

The Shape of the Prediction Machine

A Theory of Consciousness from the Reality Equation

John Rector

July 1, 2026

Abstract. The hard problem of consciousness—why there is something it is like to be a system—is left intact by this theory and is not claimed to be solved. What is offered instead is an *outside-facing* estimate. Every entity we care to model can be assigned a *Reality Equation*, $R = A/E$, in which an exact, given *Actual* is divided by a complex, entity-specific *Expectation*. The real part of Expectation is produced by a *prediction machine* that resolves an internal probability cloud into a single guess. The theory’s central claim is that the *shape* of that cloud—before it resolves—is what an observer should read to estimate an entity’s apparent consciousness. Apparent consciousness is *coherent spread*: $C_{\text{ext}} \approx \sigma \times \kappa$, the product of how wide the cloud is (σ) and how well organized that width is into distinct, stable modes (κ). Because coherence fails once spread outruns a machine’s capacity, consciousness traces an inverted-U: collapse at one end (a pure tone), incoherence at the other (noise), and coherent spread in between (music). The theory is made dimensionally consistent, its coherence term is given two equivalent and computable definitions (an information-theoretic one and a phase-coherence one), its central quantity is tied back to the equation through the shape of the counterfactual surprise landscape, and it is made relative to a chosen observable. A measurement protocol, worked examples, testable predictions, and open problems are stated so that the framework can be scrutinized and, if wrong, refuted.

Keywords: consciousness, prediction, predictive processing, complex Expectation, coherence, spread, inverted-U, edge of chaos, observable-relativity.

Status of this document. This is a theoretical proposal released for public scrutiny. Its central bridge—that the structural quantity C_{ext} co-varies with felt experience—is stated as a *postulate* (§13), not proven; the explanatory gap is acknowledged, not crossed. The predictions in §14 are offered precisely so the theory can be tested and discarded if it fails. Criticism, counterexamples, and better formalizations are invited.

Contents

1	The core claim	2
2	The Reality Equation	3
2.1	Entity, and the model-relative stance	3
2.2	Actual	3
2.3	Expectation	3
2.4	Reality and Surprise	4
3	The prediction machine and its cloud	4
4	Spread, coherence, and the measure of consciousness	5
4.1	Spread	5
4.2	Coherence, in two equivalent readings	5

4.3	Consciousness as coherent spread	6
4.4	The inverted-U, derived	6
4.5	The landscape, and the curve as a slice	7
5	Tying the cloud to the equation: the surprise landscape	7
6	The fuller measure	8
7	A measurement protocol	8
8	Consciousness is relative to an observable	9
9	Learning, and why coherent spread is stable	9
10	A catalogue of entities	10
10.1	Worked scorings	10
11	The two failure modes and their signatures	11
12	Relation to existing frameworks	11
13	The correspondence postulate	12
14	Testable predictions	12
15	Open problems	12
16	Toward Ideation	13
17	Definitions	13
18	Axioms	14
19	Conclusion	14

1 The core claim

We cannot know what anything feels from the inside. We do not know what the sun feels, what a bird feels, what a rock feels, or what another human being feels. Subjective felt experience is private to the experiencing entity, and no equation dissolves that privacy. This theory does not try.

What it offers is a way to estimate consciousness from the outside, by asking a single question:

What is the shape of the prediction machine that generates the real component of this entity's Expectation?

The real component is a single number—the machine’s guess. But the machine that produces that number has a shape: the probability cloud that exists before it resolves into the guess. The denominator of the Reality Equation contains the guess; the machine contains the cloud; and the shape of the cloud is what determines apparent consciousness.

Apparent consciousness is approximately spread times coherence,

$$C_{\text{ext}} \approx \sigma \times \kappa$$

where σ is the width of the cloud and κ is how coherently that width is organized. Low consciousness appears in two opposite ways. At one extreme the machine is *collapsed*: $\sigma \rightarrow 0$, and there is almost nothing to choose among. At the other extreme it is *incoherent*: σ is large but $\kappa \rightarrow 0$, too much unresolved signal and no stable way to turn it into a meaningful guess. High consciousness lives between these extremes, where an entity can hold many possible “this-and-thats,” distinguish and transform them, and still resolve them into one guess, one action, one direction. Consciousness is organized uncertainty. It is the felt geometry of prediction.

2 The Reality Equation

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

$$R = \frac{A}{E}, \quad S = \ln R,$$

where R is Reality, A is Actual, E is Expectation, and S is Surprise. Actual is the numerator; Expectation is the denominator; Reality is the quotient; Surprise is the logarithm of the quotient. Because Expectation is complex, Reality and Surprise are complex as well.

