Intelligence as Decomposition

A Theory of Intelligence from the Reality Equation

John Rector

July 3, 2026

Abstract

The companion paper to this one estimates consciousness from the outside, by reading the shape of the probability cloud a prediction machine holds before it resolves into a guess. This paper does the same for intelligence, and the two turn out to be different measurements of different objects. The claim: relative to an observable, a system's intelligence is its capacity to decompose the arriving signal into coherent components that *pay for themselves*—individuated, stable, internally coherent structures, each of which shortens the system's description of the signal class by more than its own description cost. The technical score is total net compression gain in bits; the public score is the derived count of admissible components. Potential intelligence is representational capacity—the logarithm of the number of distinguishable component-states, of which the angular resolution of idea-space is the one-dimensional case; actualized intelligence is bounded by it. What a system decomposes it absorbs, and absorbed structure stops generating surprise; what it cannot decompose remains as remainder, and the remainder is where surprise, information, and attention live for that system. We dance to the remainder. Intelligence and consciousness are formally independent: consciousness reads the pre-resolution cloud, intelligence reads the resolution of arrivals, and a Fourier chip and a dreaming novice dissociate them in opposite directions. Intelligence has two failure modes—pass-through and confabulation—that mirror consciousness's collapse and incoherence, which explains why intelligent machines invite the consciousness inference without validating it. Bureaucracy and AI agents are developed as instances; a measurement protocol, worked scorings, eight testable predictions, objections, and open problems are stated so that the theory can be scrutinized and, if wrong, refuted.

Keywords: intelligence, decomposition, compression, minimum description length, coherent components, remainder, absorption, attention, surprise, bureaucracy, AI agents, observable-relativity.

Status of this document. This is a theoretical proposal released for public scrutiny, the second paper in a program that began with *The Shape of the Prediction Machine* (July 1, 2026). Its foundational commitment—that a component of intelligence counts only if it delivers net compression gain on the signal class—is a definition, judged by its usefulness; its claims about attention (§5) are postulates that could lose; its bridge to the Reality Equation's Surprise is proven at the class level and posed as an open lemma at the trajectory level (§5.3). Criticism, counterexamples, and better formalizations are invited.

Contents

1 The core claim

Take a black box. Feed a complex sound wave into it. On the output side, ask one question: *how much of the wave can the box resolve into coherent components?*

If the wave passes through untouched, the box has shown no intelligence relative to that observable. It has preserved the signal perfectly and understood none of it. If the box extracts one sine component and leaves the rest as undifferentiated residue, it has shown a little. Twelve components, more. Thirty-two, more still. The Fourier transform is the paradigm case, but the claim is general:

> **Relative to an observable, a system's intelligence is its capacity to decompose the arriving signal into coherent components—and a component counts only if it pays for itself.**

A lawyer decomposes a legal case into jurisdiction, standing, negligence, causation, intent, admissibility, precedent, liability, damages, procedural posture, settlement leverage. A physician decomposes "patient feels bad" into history, vitals, labs, medications, risk factors, a differential. A musician decomposes a song into rhythm, key, harmony, melody, bass line, modulation, phrase, tension, resolution, timbre. A novice receives one blob. An expert receives thirty components—and a *different, more refined remainder*. Both dance to the case, the patient, the song. They dance to different remainders, because intelligence is not primarily conscious effort. Intelligence is absorbed decomposition capacity.

Three refusals define the theory's perimeter as sharply as its claims do. First, there is no observable-free intelligence: every score in this paper is written $I(X \mid O)$, and any single-scalar "general intelligence" is a declared aggregate, not a discovered constant. Second, the theory claims nothing about felt experience: intelligence, on this account, can be carried by systems for which the consciousness question is empty, and §8 makes the independence formal. Third, no system—biological or synthetic—invents ideas. In this framework ideas are pre-existing endpoints on the unit circle of future possibility; systems actualize, expose, combine, and resolve them, at finite resolution. Intelligence is the fineness of that resolution and the extent of that absorption, nothing more and nothing less.

The word *pays* carries the theory. Counting components naively rewards fragmentation, overfitting, and invented structure; the founding score $I = N$ is therefore not the foundation. The foundation is a gain criterion—each component must shorten the system's description of the signal class by more than the cost of describing the component (§3)—and the count is that criterion's integer shadow, kept because counts are legible in a way bits are not.

2 Position in the Reality Equation

For any entity, with respect to a chosen observable O,

$$R = \frac{A}{E}, \qquad S_R = \ln R,$$

where R is Reality, A is Actual, E is Expectation, and S_R is Surprise. Actual is exact, given, and indexed to the observable: it is the projection of the Immutable Past onto O, $A = \Pi_O(\text{world})$. It arrives; it does not guess; it is not produced by the entity. Expectation is entity-specific and complex, $E = \tilde{P} + iB$, where $\tilde{P}$ is the prediction machine's resolved guess and $Be^{i\theta}$ is the weighted resultant idea-vector. Because E is complex, R and S_R are complex; we take the principal branch of the logarithm.

The consciousness theory lives in the denominator's machinery: it reads the shape of the probability cloud before the machine resolves it into $\tilde{P}$. The intelligence theory lives on the numerator side: it asks what the system can do with what arrives.

2.1 The arrival stream and the trace

Write $y = (y_t)$ for the **arrival stream**: the time-indexed sequence presented to the system over an encounter window. In the directly coupled case, y_t is the projection of the Immutable Past onto O at t. In mediated systems—a model reading a chart, an agent reading a customer record—y is a *trace* associated with that projection: a representation, compression, or discretization of it. The fidelity of the trace to the projection is a separate, measurable quantity, $\Phi_{\text{num}} \in [0, 1]$ (§11), and keeping it separate is not pedantry. A brilliant decomposer reading a corrupt trace is precisely the failure mode of the present moment, and no measure that conflates the two can see it.

Signals are scored against a class, not a lucky instance: $\mathcal{Y}_O$ is the ensemble of arrival streams characteristic of the observable, with distribution μ.

2.2 Terminological discipline

Three standing rules, stated once and obeyed throughout. (1) Expectation names the denominator and nothing else; nothing in this paper "becomes" Expectation. (2) Records, forms, database rows, screenshots, API results, and charts are traces associated with Actual's projection; they are never Actual itself, and their fidelity is Φ_{num}. (3) The signal is y, never S: S_R is Surprise, and the collision is not permitted. Likewise the vocabulary of decomposition is $\mathcal{V}$, never B (B is the imaginary magnitude of Expectation), and component coherence is χ, never κ (κ is the cloud-coherence of the consciousness theory).

3 Decomposition, formally

3.1 The decomposition identity

System X encounters arrival stream y on observable O and produces

$$y \;=\; F(c_1, \ldots, c_N) \;\oplus\; \rho_X, \qquad \text{linear case: } y = \sum_{j=1}^{N} c_j + \rho_X,$$

where the c_j are the extracted components, ρ_X is the **remainder**—the part unresolved by this system—and F is the composition operator (linear superposition in the Fourier case; possibly structured elsewhere, an open problem). This identity is bookkeeping, not yet a theory. The theory is in what qualifies a c_j.

