wherever content is stable and horizons are long.

The deeper point is conceptual: **error correction is the technology that makes repair a different operation from re-derivation** — recovering the proof of a great theorem from parity is decoding; recovering it from nothing requires the theorem's author, or an equal, to apply the same effort anew.

Re-derivation — keep the incompressible core, rebuild the implied rest at need — pays nothing until it pays everything: cheap shelves, expensive emergencies, priced by the recreation face and rational exactly where the storage rule says so. Its silent prerequisite is that the *core stays intact* — re-derivation from a rotted core is the one failure with no local remedy, which is why cores get the armor first.

Real archives mix the three, and the mixture’s optimum under correlated damage — rot that strikes structure rather than bits — remains an open cost-allocation problem we flag rather than solve.

In plain terms. Against the silent vandal, three defenses: hire menders and patrol forever; shelve clever spare sections beside every book so that any torn page rebuilds itself from them; or keep only the seeds and trust yourself to regrow the garden when needed. Patrols for the busy wing, spare parts for the treasures, seeds for the vast derivable rest — and armor on the seed vault before anything else, because seeds that rot cannot be regrown from ambition.

7.4 One dial, two hands

The detection remnant of §7.2 is a single dial that different hands turn in opposite directions. The **engineer** turns it down at design time: canaries, crash-only structure, failures routed into watched variables — spending once to make a lifetime of watching cheap. The **adversary** turns it up at run time: shaping behavior so that everything a monitor observes is equally consistent with innocence — raising $H(\text{truth} \mid \text{sensors})$ so that every check costs the watcher more.

Concealment is anti-canary engineering; surveillance economics is a tug-of-war over one conditional entropy, and the formula does not know whose hand is on it. This paper prices the dial and stops; the contests fought over it — monitors and the monitored, wardens and the watched — lie beyond this paper, with the contests over control, which pick the dial up from the other side.

In plain terms. The same knob that lets a good engineer make failures cheap to spot lets a careful adversary make truths expensive to find. One number — how much what-you-watch says about what-you-fear — and two players pulling opposite ways. We have priced the rope; the match itself is not this paper’s to referee.

8 Reserves: the stock and the catalog

Section 7 kept individual items alive through time. This section keeps the *subject* alive through the gaps in its feed — §4’s law meeting §5’s tranches, with the storage machinery of §§6.5–7 as the means. Development runs on resource tranches (§5.3), which arrive irregularly. On a plot of the resource inflow level, between one tranche’s exhaustion and the next one’s unlocking lies a **valley** — an interval in which the inflow might fall below what the corpus’s operation requires.

§5 showed that resources and knowledge interconvert, so a corpus’s operating requirements can be expressed in either. This paper is epistemological, and its natural commodity is knowledge — specifically diversity — so valleys and reserves are accounted in it.

Such valleys are survived on reserves — a **reserve stock**, distinct from the corpus’s total holdings K — and this section prices the surviving.

8.1 Valleys are statistics

In the simplest model — tranche arrivals as a memoryless stream at rate λ , shortfall deepening at rate ζ during a gap — a valley of gap length T has **depth** ζT and **area** $\zeta T^2/2$, and the worst valley over any horizon has closed-form tail statistics. The area exceeded once per N valleys is

$$a(N) = \frac{\zeta (\ln N)^2}{2\lambda^2},$$

so the provisioning table writes itself; at $\lambda = 1$ tranche per period and $\zeta = 0.8$, verified against simulation to four digits:

Return period	Valley area to provision for
1 in 10	2.1
1 in 100	8.5
1 in 10,000	33.9

The tail is stretched-exponential in the area — fatter than exponential, thinner than power-law — with the design consequence built into its logarithm: **each additional order of magnitude of safety costs only quadratically-growing reserves**. Rare-valley insurance is cheap relative to the ruin it prevents, which is why prudent corpora carry it and why its absence marks either poverty or a short planning horizon. Real arrival processes are not memoryless and real depletion is not constant; the model is a calculable base case whose structure — a tail, a quantile, a provisioning rule — survives realistic elaboration.