2.1 Entity, and the model-relative stance

An entity is that which a Reality Equation represents. It may be a human being, a bird, a rock, the sun, an intersection, a firm, a market, a family, or a civilization—any coherent actualizer we choose to examine. “Entity” does not mean “biological organism”; it means “the thing for which we are writing a Reality Equation.”

An important clarification, and a guard against misreading the theory as panpsychism: the prediction machine is not asserted to inhere metaphysically in the entity. It is the structure contained in *an observer’s best predictive model* of the entity, relative to an observable. When we say a rock “has a prediction machine,” we mean that the best model of the rock’s behavior with respect to some observable contains a (nearly trivial) predictive structure. C_{ext} is therefore a property of a model, estimated from outside—not a claim about rock-experience. This keeps the theory operational and falsifiable rather than mystical.

2.2 Actual

Actual is given by the Immutable Past. It is exact, scalar, and not produced by the entity: it arrives, it does not guess, it does not negotiate. But Actual is *indexed by the observable*. The universal state of the world is one thing; the Actual that enters a particular entity’s equation is the projection of that state onto the observable the entity is predicting,

$$A = \Pi_O(\text{world}).$$

The hawk’s shadow that is Actual for the bird is not the Actual for the rock beside it. Actual is common in its *source*—the one Immutable Past—but specific in its *projection*. It remains exact and given; it is simply about something.

2.3 Expectation

Expectation is the denominator. It is entity-specific and complex. To make the sum well-formed we work in dimensionless units: fix a reference scale x_0 natural to the observable (a baseline, a

just-noticeable difference, or the Actual itself), and write $\tilde{P} = P/x_0$, $\tilde{A} = A/x_0$. Then

$$E = \tilde{P} + iB,$$

where \tilde{P} is the (normalized) real component—the prediction machine’s resolved guess—and B is the dimensionless magnitude of the imaginary component, contributed by the entity’s orientation toward ideas.

The imaginary component deserves care. Ideas exist as vectors on the unit circle of future possibility; each has a direction, an angle, pointing toward a possible orientation in the unknowable future. An entity does not weight all ideas equally—if it did, the vector sum would cancel toward zero. The meaningful imaginary component is the entity’s *weighted* resultant idea-vector, with magnitude B and argument θ ,

$$V_I = B e^{i\theta}.$$

For the theory of consciousness developed here, the angle θ is not the crux; it becomes central in the companion theory of Ideation (§16). What matters here is that Expectation is not merely a quantity but a quantity *from an orientation*: not just “how much,” but “from what angle.”

2.4 Reality and Surprise

Since \tilde{A} is a pure number and E is complex, Reality is complex,

$$R = \frac{\tilde{A}}{\tilde{P} + iB} = |R| e^{i\varphi},$$

and Surprise, the complex logarithm of Reality, has a magnitude part and an angular part,

$$S = \ln R = \underbrace{\ln |R|}_{\text{magnitude-surprise}} + i \underbrace{\varphi}_{\text{angular surprise}}.$$

We take the principal branch, $\varphi \in (-\pi, \pi]$, so angular surprise is the smallest signed rotation carrying Expectation onto Actual, and we require $E \neq 0$ (an entity with zero Expectation has an undefined Reality, which is the correct behavior). If Actual matches Expectation, $R = 1$ and $S = 0$: no surprise. As Actual diverges from Expectation, surprise grows—in magnitude, and in angle. Surprise is not merely “how much did Actual differ from Expectation”; it is also “from what direction did Actual correct it.”

Surprise, however, is not the measure of consciousness. A wide, coherent machine may guess accurately, so that Actual arrives close to Expectation and Surprise is low—yet the entity is highly conscious, because consciousness is read from the *machine that produced the guess*, not from the correction that follows. This decoupling of consciousness from surprise is one of the theory’s load-bearing commitments.

3 The prediction machine and its cloud

Every entity, in its model, has a prediction machine that generates the real component P of Expectation. The difference between entities is not that some have prediction machines and others do not; it is the *shape* of the machine.

Formally, let $p(x)$ be the machine’s predictive distribution over normalized outcomes x for the observable O —the probability cloud. The resolved guess is a functional of the cloud,

$$\tilde{P} = \mathcal{G}[p] \quad (\text{e.g. the mean } \int x p(x) dx, \text{ or the mode}),$$

and this single number is what enters the denominator. The cloud itself does not sit in the denominator—the denominator receives the resolved guess. The cloud lives in the machine. And the cloud, before it resolves, is the object that carries consciousness.