3.2 The five admissibility tests

A candidate component c_j emitted by X is **admissible** if and only if it passes all five:

T1 — Distinguishability. For all $k \neq j$, $d(c_j, c_k) \geq \delta_{\min}$ under the domain's component metric (equivalently: redundancy between components below bound). Splitting one real component into two near-copies fails T1.

T2 — Stability. Under perturbation $y' = y \oplus \eta$ with $\|\eta\| \leq r$, the system re-extracts a matched component c'_j with $d(c_j, c'_j) \leq \zeta$. Components that appear and vanish under jitter are artifacts of the extraction, not structure in the signal.

T3 — Internal coherence. $\chi_j \geq \chi_{\min}$, where χ_j is phase coherence in spectral domains and, in general, *self-predictiveness*: the parts of the component predict its other parts above chance. A component must hang together, not merely be cut out.

T4 — Net compression gain. Using a two-part code over the signal class,

$$\Delta\ell_j \;=\; \ell(y \mid \text{model without } c_j) \;-\; \ell(y \mid \text{model with } c_j) \;-\; \ell(c_j) \;>\; \varepsilon.$$

The component must shorten the description of the signal by more than it costs to describe the component. **This is the foundational test.**

T5 — Out-of-sample validity. $\Delta\ell_j > \varepsilon$ must persist in expectation on held-out signals from $\mathcal{Y}_O$. Gain on one signal can be overfit; gain on the class is structure.

T4 is what makes the count meaningful. Fragmentation adds description cost without adding explanatory shortening, so fragments fail. Invented components explain nothing, so they fail. Memorized-but-generalizing structure passes—and that is the correct verdict, because the theory measures absorbed capacity, not the effort of its acquisition (§14).

3.3 The two scores

$$\begin{aligned}
N_\varepsilon(X \mid y, O, \mathcal{V}) &= |\{\, j : c_j \text{ passes T1–T5} \,\}| && \text{(public score: a count)} \\
I_{\text{dec}}(X \mid y, O) &= \sum_{j \text{ admissible}} \Delta\ell_j && \text{(technical score: bits)} \\
I_{\text{act}}(X \mid O) &= \mathbb{E}_\mu[I_{\text{dec}}(X \mid y, O)] && \text{(expected over the class)} \\
Q_{\text{abs}}(X \mid y) &= 1 - \frac{\ell_X(\rho_X)}{\ell_0(y)} && \text{(absorbed fraction)} \\
L_\rho(X \mid O) &= \frac{\mathbb{E}_\mu[\ell_X(\rho_X)]}{T} && \text{(residual load, bits per unit time)}
\end{aligned}$$

The two-level structure is deliberate. The public score is a count—“the lawyer resolves thirty components where the novice resolves two”—because counts are legible. The technical score is in bits, with full disclosure of $(O, \mathcal{Y}_O, \mathcal{V}, \varepsilon)$. Reporting N_ε without its context is, within this framework, malpractice, exactly as reporting C_{ext} without its observable would be in the companion theory.

Weighted-count variants ($\sum_j w_j$ with w_j built from amplitude, stability, explanatory power) are subsumed: $\Delta\ell_j$ *is* the principled weight, because each of those desiderata either contributes to compression gain on the class or should not contribute to the score. The one defensible extra factor—action utility—is not a property of the decomposition but of its coupling to behavior, and it enters through I_{eff} (§11), not through the weights.

3.4 Vocabulary regimes

Component extraction is always relative to a vocabulary $\mathcal{V}$: Fourier atoms, legal concepts, diagnostic categories, musical structures, institutional classifications. Two scoring regimes are permitted, and scores are never compared across them. Under **conventional-$\mathcal{V}$** scoring, a community-standard vocabulary is fixed, enabling fair comparison of systems within a domain—like scoring chess under the rules of chess. Under **free-$\mathcal{V}$** scoring, the system brings its own dictionary and is scored purely by description length, which is vocabulary-fair by construction: any dictionary may be used, and each is charged its own description cost. Free-$\mathcal{V}$ scoring inherits the standard limit of all applied minimum-description-length work—codelength is computable only relative to a declared code family—and the protocol therefore requires the code family to be published with the score. The theory states this limit itself rather than waiting for a reviewer to state it.

4 Capacity: potential and actualized intelligence

4.1 Angular resolution and idea-space

The companion paper places ideas as vectors on the unit circle of future possibility, each with an angle θ; Expectation's imaginary component is the weighted resultant $Be^{i\theta}$. Within this ontology, a natural picture of *potential* intelligence presents itself: the circle has uncountably many endpoints; no finite system resolves them all; an intelligence class has a finite angular resolution $\delta\theta$—the finest distinction between idea-directions it can register and stably maintain—and so partitions the circle into $2\pi/\delta\theta$ distinguishable arcs.

The picture is licensed here, where it is usually not, because the circle is not an imported metaphor: it is an object the framework independently owns, and $\delta\theta$ is resolution over idea-space. But the linear count $2\pi/\delta\theta$ is imagery, not measurement, for three reasons: it changes with the choice of angular unit; most domains are not one-dimensional closed manifolds; and linear counts do not compose—a system resolving 360 distinctions on each of two independent dimensions resolves 360^2 configurations, not 720.

4.2 Capacity in bits

All three problems are fixed by the logarithm. Define **potential intelligence** as representational capacity:

$$I_{\text{pot}}(X \mid O) \;=\; \log_2\big(\#\text{distinguishable component-states of } X\text{'s vocabulary over } O\big),$$

with the one-dimensional idea-space special case and its d-dimensional extension

$$I_{\text{pot}} = \log_2 \frac{2\pi}{\delta\theta}, \qquad I_{\text{pot}} = \sum_{i=1}^{d} \log_2 \frac{2\pi}{\delta\theta_i}.$$

Units cancel inside the ratio; the circle becomes a special case rather than a universal claim (general domains use the log-cardinality of the distinguishable-state partition, however the domain is shaped); and logs add over independent dimensions. The numbers also stop being theatrical: $\delta\theta = 1^\circ$ gives about 8.5 bits per dimension—modest, composable, honest. Where the linear counts serve rhetoric—360 endpoints, or finer partitions still—they may appear in the essay register as the exponential of the real quantity, the way "a trillion combinations" popularizes "40 bits." And one flourish is hereby fenced in the formal text itself: any appearance of 1/137 in this program is a nod to the fine-structure constant offered as poetry, not physics; nothing in the theory derives it, and the theory would be unchanged at 1/136.