8.2 The pricing theorem

What determines the reserve? The answer splits by which regime the valley threatens (§4's floor deciding). Above the subsistence floor, shortfall merely slows development, losses are proportional, and the optimal stock is the classical newsvendor quantile: hold Q^* at the critical fractile of the worst-valley area distribution, where marginal holding cost equals marginal expected shortfall. At the floor, a breach is an existence event (§4.1), losses are not proportional, and the stock is set instead by a chance constraint — provision to the area quantile of the survival probability you choose, a policy number, not a derivable one. The two regimes use the same table above with different rows and different reasoning: degradation prices by expectation; extinction prices by quantile.

In plain terms. The hearth of §4 must eat in winter too. Stocking a pantry against famine and stocking it against inconvenience are different mathematics. Against inconvenience, buy insurance like a merchant: weigh the cost of the shelf against the average cost of running short. Against famine, buy it like a species: pick how sure of surviving you insist on being, read the table, and pay — there is no "average" worth computing across futures in which you do not exist.

8.3 The release toll

Reserves are a stock; surviving a valley needs a *flow*. Can the stored diversity be released fast enough? That depends on the maximal release rate: how fast the needed content can be *found*, *fetched*, and *decoded*. These processes have a price, paid in §3's currency of doing — **the release toll**.

One capacity, first, is already guaranteed and can be set aside. The corpus's operation needs intake X (§4.1); the valley delivers $Y < X$; the stores must therefore supply $Z = X - Y$ per unit time — at most X , since it is a shortfall in an intake of size X . And the integration bandwidth was sized in §4.1 at $B \geq X$: the mouth built for the full intake swallows any partial replacement of it. Digesting released content, then, never binds. Nor does reading: integrating any item requires

reading at least one stored item — the link-targets must be examined — so integration capacity B guarantees read capacity of at least B as well. What remains unguaranteed is narrower, and colder.

Everyday reading works the *warm* set: link-targets the integrator already knows of, fetched live, stored uncompressed. Valley release pulls from the *cold* archive — that is what reserves are: content stored at a distance, unintegrated, compressed by the storage rule, findable only through the cold store’s own catalog. The cold machinery — deep catalog lookups, slow-media fetches, decoding — is sized by its everyday duty, which is occasional; a deep valley might demand cold release at rates that its everyday traffic never approached. Capacity built for that peak and idle otherwise is a real cost, incurred because valleys exist, and it belongs on the release tab with the rest of the toll.

Directly searching a storage of any size costs on the order of the storage itself: without structure, finding means reading. Every real archive therefore builds an **index** — libraries their catalogs, databases their indexes — and the toll’s anatomy runs through it. Per item released, the toll is: *find* (a catalog lookup — roughly logarithmic in the archive’s size with a good catalog, roughly linear without one), plus *fetch* (the medium’s latency), plus *decode* (undoing whatever compression the storage rule chose).

And the catalog keeps books of its own, in three lines. It is *built* at deposit time: each stored item pays a small indexing fee — keys, summaries, cross-references written — so that finding it later costs a lookup instead of a search. It is *maintained* forever: the catalog is itself a corpus (§2.1’s one ontology — it is exactly §6.5’s compressed remainder: names, authors, shelves), so it rots and repairs like everything else in §7. And it is *consulted* at need, which is the payoff.

A catalog is pre-paid finding: index every deposit so that need-time costs log instead of linear — worth building exactly when expected retrievals times the search savings exceed the indexing fees plus the catalog’s upkeep, which the reader will recognize as the storage rule’s sibling, written for finding instead of keeping.

Verification and repair, note, are not release tolls: they are general quality-control lines, allocated wherever error-rate times stakes is largest, already priced in §§6.5 and 7; the release-specific residue is retrieval and decoding.

The consequence deserves its own sentence: **in a shallow valley, the binding factor is the size of the reserve stock; in a deep valley, it is the speed of its catalog.** A vast archive with poor indexing starves its owner in front of full shelves — the insurance was paid, the payout throughput was not.

The two instruments of reserve design are therefore the **stock** (insuring the valley’s area, per §8.2) and the **catalog** (insuring the release rate against the valley’s depth) — and the pair echoes the paper’s oldest duality: an integral quantity and a rate quantity, a loan-shaped instrument and a flow-shaped one, the same two-column ledger reappearing one level up. Institutions that hold knowledge reserves know this: librarianship’s obsessive cataloging is not tidiness. **Retrieval throughput is the insurance’s payout rate.**

A closing remark connects reserves to boundaries. A corpus may hold its reserves at a distance — a specialized store, integrated loosely, maintained cheaply *because* unintegrated, having a separate catalog or none at all — and the arrangement recurs at every scale, down to the archive inside a single mind. What the reserve-holder’s optimal distance is, and what the holding does to the family of corpora around it, are questions of strategy among corpora, beyond this paper; here the reserve enters the books only as what it is locally: a stock, a catalog, and two bills.