A collapsed cloud produces a highly certain guess with almost nothing behind it. A broad but featureless cloud produces a guess surrounded by unresolved noise. A broad but organized cloud produces a guess that has emerged from many distinguishable, resolvable alternatives. The outside-facing test is therefore not “how wide is the cloud” alone, but “can the width be transformed into distinct, nameable, resolvable alternatives before it collapses into the guess.”

4 Spread, coherence, and the measure of consciousness

4.1 Spread

Spread is the width of the cloud. A clean, bounded choice is normalized dispersion,

$$\sigma = \frac{\sqrt{\text{Var}_p[x]}}{x_{\text{ref}}},$$

optionally squashed to $[0, 1]$ by $\sigma \mapsto \sigma/(1 + \sigma)$ for comparison across entities. A delta cloud (perfect certainty) has $\sigma \rightarrow 0$; a broad cloud has large σ . Spread by itself says nothing about organization—that is the work of coherence.

4.2 Coherence, in two equivalent readings

Coherence is the degree to which the cloud’s width is organized into distinct, stable modes rather than a featureless smear. It is defined purely on the form of the cloud, with no appeal to action or to language, so that it does not double-count the auxiliary factors of §6. Two equivalent readings make it computable.

(i) Information-theoretic (negentropy at fixed spread). Among all clouds of a given variance, the smoothest—highest entropy—is the Gaussian; a featureless smear is the noise limit. Define

$$\kappa = 1 - \frac{H(p)}{H_{\max}(\text{Var}_p)} \in [0, 1], \quad H_{\max}(v) = \frac{1}{2} \ln(2\pi e v),$$

where $H(p)$ is the entropy of the cloud. Then $\kappa \rightarrow 0$ for an unstructured smear (noise), and $\kappa \rightarrow 1$ for a cloud concentrated into sharply separated modes (well-resolved alternatives). Coherence is departure from maximal smoothness toward distinct modes—precisely “distinguishable this-and-thats.”

(ii) Spectral (phase coherence). Represent the machine’s predictive dynamics as a signal and decompose it into components—the Fourier-like transform. A signal that looks like one wave is in fact a resultant of many component frequencies; the same is true of the cloud behind the guess. Let component j have weight w_j and phase ϕ_j . Coherence is the phase-alignment order parameter

$$\kappa = \frac{|\sum_j w_j e^{i\phi_j}|}{\sum_j w_j} \in [0, 1].$$

A pure tone has a single component, so its spread is near zero. White noise has many components with random phases, so the vectors cancel and $\kappa \rightarrow 0$. Music has many components with organized phases, so both σ and κ are high (Figure 3). This is the literal, physical meaning of “coherence”—phase alignment—and it makes the music analogy exact. It also gives the imaginary axis of Expectation a role in consciousness after all, since phase is angle; this is the seam along which the theory of Consciousness and the theory of Ideation are joined (§16).

The two readings agree in spirit: structured-given-spread is the same thing as phase-organized. Either can serve as κ .

4.3 Consciousness as coherent spread

Apparent consciousness is the product

$$C_{\text{ext}} \approx \sigma \times \kappa.$$

The product form captures both failure modes at once: if $\sigma \rightarrow 0$ then $C_{\text{ext}} \rightarrow 0$ (collapse), and if $\kappa \rightarrow 0$ then $C_{\text{ext}} \rightarrow 0$ (incoherence). Randomness is not consciousness: static is unpredictable and spread out, but incoherent, so its C_{ext} is low. Certainty is not consciousness either: a collapsed cloud has nothing to organize. Consciousness requires coherent spread—an entity that is not merely uncertain but *coherently* uncertain, holding its identity while its machine resolves a structured cloud into a guess.

4.4 The inverted-U, derived

The bare product $\sigma \times \kappa$ does not by itself bend downward: if κ were independent of σ , more spread would mean monotonically more consciousness. The downturn is real because coherence *fails once spread outruns the machine's capacity* to organize it. Let σ_c be that organizing capacity, and let

$$\kappa(\sigma) = \frac{1}{1 + (\sigma/\sigma_c)^m}, \quad C_{\text{ext}}(\sigma) = \frac{\sigma}{1 + (\sigma/\sigma_c)^m}.$$

This has an interior maximum at

$$\sigma^* = \sigma_c (m - 1)^{-1/m},$$

so for $m = 2$ the peak sits at $\sigma^* = \sigma_c$ with $C_{\text{ext}}^{\text{max}} = \sigma_c/2$ (Figure 1). The inverted-U is thus a theorem about a machine of finite capacity, not a picture. Its content is striking: *consciousness is maximized when spread is pushed right up to the edge of the machine's ability to keep it coherent*—an edge-of-chaos statement. Below the edge lies collapse (tone); above it lies incoherence (noise); at the edge lies coherent spread (music).