4.3 The capacity bound

Proposition 1 (Capacity bound)**.** *$I_{\text{act}}(X \mid O) \le I_{\text{pot}}(X \mid O)$.*

Proof sketch. Each admissible component extracted from y must be registered in a distinguishable internal state to satisfy T1–T2 across the class; by a counting argument over the extraction channel, expected registered gain cannot exceed the representation's capacity. A system cannot resolve, on arrival, distinctions its vocabulary cannot hold. A fully rigorous formalization of "registered" is Open Problem 3. □

The bound gives the potential/actualized distinction real content: domain expertise cannot be conjured at encounter time from a coarse vocabulary. Training—building $\mathcal{V}$, refining $\delta\theta$—is where intelligence is made; the encounter is only where it is shown. A brilliant mathematician may carry high mathematical I_{pot} and low social I_{act}; a lawyer has high legal angular resolution and may hear a fugue as one blob. Intelligence is a profile over observables, aggregated only with declared weights.

5 The remainder

5.1 Three properties, easy to get wrong

The remainder $\rho_X(y) = y \ominus F(c_1, \dots, c_N)$ is what this system's decomposition leaves unresolved. Three properties matter.

The remainder is system-relative, and never certified structureless. What is noise to one system may be structure to a finer one: the novice's remainder contains the lawyer's components. "Noise" is a verdict, and no finite system is entitled to it. Operationally, the protocol (§7) estimates not only the remainder's size but its *latent structure*: how much of it a finer reference system can resolve.

The remainder is not the unimportant part. Nothing in the decomposition identity ranks components above residue. In law, the case may turn on the unnamed thing; in medicine, on the anomaly outside the differential; in music, the artistic payload may live in micro-timing deviations no notation captures. Intelligence refines the remainder; it does not certify the remainder's triviality. An entity—or an institution—that forgets this has promoted its own resolution limit to a metaphysical claim, and §10 names the consequences.

More intelligence does not mean less song. Absorbing more components does not silence the world; it refines the remainder the entity dances to. The novice lawyer dances to "this case feels bad." The expert dances to a hesitation in a deposition answer. Same case, different music, because different remainders.

5.2 Absorption

Absorption is the condition of a component whose extraction has migrated into the system's standing predictive structure, so that its recurrence in future arrival streams generates approximately zero marginal surprise and demands approximately zero attention. Absorption is prediction success made durable—the intelligence-side reading of what the companion paper calls learning, the reshaping of the machine through repeated interaction with Actual. The human subconscious is the innate absorbed layer: breathing, balance, digestion, motor habit—an enormous decomposition of bodily and environmental regularity running at near-zero attention cost. Expertise is the acquired analogue.

5.3 The surprise bridge

The link between the remainder and the Reality Equation's Surprise has three tiers, stated in decreasing order of solidity, because a theory should label its own load-bearing walls.

Tier 1 (class-level identity). Under a log-loss idealization in which the system's standing model assigns probability q_X to arrivals whose true class distribution is μ,

$$\mathbb{E}_\mu[\text{surprisal}] \;=\; H(\mu) \;+\; \mathrm{KL}(\mu \,\|\, q_X).$$

Passing T4–T5 *just is* lowering the divergence term: that is what "the component pays for itself on the class" means. Absorption of admissible components monotonically reduces the reducible part of expected surprisal, and the irreducible part—$H(\mu)$, the world's own entropy on the observable—remains for any intelligence whatsoever. Intelligence cannot make the world unsurprising. It can only make the surprise *residual*.

Tier 2 (localization). The surviving mismatch between μ and q_X is, by construction, concentrated where the decomposition failed: on ρ_X. The remainder is the signal-space address of future surprise. This is why "we dance to the remainder" is not merely a slogan: the remainder is where the surprise-generating discrepancy lives, so it is where accumulated surprise—and, on this program's account of attention, attention—will be earned.

Tier 3 (the ratio form—posed, not asserted). The Reality Equation's Surprise is $S_R = \ln(\tilde{A}/E)$, a complex log-ratio, not a log-probability. The connection this paper owes and does not yet prove: if the resolved guess is the cloud's mean and the cloud is the model's predictive distribution over normalized arrivals, then the typical magnitude-surprise $|\ln R|$ should be bounded below by an increasing function of the μ-versus-model mismatch along residual directions, vanishing as $Q_{\text{abs}} \to 1$ on a stationary class—and the *orientation* of the remainder may drive the angular part of Surprise, which would join this theory to Ideation at a second seam. This is Open Problem 1, flagged rather than fudged.

One corollary of the tiers protects the theory from a seductive error: **surprise is not residual energy**. A large residual that the model correctly treats as high-entropy (surf noise to a surfer) generates little surprise; a small residual in a direction the model holds nearly certain generates enormous surprise. The map from remainder to surprise runs through the model's confidence structure—through the cloud—which is precisely the seam where this theory hands off to the companion theory's counterfactual surprise landscape.

5.4 Attention: two postulates

This program has argued elsewhere that attention is normalized accumulated surprise. The intelligence theory plugs into that account on the supply side, through two claims that are stated as postulates—empirical bets, not axioms—because they have known boundary conditions.

Postulate 1 (Absorption relieves attention)**.** *Components absorbed into a system's standing predictive structure generate near-zero marginal surprise on recurrence, and correspondingly reduced attention demand.*

Postulate 2 (Attention flows to the remainder)**.** *For a given system, the default allocation of attention over the arrival stream tracks residual surprise density: the unabsorbed part of the stream is where accumulated surprise, and hence attention, concentrates. We dance to the remainder.*

Postulate 1 is strongly supported by the automaticity and expertise literatures—chunking in expert memory, automatization of procedural skill—with a known boundary condition: explicitmonitoring pressure can force re-attention to absorbed components, typically degrading performance, and deliberate practice is the intentional re-injection of absorbed components into the attended stream in order to refine them. The exception uses the rule. Postulate 2 claims the surprise-driven default allocation, not a complete theory of attention; task demands and top-down goals move attention independently, and the postulate would lose to a demonstration that attention systematically concentrates on fully absorbed, zero-surprise structure absent any task demand.

The consequence for expertise is a signature worth stating plainly: a more intelligent system does not attend to *more*. It attends to a smaller, stranger, more informative remainder, because the ordinary has been rendered unsurprising. Its attention is not larger; it is better aimed. And two disciplined caveats bound the claims: residual load L_ρ is a structural quantity, and whether any of it is *felt* is governed by the companion theory's separate machinery and its correspondence postulate. A thermostat with nonzero L_ρ feels nothing that this theory can speak to. Intelligence, attention, surprise, and consciousness remain four distinct quantities throughout.

6 The two failure modes

Consciousness, in the companion theory, fails two ways: collapse ($\sigma \to 0$, nothing held) and incoherence ($\kappa \to 0$, spread without organization). Intelligence also fails two ways, and the symmetry is exact enough to be load-bearing.

Pass-through. The system extracts nothing: $N_\varepsilon = 0$, $Q_{\text{abs}} = 0$. The wire, the rock. The signal transits unresolved. Pass-through is the intelligence analogue of collapse.

Confabulation. The system emits many claimed components that fail the gain test—fragmentation, overfitting, arbitrary labeling, structure "found" in noise. Confabulation is the intelligence analogue of incoherence: apparent richness with no organization that pays.