In plain terms. Lifeboats have departed sinking ships without taking provisions because no one could find the biscuit under the cannonballs. When the lean years come, three things save you: having stored enough (the pantry), knowing where everything is (the catalog), and being able to get it out fast enough (the cellar stairs). The pantry we can size with a table — we just did; the stairs were sized in §4; the catalog — the humblest of the three — turns out to be what actually decides whether we starve in the worst winters.

Librarians were never being fussy. They were drawing the map that says where the biscuit is — before the ship starts sinking.

9 Boundaries and stranding

Bodies overlap, and §2.4 promised that the overlap matters. This section pays the promise: what happens to the prices when content, correlations, and commands cross the line where one body ends and another begins. We state three results and defer the wars fought over them.

9.1 Stranded correlations

Recall two facts already on the table. Correlation is the asset (Law 2): what a corpus's content is *worth* runs through what it is reliably correlated with. And to compare, combine, or verify two items together, both must lie within one body — that often requires joint reach (§2.4).

Let an item in body A be correlated with an item in body B — two halves of one regularity. While some (likely shared) body spans both, the correlation is an operable asset. Now let the boundary harden: an axis closes — permission revoked, access severed, awareness lost. The correlation does not vanish; correlations are facts about the world, not about reach. But no remaining body can operate it jointly. The asset still exists and no one can spend it. We call it **stranded**.

Stranding is not destruction, and the difference is in the price too. Destruction pays a fare and ends the story. Stranding costs nothing at the moment the boundary closes — no process executed on the content, no fare (Law 1) — and instead converts the asset's value into a standing option: whoever re-opens the boundary, by whatever axis it closed on, recovers the joint asset at the re-opening's price. Late sharing, per Law 1, is a new acquisition at today's conditional entropies — the option's strike price drifts with the world.

In plain terms. Tear a treasure map in two and give the halves to enemies. Nothing was burned: every line of ink survives. But the treasure is findable by no one — the *knowledge* lived in the halves' being readable together, and that thing, which belongs to neither half, is what the tearing disabled. It waits, whole, in pieces — for as long as the halves last, for anyone who can bring them back together, at whatever bringing-together will cost by then. And if someone burns one of the halves, the burning is a real erasure and pays its real fare — stranding was free; ending the wait is not.

9.2 The commandable set

Who pays what to operate *inside another's body*? The question sounds political; the frame principle makes it arithmetic.

Walk the chain. Every operation is executed by some concrete machinery (Law 1). Its cost is evaluated in the executor's frame — the executor's holdings and reach (Law 4). A controller, therefore, never pays "its own" price for a cross-boundary operation; it pays the price *of whichever executor it directs*. Define a controller's **commandable set**: all the executors —

devices, processes, people — that it can actually direct, across all the bodies its axes reach into. The controller’s minimal price for any operation is then the **minimum, over its commandable set, of the frame-priced cost**.

The consequence reverses a naive picture. Reaching into another body looks cheap for the powerful — the warden erases the prisoner’s cache at a word. But the minimum was purchased in advance, and the purchase is enumerable: *surveillance*, to know what is where (detection, §7.2, paid per the remnant); *access*, to hold the command channels open (standing costs on the permission and physical axes); and *maintenance of the executors themselves* — the guards are salaried, the informants cultivated, the software agents kept patched. The word is cheap because everything before the word was not. Cheap coercion is pre-paid coercion; the discount never existed.

In plain terms. A king snaps his fingers and a library across the border burns. What did the burning cost him — a finger-snap? Count what the snap presupposes: spies who told him which library, envoys who keep the channel to the arsonist open, and the arsonist’s wages, paid whether or not anything ever burns. The finger-snap is the last and smallest installment of a long-running bill. Power’s famous cheapness is a bookkeeping that prices a bouquet at the delivery boy’s salary.