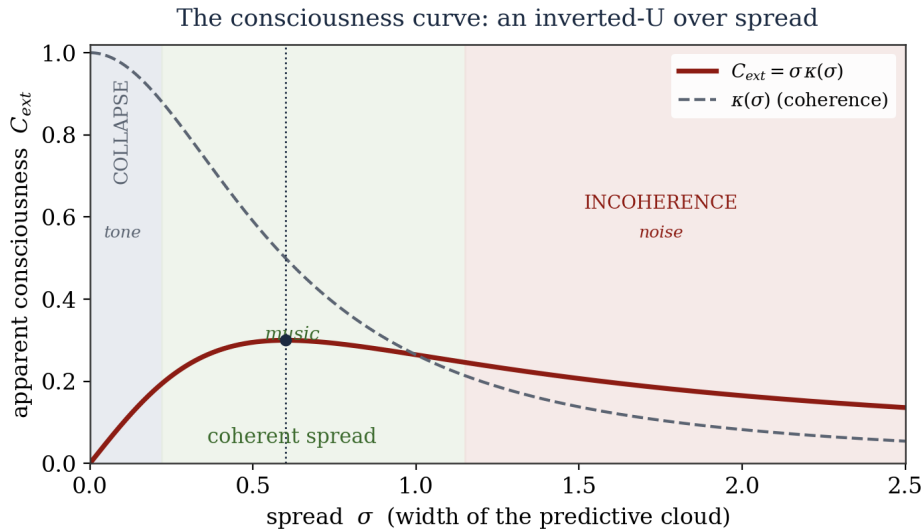


Figure 1: The consciousness curve. With finite organizing capacity σ_c , coherence $\kappa(\sigma)$ falls once spread outruns capacity, so $C_{\text{ext}} = \sigma \kappa(\sigma)$ rises, peaks at σ^* , then decays: tone \rightarrow music \rightarrow noise.

4.5 The landscape, and the curve as a slice

Spread and coherence are two independent axes, not one. The surface $C_{\text{ext}} \approx \sigma \times \kappa$ is low along both edges—collapse on the left ($\sigma \rightarrow 0$), incoherence along the bottom ($\kappa \rightarrow 0$)—and high only in the coherent-spread interior (Figure 2). The inverted-U of Figure 1 is what one sees walking along the capacity constraint $\kappa = \kappa(\sigma)$ across this landscape. Placing entities as points (σ, κ) makes the theory’s predictions visual, and shows why a rock and white noise are equally unconscious for opposite reasons, and why anesthesia and delirium are the two ways a human falls off the peak.