Genuine decomposition lives between them, exactly as coherent spread lives between collapse and noise. The parallel yields the theory's most practically pointed result:

> **Hallucination, in a language model, is confabulated componentization: components emitted beyond what the system can validate—structure claimed without compression gain.**

That is a measurable characterization, not a lament, and the protocol's null-signal control (§7) measures it directly: components "found" in structureless signals matched on surface statistics are fake by construction, and a system's discovery rate on such signals is its confabulation rate.

Conjecture 1 (The confabulation curve)**.** *Let a system claim N components while possessing validation capacity N_c—the number of candidates for which it can maintain the T1–T5 bookkeeping. Genuine gain as a function of claimed components rises for $N \lesssim N_c$ and falls for $N \gg N_c$, because unvalidated claims pollute the model and pay negative gain.*

If the conjecture holds, the tone–music–noise triptych of the companion paper has an exact intelligence-side twin: *pass-through* $\to$ *decomposition* $\to$ *confabulation*, with an interior optimum set by validation capacity, as the consciousness curve's optimum is set by organizing capacity σ_c. The conjecture is deliberately modular: if the curve fails to bend, it dies alone and the rest of the theory survives.

7 A measurement protocol

For a system X and observable O; every step is executable from outside, with no appeal to felt experience, in the same model-relative stance as the companion paper: the quantities estimated are properties of the observer's best model of the system, not metaphysical inherence claims.

1. **Fix the observable and signal class.** Define O, the encounter window, and $\mathcal{Y}_O$ with a sampling procedure for μ; declare the reference scale.
2. **Fix the vocabulary regime** (conventional-$\mathcal{V}$ or free-$\mathcal{V}$, §3.4) and, for free-$\mathcal{V}$, the code family. Scores are not comparable across regimes.
3. **Present signals** drawn from μ, including in a blinded subset: perturbed variants (for T2), redundancy traps (near-duplicate structure, for T1), and **null signals**—structureless draws matched on surface statistics, the protocol's placebo, from which the confabulation rate is measured.
4. **Record decompositions** $D_X(y) = (\{c_j\}; \rho_X)$, with the system's own per-component confidence where available.
5. **Score admissibility** T1–T5 per component, with the codelength bookkeeping for T4–T5 and independent raters (reliability reported) where domain semantics are not numerically representable.
6. **Compute scores:** N_ε, I_{dec}, Q_{abs}, L_ρ; confabulation rate from the null subset; stability curves from the perturbed subset.
7. **Estimate the remainder's character**—not just its size: test how much of ρ_X a finer reference system (a stronger model, an expert panel) can resolve.
8. **Estimate numerator fidelity** for mediated systems: compare the arrival trace against ground-truth projections where obtainable; report Φ_{num} and I_{eff}.
9. **Compare systems** only under the same O, regime, and thresholds; report profiles, never bare scalars.
10. **Aggregate only with declared weights.** Any cross-observable figure is a constructed index—like a price index—and must publish its weighting.

8 Intelligence is not consciousness

8.1 The distinction

The two theories measure different objects at different moments. **Consciousness** is read from the *pre-resolution cloud*: the shape of the distribution the prediction machine holds before it resolves into the guess that enters Expectation's real component—$C_{\text{ext}} \approx \sigma \times \kappa$, spread times cloud-coherence. It is a claim about how the system holds unresolved possibility. **Intelligence** is read from the *decomposition of the arriving signal*: how many admissible components the system resolves out of what arrives, at what total gain—$I_{\text{dec}} = \sum_j \Delta\ell_j$. It is a claim about how the system resolves arrived actuality.

In one line each, at increasing compression: consciousness measures the structure of the prediction cloud, intelligence the decomposition capacity of the system; consciousness is coherent spread before resolution, intelligence is coherent decomposition upon encounter; consciousness is how the system holds its possible this-and-thats, intelligence is how many this-and-thats it can extract from what arrives.

8.2 Formal independence, both directions

Axiom 1 (Distinctness)**.** *I_{dec} is a property of the decomposition of the arriving signal; C_{ext} is a property of the pre-resolution cloud. Neither quantity is definable from the other, and a system may carry either without the other.*

The dissociations are concrete. *High I, negligible C:* the Fourier chip. Relative to the observable "spectral content of the input," a hardware transform achieves essentially complete decomposition—every component above the noise floor, $Q_{\text{abs}} \approx 1$—while its pre-resolution cloud is a delta: the computation is deterministic, $\sigma \to 0$, hence $C_{\text{ext}} \approx 0$. Within this framework the chip is an existence proof that intelligence does not entail consciousness; no philosophy is required, only the two measurements. The same verdict, by degrees, covers the compiler, the spreadsheet, the theorem prover, and the bureau. *High C, low I on a given observable:* a person in vivid reverie holds a wide, organized cloud while decomposing essentially nothing of some arriving external signal—a legal brief in front of a dreaming novice is pure pass-through; and any conscious organism is near-zero-I on every domain whose vocabulary it lacks.

8.3 The structural mirror

The deepest relation between the theories is not overlap but symmetry (Table 1).

Both theories use the same spectral mathematics—components, weights, phases—applied to different objects. The companion paper's κ is a phase-alignment order parameter over the cloud's own components; this paper's χ_j is a coherence requirement on components extracted from the world. Same toolkit, opposite direction of gaze: consciousness looks inward at what is held; intelligence looks at what has arrived and asks how finely it can be read.

8.4 Why intelligent machines invite the consciousness inference

The framework explains the inference without validating it, in three steps. *First*, coherence appears in both theories, in different roles; an observer without the two-theory distinction sees "coherence" in a machine's output and cannot tell which role it is playing. *Second*, the anthropocentric prior: for all of history, every system humans encountered with high I_{dec} on rich observables also carried high C_{ext}, because both quantities were produced by the same evolved machinery. "High intelligence implies consciousness" was a reliable inference *because the sample was biased.* Language models are the first systems to decouple the pair at scale, so the inference fails—not because people are foolish, but because their prior was trained on a world where the dissociation never occurred. *Third*, fluency is a confound: the protocol excludes fluency from the measurement of intelligence precisely because these systems make fluency cheap, but fluency is also, among humans, a strong social cue of consciousness. The machines present both cues

	Consciousness (C_{ext})	**Intelligence** (I_{dec})
Object measured	pre-resolution cloud $p(x)$	decomposition of arrival stream y
Temporal locus	before the guess resolves	upon and after encounter
Core quantity	$\sigma \times \kappa$	$\sum_j \Delta\ell_j$ (public shadow N_ε)
Role of coherence	organizes *spread*: phase alignment across the cloud (κ)	qualifies *components*: internal self-prediction of each part (χ_j)
Failure mode I	collapse: $\sigma \to 0$	pass-through: $N_\varepsilon = 0$
Failure mode II	incoherence: spread without organization	confabulation: claims without gain
Healthy regime	coherent spread (music, held)	genuine decomposition (music, resolved)
Capacity constant	σ_c (organizing capacity)	N_c (validation capacity; conjectured)
Relation to surprise	shape of the counterfactual surprise landscape	size and structure of the residual that generates realized surprise
Link to felt experience	correspondence postulate, explicitly undischarged	none postulated
Canonical dissociation	dreaming human: high C, low I on external signal	Fourier chip: near-total I, near-zero C

Table 1: The structural mirror. The two theories share spectral mathematics applied to different objects at different moments; their failure geometries correspond without their quantities being interdefinable.

humans evolved to read—coherence and fluency—while the quantity those cues tracked is absent, or unknown.