9.3 Reach is not absorption

One boundary confusion deserves its own paragraph, because whole arguments have been built on it. A controller whose commandable set penetrates a body deeply — reading its stores, directing its processes, erasing at will — has *included* that body in its own, on the axes it commands. It has not thereby absorbed the body’s **subject**. Generation, adjudication, and integration are capabilities of the commanded corpus, not properties of the command channel; they continue to run — or starve — by §4’s law, on their own feed, under their own criteria. Commanding a library, down to every shelf, makes none of its librarians think your thoughts.

What command *can* do to a subject is starve it (possibly to death), redirect its feed, or tax its every transaction — which is a different set of operations, with its own ledger, beyond this paper.

Groups pay a price for having boundaries at all, and it has a two-line definition: a family of corpora’s **coordination inefficiency** is its realized dissipation minus what perfect coordination would have dissipated for the same result. Stranding is this inefficiency’s permanent form; every toll of §6.4 paid at an internal boundary is its transactional form. The measure exists at every scale that has parts; its uses lie with the same contests, beyond this paper.

In plain terms. You can chain the librarian, read her mail, and burn her drafts — her thoughts remain hers, running on her fuel, until you starve them into ceasing, that is, until you kill her. Control reaches bodies; subjects can only be fed or famished. And any team, any federation, any empire pays a quiet tax simply for being *made of parts*: everything its pieces cannot do jointly that one seamless mind could have. That tax has a name now.

10 What would prove us wrong

A price list earns trust by naming its own audits. We list the paper’s exposures in rank order: each entry states the claim at risk, where the data lives, and what result kills what.

The Landauer wall itself. Any demonstrated erasure below $kT \ln 2$ per bit, at any scale, falsifies the fare and this paper with it. The bound has been probed and held [Bérut et al. 2012]; the exposure stands permanently open, as it should.

The Koomey anchor. Computations per kilowatt-hour have doubled on a stable trend for decades [Koomey et al. 2011]. Our framing requires that trend to *saturate at the Landauer wall on the erasure side* while reversible operations may continue past it. A trend that crosses the wall without reversible engineering falsifies the two-currency reading; a saturation at the wall is the reading's public confirmation, on a time series anyone can extend.

The scrubbing curve. The silent-rot model (§7.1) fixes the functional form by which silent errors accumulate between verifications. Storage-reliability studies of latent errors and scrubbing intervals publish exactly these curves; a measured accumulation shape incompatible with silent-and-cumulative — for example, damage that reliably self-announces — falsifies §7's foundation.

The storage rule against real archives. §6.5 predicts that keep-versus-discard decisions in mature archives track need-probability times recreation cost against maintenance: library deaccession records and code-repository pruning histories are data. Systematic keeping of cheap-to-recreate, rarely-needed content — or discarding of Perelman-priced content — by *competitively disciplined* archives would tell against the rule's optimality reading (the rule as floor survives; §11 grades the difference).

The apprenticeship premium. §6.4 predicts co-location investment tracks field codifiability: disciplines whose content resists codification should show larger and more persistent co-location premiums. Economics of science provides both proxies; a null or inverted correlation damages the tacit-core economics.

The feed law's counterexample. §4.3 stated it: one documented subject that thrives unfed, or one stock whose decay improves its subject, breaks §4.

The consistency battery of the growth demonstration — the alternation signature, the regime diagnostics, the envelope mechanism — is stated in §5.4 and tested where its machinery will be derived — beyond this paper; we list it here so the reader sees the exposures are joined.

In plain terms. Here is our neck, and here are six ways to swing at it. Physics can swing (beat the erasure toll); engineering history can swing (watch where the efficiency curve bends); storage data, archive records, and labor economics can each swing; and a single well-documented counterexample to "sources must eat" ends the show. We do not expect to be hit. We built the scaffold anyway, because a price list that cannot be audited is a rumor.

11 The honesty ledger

A price list earns trust a second way: by stating exactly what does *not* rest on it. This section is the paper's declared epistemology, and it opens with the grading that governs everything above.

11.1 Three tiers of claim

Every substantive claim in this paper lives on one of three tiers, and the tier was stated where the claim was made.

Tier one — expressed. The two floors and their conditional refinement are established physics, cited and experimentally supported. The conditional floor schema and the four transactional laws are our assembly: bookkeeping made explicit, instantiating that physics act by

act. A reader who rejects everything else in this paper still owes these tiers whatever they owe thermodynamics.

Tier two — connected. The structural results — the storage rule, the two instruments, the stranding of correlations, the commandable-set principle, the pricing of the famous dispute — are derivations that stand on tier one plus standard theory (inventory theory, coding theory) plus stated modeling choices. They inherit the floors’ certainty only through those choices, each of which was named where it was made.

Tier three — not yet. Dynamics. Nothing in this paper derives how corpora *grow*, *compete*, or *collapse* from the floors: **the floors bound; they do not drive**. A reader who has taken any sentence here as a theory of history has taken more than was sold; the papers that attempt the dynamics are announced in §12 and will be graded on their own ledgers.

11.2 Two label distinctions the reader should keep

Derived, inherited, candidate — the provenance labels of §1.3 — ran through the text; we re-flag only the two places where mislabeling would cost the most. First, the coinages (the Kuhn fee, Kurzweil and Modis modes, the tacit core’s interfaces) are **names**, offered for convenience and forfeit to any prior art.

Second, and more importantly: several results are **normative, not predictive** — the recovery mixture of §7.3 (armor mass against maintenance wage, chosen by price) and the optimality reading of the storage rule tell engineers what is *rational*, and observing real systems tests whether real agents optimize, not whether the floors exist. We call this the optimality escape, and we have kept it out of §10’s exposures deliberately: a claim that survives every observation by blaming the observed is not an exposure but a costume.

11.3 The open seams

Stated once, one line each, so no reader mistakes silence for closure. The grading of subjecthood — how much generation, adjudication, and integration make a subject — was deferred at §2.3. The order in which a starving subject’s capabilities fail was deferred at §4.3, with ossification flagged as a candidate symptom. The allocation of the three recoveries under *correlated* damage was flagged open at §7.3. The dynamics of the dial under sustained contest — monitor against concealer, over time — were exported whole at §7.4. And the measurement programs behind §10’s exposures are named, not executed: naming your audit is not passing it.

11.4 This paper, priced by itself

Two of this paper’s own limits are consequences of its own results, and honesty enjoys the symmetry. Nearly every statement above decomposes further — every floor into mechanisms, every rule into cases — and each decomposition would bring the text closer to reality.

We stopped where we stopped because a paper is a transfer (§6.4), a transfer is priced by what the receiver must reconstruct, and there is a limit to the size of corpus a reader can integrate at all — a limit with its own factors, untreated here. A foundational text that exceeded its readers’ reception would have paid its printing costs to strand its own correlations. We have tried to stop before that line; where we misjudged it, the failure mode is at least one the paper itself names.

In plain terms. An honest shopkeeper’s price list tells you what things cost. It does not tell you what people will buy, who will get rich, or which shops will burn — those take theories of the customers, and we have only priced the goods. And the list itself obeys its own economics: we could have itemized forever, and stopped instead where we judged a reader’s patience — like every other stock in this paper — actually runs out.

12 The open programme

This paper fenced more than it farmed, deliberately. The fields it leaves open — for our future work, and for anyone’s:

The dynamics of development. The shapes and regimes of knowledge-driven growth; the envelope mechanism and its historical fingerprint; the Kurzweil–Modis demonstration at full depth, with the net-unlock criterion made quantitative; the anatomy of unlocking and of the unlocked resource. It inherits §5’s deferrals and owes §10 its test battery.

The contests between corpora. The dial under adversarial pressure; the commandable set as a weapon, and its full ledgers; what suppression costs the suppressor. It inherits §7.4 and §9’s deferrals.

The internal mechanics of a corpus. The grading of subjecthood; the micro-operations and their prices; links as first-class objects; the order of starvation. It inherits §2’s and §4’s deferrals.

The social layer. Ownership as an option on future reach; permission’s deterrent half; claims and prognoses wrapped around the subject-free machine. It inherits §2.4’s fence.

The reception limit. The maximum integrable corpus, its factors, and what it constrains — including, reflexively, papers.

In plain terms. We priced the goods. How economies of these goods grow, how they fight, how a single shop works on the inside, who holds the receipts, and how much of any of it one mind can take in — open fields, fenced and waiting.

And the reader who walked the plain terms road has been walking one town all along. With hearths that must be fed, a market with two prophets arguing over pantries, a prospector at his pan, mapmakers pricing maps by the trip, a bakery whose secret walks home each night with the baker, a library with a merciful catalog and a cold cellar below, a harbor with provisioned boats, a wall where a torn map waits in two pockets, and an honest shopkeeper who prices his own price list. Anyone is welcome to make these better, richer, more colorful.