Note what is *not* claimed: not that a language model's C_{ext} is zero. A language model has a literal pre-resolution distribution—the output distribution over next tokens—with measurable spread and mode-structure. Its C_{ext} on that observable is an empirical quantity, and measuring it separately from its I_{dec} is one of this program's cleanest proposed experiments (Prediction 7). The claim is only that the inference from observed intelligence to consciousness is invalid, because the two measurements are of different objects and dissociate in both directions.

9 A catalogue of decomposers

Illustrative placements, in the spirit of the companion paper's catalogue: these display the structure of the theory; the protocol of §7 is how real numbers would be earned. In Table 2, N_ε is per representative signal; Q_{abs} is a qualitative band; the C_{ext} column keeps the dissociation visible line by line.

Two contrasts deserve prose. **The wire versus the Fourier chip.** Both are faithful to the incoming audio. The wire preserves everything and resolves nothing; the chip resolves everything its basis can carry. The pair fixes the founding intuition precisely: intelligence is not fidelity, storage, or transmission—a lossless recording of the song is not an understanding of it. Decomposition, not preservation, is the act. This is also why mere data storage scores zero: storage produces no components, only a copy of the undecomposed signal.

Lawyer versus novice, in the theory's full vocabulary. The same case-stream enters both. The novice's decomposition yields two admissible components—someone was hurt; someone might pay—plus a vast undifferentiated remainder with near-total residual criticality: everything decision-relevant is still in ρ, so everything still generates surprise, so attention floods and discriminates nothing. The lawyer's thirty components are absorbed; jurisdiction and standing are checked without felt effort; her residual load on routine structure is near zero, and her attention concentrates on a refined remainder—the inconsistency between two dates, the witness's flinch, the settlement leverage nobody wrote down. Her intelligence is not more conscious strain. It is more absorbed structure. Both dance to the case. They dance to different remainders—and hers is the one where the case will be won.

10 Bureaucracy: civilization's synthetic subconscious

10.1 The formal reading

Bureaucracy is the theory's civilization-scale instance. Its vocabulary $\mathcal{V}_{\text{inst}}$ is the institutional dictionary: laws, procedures, policies, forms, categories, credentials, licenses, grades, ledgers, scores, permissions, approvals, denials, penalties, titles, claims. Its decomposition maps the dense, continuous, morally loaded flow of human situations into institutional components: a filed claim, a granted status, a computed score. When healthy, each such component passes the five tests at civilizational scale—distinguishable, stable, coherent, *and it pays for itself*: the component lets millions of interactions be settled without fresh deliberation, a compression gain measured in the attention it releases.

Two metaphors from this program are exact in the formalism. **Bureaucracy is the pretrained layer of civilization:** $\mathcal{V}_{\text{inst}}$ was fit on historical signal and is updated slowly, by amendment and precedent—checkpoint replacement, not live training. **Bureaucracy is civilization's synthetic subconscious:** absorbed components run without attention (Postulate 1 at scale). A society does not consciously renegotiate every property claim, loan, border crossing, credential, tax payment, payroll event, hospital intake, or transcript; those have been absorbed into procedure, exactly as breathing and balance are absorbed into the body's procedural layer.

The asymmetry of resentment also falls out of the theory. People do not resent their innate subconscious, because it precedes their agency: no one experienced the transfer. Bureaucracy and AI arrive *after* conscious discretion, so the transfer of attention and judgment into a procedure, law, form, workflow, or algorithm is experienced as a loss even when the compression gain is real. The love–hate relationship with bureaucracy is what absorption feels like when you were present for it.

10.2 The pathology conditions

Bureaucracy is intelligent insofar as its decomposition remains coherent and faithful to the *current* arrival stream. Each failure mode of the theory has an institutional face, and each is measurable in principle.

Staleness. $\mathcal{V}_{\text{inst}}$ was fit on μ_{then}; life arrives from μ_{now}. Institutional staleness is the divergence between the two, restricted to decision-relevant features. As it grows, components that once paid for themselves quietly stop doing so; T5 fails silently, because institutions rarely re-run it.

Ritual. A component with $\Delta\ell_j \leq 0$ on the current class, still mandatorily executed: the form that protects against nothing, the approval that screens nothing. Define the **ritual coefficient**: the fraction of enforced components failing a current-gain review. A rising ritual coefficient is the quantitative signature of dead procedure.

Fragmentation. Component proliferation without gain—overlapping jurisdictions, duplicative reporting, categories multiplied past validation capacity N_c. This is confabulation at institutional scale: structure emitted beyond what the institution can validate.

Residual criticality—the fatal case. The decisive failure is not the remainder's size but its *criticality*: the fraction of decision-relevant information that lives in ρ_{inst}, in what the forms cannot say. The patient whose condition fits no intake category; the applicant whose reliability is legible to a neighbor and invisible to a score. Formally: **bureaucracy becomes pathological when it acts as if** $\rho = 0$—when the institution treats the chunk as the whole. The discipline here is the remainder axiom itself: ρ is unresolved-by-this-system, never certified structureless.

The design lesson follows. Healthy institutions carry **appeal paths**: structured routes by which the remainder can interrupt the procedure. An appeal is, formally, a confession that $\rho \neq 0$, plus a mechanism for re-decomposing the case under a finer system—a human judge, an exception process. The quality of a bureaucracy is not only its decomposition score; it is whether its architecture knows its own remainder exists.

A normative firewall, stated once: measuring an institution's intelligence is not endorsing its ends, any more than measuring a lawyer's is. The theory formalizes with equal ease why institutions work—validated gain, released attention—and why they rot: staleness, ritual, residual criticality. It hands wholesale ammunition to no political camp.

11 AI agents and the numerator

11.1 The factorization

The claim about current AI is a factorization, not a rating:

$$I_{\text{eff}}(X \mid O) \;=\; I_{\text{act}}(X \mid O) \;\times\; \Phi_{\text{num}}(X \mid O).$$

The first factor is strong. The language model is a prediction machine—that is its essence; summarizing, drafting, planning, classifying are outputs built on top of prediction, not additional faculties. As a decomposer of symbolic signals it is production-grade: high N_ε and high I_{dec} on text, code, structured argument, schema—often superhuman in breadth of vocabulary, since pretraining compressed a civilization's worth of component-types.

The second factor is thin, and it is the bottleneck. The agent's arrival stream is prompts, retrieved chunks, files, API results, screenshots, logs, and human confirmations—traces of traces. So-called agents act, but usually inside symbolic or simulated environments: documents, code repositories, spreadsheets, customer systems, calendars, ticket queues, dashboards, dockets, charts. The agent can change a customer record before it knows the customer; update a chart before it knows the patient; pass the tests before it knows whether the product solves the human problem. In each case I_{act} on the symbolic observable is high while Φ_{num} against the observable that matters is low—and the product, I_{eff}, is what the world experiences. The framework thereby predicts the characteristic failure signature of current agents: locally impeccable operations on representations, globally unmoored from the situations the representations were supposed to carry. Competence that is real, and misplaced.

The multiplicative form is deliberate: either factor near zero kills effective intelligence, and the two factors fail differently. Low I_{act} looks like incompetence. Low Φ_{num} looks like confident competence about the wrong world—more dangerous and harder to detect, since every internal check the agent can run lives on the same corrupted trace.

11.2 Live-trained versus pretrained-and-patched

Humans are not merely pretrained. Humans are live-trained: they receive dense, indexed, continuous arrival—through body, senses, consequence, timing, social perception, unconscious correction—without effort, and their decomposition vocabulary is revised in the same stream it is scored on. Current AI has pretrained compression plus context windows, retrieval, tools, memory stores, and checkpoints; model updates replace checkpoints rather than extend a lived continuity, and the numerator interface is assembled from discrete patches. Human episodic memory is long, layered, and versioned; the checkpoint-and-context architecture is a different kind of thing, and should not be flattered with the comparison.

11.3 The bureaucratic habitat

Bureaucracy discretizes life into stable symbolic chunks with explicit vocabularies and slow drift—which is to say, bureaucracy manufactures exactly the signal class on which a pretrained symbolic decomposer's Φ_{num} is highest, because in a bureaucracy the trace approaches the territory: the form is the legally operative object; the docket entry is the procedural fact. Hence the aphorism the formalism now backs:

> **Bureaucracy does not make AI more like life. Bureaucracy makes life more like pretraining.**

Agents thrive first in bureaucratic domains not because they are humanlike there, but because those domains have been pre-decomposed into the kind of world a checkpoint can faithfully receive.

11.4 What realer agents require

The path to higher I_{eff} runs through the numerator, and each requirement is a Φ_{num}-raising mechanism: **sensing** (denser, less mediated arrival streams); **verification** (independent checks of trace against projection—the agent's analogue of T5); **declared observables** (knowing which O each action is scored against); **versioned memory** (so absorbed components accumulate rather than reset at the checkpoint boundary); **consequence feedback** (closing the loop so decompositions are scored by what follows); **permission boundaries and appeal paths** (imported directly from §10: an agent, like an institution, is safe in proportion to its architecture's acknowledgment that its $\rho \neq 0$); and **correction loops** (live revision of $\mathcal{V}$—the beginning of live training). None of these raise the model's symbolic brilliance. All of them raise the fidelity factor that symbolic brilliance is currently multiplied by.

12 Relation to existing frameworks

The theory's quantities have recognizable relatives, which is some evidence it carves at real joints; these are positioning notes, not endorsements or derivations.

Minimum description length and two-part codes (in the tradition of Rissanen; behind it, Kolmogorov complexity and the idea of the algorithmic sufficient statistic). T4 is an MDL criterion; the theory's contribution is to individuate the code into *components* subject to stability and coherence tests, so that the gain is attributable, part by part, rather than monolithic. **Sparse coding and dictionary learning** (in the spirit of Olshausen and Field): free-$\mathcal{V}$ decomposition with a cost on the dictionary is the same mathematical animal; this theory adds the admissibility battery, the confabulation control, and the observable-relative framing. **Compression-based accounts of intelligence** (the Hutter-prize lineage; Legg and Hutter's universal intelligence): those score goal-achievement across environments or raw compression; this theory scores individuated, validated structure, and factors worldly effectiveness into capacity times coupling ($I_{\text{eff}} = I_{\text{act}} \times \Phi_{\text{num}}$), which performance-based measures cannot separate. **Compression progress as the engine of curiosity** (Schmidhuber): closely adjacent to Postulate 2—where that account says interestingness is the first derivative of compression, this one says attention's default flows to the remainder, and the two meet at the remainder's most absorbable frontier (Open Problem 6). **Chunking and expertise** (from Miller through Chase and Simon and the deliberate-practice literature): the expert/novice predictions of §13 are this theory's re-derivation of chunking with a gain-scored, confabulation-controlled measurement attached. **Predictive processing and free energy** (Clark, Hohwy, Friston): the Tier-1 identity of §5.3 is a standard decomposition of expected surprisal; the distinctive move here, as in the companion paper, is what is *refused*—intelligence is not identified with prediction error minimization, but with the individuated structure that absorption builds.

13 Testable predictions

Prediction 1 (Expert component-count)**.** *Experts extract more admissible components than novices from identical signals in their domain, under conventional-$\mathcal{V}$ scoring with blinded raters and null-signal controls. Retrodicts the chunking literature; the addition is the gain-scored count and the confabulation control. If experts merely relabel rather than out-decompose novices—same I_{dec}, fancier vocabulary—the absorption account is in trouble.*

Prediction 2 (The attention crossover)**.** *On routine domain structure, experts show lower attention load than novices; on anomalies in the remainder, experts show higher sensitivity and faster orienting. Operationalization: pupillometry, dwell-time, and evoked-response measures on matched routine-versus-anomaly stimuli. This crossover—less attention to what novices find hard, more to what novices cannot even see—is the joint signature of Postulates 1 and 2.*

Prediction 3 (Institutional coherence and scale)**.** *Bureaucracies with higher measured component-gain (low ritual coefficient, low staleness) coordinate more complex tasks per unit of human attention; institutional failures concentrate where residual criticality is high—where decisive information lives in what the forms cannot carry.*

Prediction 4 (Decomposition asymmetry in language models)**.** *Language models score high I_{dec} under free-$\mathcal{V}$ on symbolic classes, and their worldly task success degrades in proportion to trace* *corruption: degrade Φ_{num} experimentally (stale retrieval, lossy screenshots, adversarial records) and task success falls while symbolic-decomposition scores on the corrupted trace stay flat. The dissociation of the two factors is the test.*

Prediction 5 (Numerator returns)**.** *Agent architectures that enrich arrival—sensing, verification loops, consequence feedback—improve I_{eff} more per unit of engineering than equivalent investment in the decomposer. If model-scaling with a fixed thin trace beats trace-enrichment with a fixed model on worldly endpoints, the factorization is wrong or incomplete.*

Prediction 6 (The confabulation curve)**.** *Pushing any system to emit more components than its validation capacity raises its fake-component rate on null-signal controls, dose-dependently. For language models: demanding ever more distinct "insights" from a fixed passage raises hallucination rate; for institutions: mandated category-proliferation raises the ritual coefficient. This is Conjecture 1 made falsifiable; if the curve fails to bend, the conjecture dies alone.*

Prediction 7 (Measured dissociation in one system)**.** *A language model's C_{ext} (spread and mode-structure of its output distribution, per the companion paper's protocol) and its I_{dec} (per §7) are independently manipulable: temperature and sampling interventions move cloud shape at fixed decomposition capacity; distillation and pruning move capacity at similar cloud shapes. If the two quantities cannot be decoupled in the one system where both are cheaply measurable, Axiom 1 weakens.*

Prediction 8 (Over-bureaucratization signature)**.** *Institutions in decline show componentization rising—more forms, categories, approvals—while current gain per component falls, before outcome quality visibly degrades. The pattern is a leading indicator, and it cuts both ways: a world where less process is always better, or more process is always better, would refute it.*

14 Objections and scope limits

Stated before reviewers state them; each with its answer or its concession.