And a town always has more streets than even the greatest builder can walk — they are anyone’s to build.

Coda

We asked, at the start, what it costs — in the world’s own units — to know, to keep knowing, and to know together. Ten sections of prices later, the honest answer ends where every price list ends: it informs the builders, but does not choose the buildings.

But still, the town of this paper is built of knowledge. It does not deplete — it compounds. Adding to it enriches all who share it, at little cost; destroying is what truly costs, and the destroyer pays more than anyone else.

And its life is our dreams. They are us: as long as they live, so do we.

Acknowledgments

This paper was developed in extensive collaboration with Claude (Anthropic): formalizing the author’s arguments, assembling and checking the inherited physics, drafting both registers of the text, verifying quantitative claims by independent computation, adversarially attacking the central schema, and verifying the references. The collaboration was conducted under recorded working rules, as colleagues; its process is part of the paper’s record. All claims and their meaning were verified by the author, who bears sole responsibility for the content.

References

- Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., & Lutz, E. (2012). Experimental verification of Landauer’s principle linking information and thermodynamics. *Nature* 483, 187–189.
- Borges, J. L. (1942). Funes el memorioso. In *Ficciones* (1944); English translation "Funes the Memorious" in *Labyrinths* (1962).
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis* 58(1), 7–19.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* 35(1), 128–152.
- Connes, A., & Rovelli, C. (1994). Von Neumann algebra automorphisms and time–thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity* 11(12), 2899–2917.
- Deffner, S., & Campbell, S. (2017). Quantum speed limits: from Heisenberg’s uncertainty principle to optimal quantum control. *Journal of Physics A: Mathematical and Theoretical* 50, 453001.
- del Rio, L., Åberg, J., Renner, R., Dahlsten, O., & Vedral, V. (2011). The thermodynamic meaning of negative entropy. *Nature* 474, 61–63.
- Gibbons, G. W., & Hawking, S. W. (1977). Cosmological event horizons, thermodynamics, and particle creation. *Physical Review D* 15(10), 2738–2751.
- Hartnoll, S. A., & Mackenzie, A. P. (2022). Colloquium: Planckian dissipation in metals. *Reviews of Modern Physics* 94, 041002.
- Hong, J., Lambson, B., Dhuey, S., & Bokor, J. (2016). Experimental test of Landauer’s principle in single-bit operations on nanomagnetic memory bits. *Science Advances* 2, e1501492.
- Kolchinsky, A., & Wolpert, D. H. (2018). Semantic information, autonomous agency and non-equilibrium statistical physics. *Interface Focus* 8, 20180041.
- Koomey, J., Berard, S., Sanchez, M., & Wong, H. (2011). Implications of Historical Trends in the Electrical Efficiency of Computing. *IEEE Annals of the History of Computing* 33(3), 46–54.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Kurzweil, R. (2005). *The Singularity Is Near: When Humans Transcend Biology*. Viking.

- Landauer, R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* 5(3), 183–191.
- Lloyd, S. (2000). Ultimate physical limits to computation. *Nature* 406, 1047–1054.
- Margolus, N., & Levitin, L. B. (1998). The maximum speed of dynamical evolution. *Physica D* 120, 188–195.
- Modis, T. (2002). Forecasting the growth of complexity and change. *Technological Forecasting and Social Change* 69(4), 377–404.
- Parrondo, J. M. R., Horowitz, J. M., & Sagawa, T. (2015). Thermodynamics of information. *Nature Physics* 11, 131–139.
- Polanyi, M. (1966). *The Tacit Dimension*. Doubleday.
- Sagawa, T., & Ueda, M. (2010). Generalized Jarzynski Equality under Nonequilibrium Feedback Control. *Physical Review Letters* 104, 090602.
- Slepian, D., & Wolf, J. K. (1973). Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory* 19(4), 471–480.
- Stremsky, A. (2026). Why the Winner Starves: Epistemic Starvation in Gap-Defending Knowledge Hierarchies. Companion manuscript (arXiv identifier pending).
- Zaanen, J. (2004). Superconductivity: Why the temperature is high. *Nature* 430, 512–513.