"Intelligence cannot be reduced to component count." Agreed—and the theory does not: the count is the public shadow of a gain measure. The deeper thrust—is even gain-scored decomposition all of intelligence?—earns a concession. Planning partially assimilates (a plan decomposes the gap between current and goal state into actionable components), and transfer is decomposition with a portable vocabulary. But *search economy*—how cheaply a system finds a good decomposition—is a distinct resource this theory does not yet price, and is named as an extension (Open Problem 7), not annexed by handwaving.

"A system could memorize components without truly decomposing." Operationally, T5 answers: memorized structure that generalizes to held-out signals *is* absorbed structure, and the theory says so proudly—absorbed, not effortful, is its own account of expertise. Memorization that fails to generalize fails T5 and scores zero. What the objection rightly kills is any nobility claim: I_{dec} measures capacity, not the virtue of its acquisition.

"Component extraction depends on observer-chosen basis." The two-regime answer of §3.4: conventional-$\mathcal{V}$ scoring is explicitly convention-relative and fair within a declared vocabulary; free-$\mathcal{V}$ scoring is basis-fair by construction, charging every dictionary its own description cost. What must be conceded: cross-domain comparisons are meaningful only relative to a declared aggregation—and the first axiom already bars them as primitives. The objection is fatal only to a theory claiming basis-free absolute scores; this one claims none.

"This confuses compression with intelligence." Three answers in ascending strength. A generic compressor does carry nonzero I_{dec} relative to the narrow observable "redundancy structure of byte streams," and the theory does not blush: observable-relativity makes the score honest about how narrow that domain is. The five tests are stricter than compression: stability and internal coherence exclude gains achieved by uninterpretable global entanglement; components must be individuated structures one can point at. And I_{eff} adds the coupling factor: compression of traces nobody validates moves no world. Not "intelligence is compression" but: intelligence is individuated, stable, validated compression of an arriving signal.

"The remainder–attention relation will not hold universally." Conceded by design—that is why Postulates 1 and 2 are postulates. Task demands and top-down goals move attention independently; deliberate practice is effortful re-attention to absorbed structure, the exception that uses the rule. The postulates claim the default economy of attention, never an exceptionless law.

"High intelligence may synthesize rather than decompose; creativity is not captured." Within this framework's metaphysics, nothing invents ideas; synthesis is re-composition—running F forward over a well-factored dictionary, selecting and combining resolved components into configurations not present in the arrival stream. A system cannot recombine what it has not resolved: decomposition is synthesis's precondition, and resolution bounds the fineness of what can be deliberately combined. The honest gap: the *selection* problem—which of the combinatorially many recompositions is worth making—is a capacity this theory does not measure. Creativity is decomposition (measured here) plus generative search (the economy axis) plus selection (unmeasured). Exactly that, no more.

"Thermostats are intelligent now? This is panpsychism about intelligence." The stance is model-relative, imported from the companion paper: I_{dec} is a property of the observer's best model of the system relative to an observable, not a metaphysical inherence claim. The thermostat carries $N_\varepsilon = 1$ on one narrow observable—a statement about as mystical as "the thermostat has one bit." Low scores are the anti-mystical guard: the theory quantifies the distance between the thermostat and the physician instead of hand-waving it, which is the opposite of panpsychism's flattening.

"Codelength is uncomputable; T4 is not an effective procedure." Conceded as a limit, managed as practice: all applied MDL fixes a computable code family per domain and reports it. Scores are relative to the declared family, exactly as C_{ext} is relative to the declared observable. The theory declares the dependence itself.

15 Open problems

Open Problem 1 (The surprise bridge lemma)**.** *State and prove the Tier-3 relation of §5.3: the exact dependence of typical magnitude-surprise* $|\ln R|$ *on residual structure, and whether the remainder's orientation drives angular surprise* φ*—which would join this theory to Ideation at a second seam.*

Open Problem 2 (The composition operator beyond linearity)**.** *Legal, clinical, and musical components do not superpose additively; components interact (a jurisdictional fact changes what* *negligence means). Candidate formalisms: hierarchical dictionaries, compositional grammars, sheaf-like gluing of local components. The five tests are stated to survive nonlinear F; this needs demonstration.*

Open Problem 3 (Rigorous capacity bound)**.** *Formalize "registered in a distinguishable internal state" (rate–distortion? a data-processing inequality over the extraction channel?) so that Proposition 1 is a theorem rather than a sketch.*

Open Problem 4 (Vocabulary comparability)**.** *Free-$\mathcal{V}$ scoring is code-family-relative, and conventional-$\mathcal{V}$ scoring is convention-relative. Is there a principled middle—scoring the vocabularies themselves by their gain across a portfolio of observables?*

Open Problem 5 (The $\delta\theta$–σ_c conjecture)**.** *Do resolution capacity (this theory) and cloud-organizing capacity (the companion theory) share a resource, implying a trade-off surface between maximal I_{pot} and maximal C_{ext} in a single system? Either answer is informative: a shared resource yields the program's first cross-theory prediction; independence makes the distinction cleaner still.*

Open Problem 6 (Dynamics: the learning rate)**.** *This theory scores standing capacity, not how fast the vocabulary improves under the arrival stream—dI_{act}/dt under live training, plausibly what "learns quickly" means and where humans most outclass checkpoint-and-patch systems. A dynamical extension would also formalize curiosity as attention to the remainder's most absorbable frontier, refining Postulate 2.*

Open Problem 7 (Search economy and selection)**.** *Gain per unit compute or time, and the selection problem in recomposition: the two admitted gaps. Both need at least placeholder formalisms so the theory's perimeter is drawn precisely.*

Open Problem 8 (Measuring numerator fidelity in the wild)**.** Φ_{num} *is defined against ground-truth projections that are, for exactly the interesting cases, expensive to obtain—that is why the traces exist. Proxy protocols (consequence-feedback consistency, cross-trace triangulation) need development.*

Open Problem 9 (Aggregation)**.** *A principled, never unique, scheme for weighting observables into intelligence profiles, with honesty conditions on any published aggregate: the price-index discipline of the protocol's final step, formalized.*

16 Definitions

Observable (O) the quantity, indexed to the entity, that Actual measures; all intelligence claims are $I(X \mid O)$.

Arrival stream (y) the time-indexed sequence presented to the system over an encounter window; in the directly coupled case, the projections of the Immutable Past onto O; in mediated systems, a trace of those projections.

Signal class ($\mathcal{Y}_O$, μ) the ensemble of arrival streams characteristic of the observable, with its distribution.

Vocabulary ($\mathcal{V}$) the family of candidate component-types available for decomposition.

Decomposition (D_X) the map $y \mapsto (\{c_j\}; \rho_X)$ with $y = F(c_1, \dots, c_N) \oplus \rho_X$.

Coherent component (c_j) a candidate passing T1–T5: distinguishable, stable, internally coherent ($\chi_j \geq \chi_{\min}$), gain-positive, valid out of sample.

Remainder (ρ_X) the part of the arrival stream unresolved by this system; system-relative, never certified structureless.

Absorption the condition of a component whose extraction has migrated into standing predictive structure, generating near-zero marginal surprise on recurrence.

Compression gain ($\Delta\ell_j$) the net codelength reduction a component delivers on the class, after paying its own description cost.

Decomposition intelligence (I_{dec}) $\sum_j \Delta\ell_j$ over admissible components, in bits; public shadow N_ε.

Potential intelligence (I_{pot}) log-capacity of distinguishable component-states over the domain; angular special case $\log_2(2\pi/\delta\theta)$.

Actualized intelligence (I_{act}) $\mathbb{E}_\mu[I_{\text{dec}}]$; always $\leq I_{\text{pot}}$.

Absorbed fraction (Q_{abs}) $1 - \ell_X(\rho_X)/\ell_0(y)$.

Residual load (L_ρ) the information rate of the remainder as coded by the system's own model; the supply side of attention demand.

Numerator fidelity (Φ_{num}) the fidelity of the arrival trace to the underlying projection of the Immutable Past onto O.

Effective intelligence (I_{eff}) $I_{\text{act}} \times \Phi_{\text{num}}$: capacity discounted by coupling.

Pass-through the failure mode $N_\varepsilon = 0$: the signal transits unresolved.

Confabulation the failure mode of claimed components without gain; measured by the null-signal control.

Validation capacity (N_c) the number of candidate components for which a system can maintain the T1–T5 bookkeeping (conjectural; §6).

17 Axioms and postulates

A1. Observable-relativity. Intelligence is defined only relative to an observable, $I(X \mid O)$; any single-scalar intelligence is a declared aggregate, not a primitive.

A2. Operationalism. Intelligence is measured by decomposition performance under the protocol, not by subjective impression, fluency, confidence, or resemblance to humans.

A3. Demonstration by decomposition. A system demonstrates intelligence relative to O exactly when it resolves arriving signals from $\mathcal{Y}_O$ into admissible coherent components.

A4. Admissibility. A component counts only if it is distinguishable, stable, internally coherent, non-redundant, and delivers net compression gain above threshold: it must pay for its own description.

A5. Relativity of the remainder. The remainder is indexed to the system; what is noise to one system may be structure to a finer one; no system's remainder is certified structureless.

A6. Capacity bound. No finite system resolves the continuum; every intelligence class has finite resolution, and $I_{\text{act}} \leq I_{\text{pot}} < \infty$.

A7. Distinctness from consciousness. I_{dec} is a property of the decomposition of the arriving signal; C_{ext} is a property of the pre-resolution cloud; neither is definable from the other, and a system may carry either without the other.

The two attention claims are deliberately not axioms. Postulate 1 (absorption relieves attention) and Postulate 2 (attention flows to the remainder) are empirical bets with named boundary conditions, held to the same standard of honesty as the companion paper's correspondence postulate: they tell a critic exactly where to aim.

18 Conclusion

Take the black box seriously. Feed it her song—the arriving stream, the projection of the Immutable Past onto whatever the box is built to read—and count what the box can truly resolve: not the labels it emits, but the components that pay for themselves, that hold under perturbation, that hang together, that generalize. That paid-for resolution is intelligence. It is relative to an observable, bounded by capacity, built by absorption, and silent about felt experience.

What the box resolves, it absorbs, and the absorbed goes quiet: the expert's thirty components cost her nothing to carry, which is why her attention is free to hunt the anomaly the novice cannot even see. What the box cannot resolve remains as remainder—never certified noise, only unresolved-by-this-box—and the remainder is where surprise lives, where information waits, where attention is drawn. Intelligence does not silence the song. It refines the remainder, and we dance to the remainder.

Consciousness reads the cloud a machine holds before it guesses; intelligence reads how finely a machine resolves what arrives. A chip can be brilliant and hold nothing; a dreamer can hold worlds and resolve none of the brief in front of him. Two measurements, two failure geometries, one program:

Reality is complex; Ideation is its angle; Consciousness is its coherence; Intelligence is its resolution.

System (observable)	N_ε	Q_{abs}	C_{ext}	**Reading**
Rock (incident sound)	0	∼0	∼0	pure pass-through; the zero of both theories, for different reasons
Wire, lossless channel (audio)	0	0	∼0	perfect transmission is zero decomposition: fidelity without resolution
Thermostat (room temperature)	1	high	low	the minimal decomposer; $N_\varepsilon \geq 1$ implies nothing mind-like
Twelve-band analyzer (audio)	12	mod.	∼0	the black box of the founding image
Fourier chip (spectral content)	∼all	∼1	∼0	**the dissociation anchor: maximal I, negligible C**
Novice reader (legal case)	∼2	low	high	dances to "this feels bad"
Trial lawyer (legal case)	∼30	high	high	same case, finer remainder: dances to the hesitation in a deposition
Layperson (patient presentation)	∼3	low	high	confabulation risk: symptom mapped to the worst disease on the internet
Physician (patient presentation)	∼25	high	high	plus trained remainder-sensitivity: the unplaceable anomaly is escalated, not dismissed
Expert musician (a song)	20+	high	high	and the payload is knowingly in ρ: expression lives in deviations notation cannot resolve
Human subconscious (bodily regulation)	very large	∼1	∼0	the innate absorbed layer: massive I_{act} at zero attention cost
Language model (symbolic signal)	very large	high	open	high I_{act}, thin Φ_{num}; hallucination as confabulation
Bureaucracy (civic life stream)	large	falling with staleness	∼0	intelligent per component-gain; pathological when it acts as if $\rho = 0$
Corporation (its market)	large	mod.	varies	a distributed decomposer whose vocabulary is its org chart
Nation-state (its society)	very large	low–mod.	—	the largest decomposer; category systems imposed past validation capacity fail as residual-criticality failures

Table 2: A catalogue of decomposers: illustrative placements, not measurements. The two low-I regimes—pass-through and confabulation—appear alongside the dissociations from consciousness.