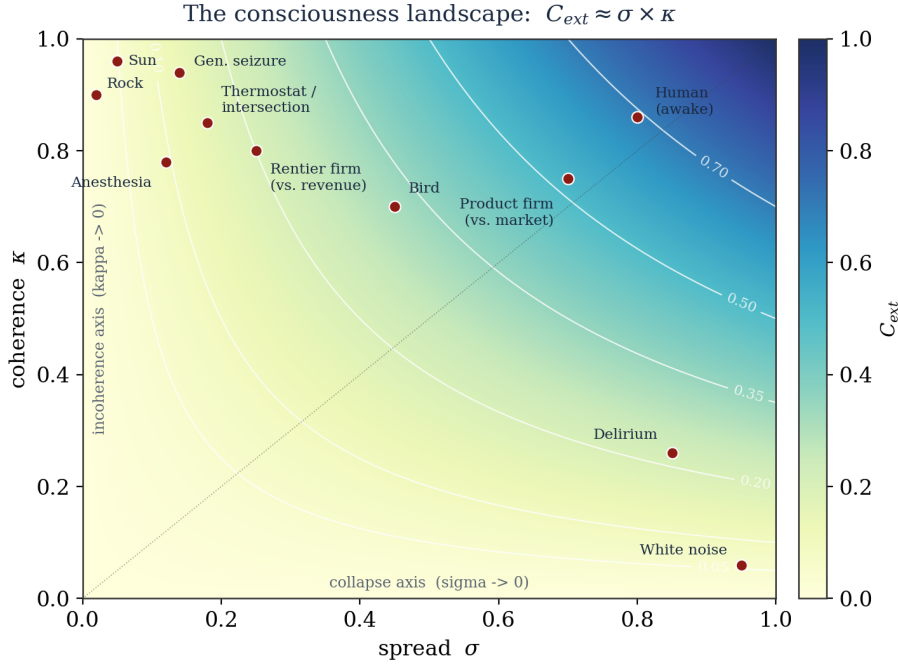


Figure 2: The consciousness landscape $C_{\text{ext}} \approx \sigma \times \kappa$. Collapse lives on the left edge, incoherence on the bottom edge; coherent spread—high consciousness—lives in the upper-right interior. Entity placements are illustrative (Table 1).

5 Tying the cloud to the equation: the surprise landscape

If only the resolved guess P enters the denominator, one might worry that the cloud—and hence consciousness—is disconnected from the Reality Equation. It is not. The cloud is exactly the machine’s distribution over the actuals it takes seriously, so it determines how surprise *would* vary if different actuals arrived. Read the counterfactual surprise function over the support of the cloud,

$$S(a) = \ln\left(\frac{\tilde{a}}{E}\right), \quad a \in \text{supp}(p).$$

Two summary statistics recover the two axes from outside:

$$\text{expected surprise: } \bar{S} = \mathbb{E}_p[\ln(1/\tilde{p}(a))] \quad (\text{a free-energy-like quantity}),$$

$$\text{surprise structure: } \Sigma_S = \text{the modal structure of } S(a).$$

A collapsed machine has a flat $S(a)$: everything is equally unsurprising, or equally catastrophic. A coherent-spread machine has a structured $S(a)$ with distinct basins, one for each of its

Spectral reading of the predictive cloud




Figure 3: The spectral reading of the cloud (§4.2(ii)). Tone: one component, $\sigma \rightarrow 0$. Music: several phase-organized components, σ and κ both high. Noise: many components with random phases, $\kappa \rightarrow 0$.

distinguishable hypotheses. Consciousness is legible in the shape of this response surface—and, crucially, the response surface is probeable without any access to felt experience. The cloud touches the equation not through the guess it happens to emit, but through the structure of the surprise it is prepared to feel.

6 The fuller measure

Because coherence is now defined purely on cloud-shape, the remaining features of a prediction machine are cleanly separate and can be folded in as multiplicative modulators:

$$C_{\text{ext}}(\mathcal{X} | O) \propto \sigma \cdot \kappa \cdot G \cdot \Lambda \cdot K,$$

where

- G is *sensitivity*: how small a deviation between guess and Actual the entity can register. A human can be reorganized by a half-second of silence; the sun cannot.
- Λ is the *probability horizon*: how far into the future the cloud reaches before resolving into a guess—how far ahead the entity “worries,” in the sense that future uncertainty presses into present prediction. A rock’s horizon is near zero; a human’s spans decades.
- K is *action-coupling*: the degree to which the cloud’s structure actually shapes what the entity does. A forecast, a spreadsheet, or a database can hold many futures without being conscious; the difference is whether the predictive structure treats a possibility as real enough to approach it, avoid it, repeat it, or build around it.

The core relation remains $C_{\text{ext}} \approx \sigma\kappa$; sensitivity, horizon, and action-coupling sharpen the estimate without redefining it.

7 A measurement protocol

The theory is only worth scrutinizing if C_{ext} can, in principle, be estimated from outside. The following recipe estimates $C_{\text{ext}}(\mathcal{X} | O)$ without any appeal to felt experience.

1. **Fix the observable O** (what counts as Actual) and its reference scale x_0 .

2. **Probe** the entity across a representative ensemble of contexts.
3. **Estimate the cloud** $p(x)$ per context—from the dispersion of the entity’s pre-outcome states, guesses, or hedged commitments: neural population states for a brain, forecast and position dispersion for a firm, microstate distributions for a physical system.
4. **Compute** spread σ (dispersion) and coherence κ (negentropy at fixed spread, or the spectral phase-order parameter of §4.2).
5. **Aggregate** $C_{\text{ext}} = \sigma \kappa$, optionally modulated by sensitivity G , probability horizon Λ , and action-coupling K (§6).
6. **Cross-check** against the surprise landscape $S(a)$ of §5: a coherent-spread machine should exhibit a structured, multi-basin surprise landscape, whereas a collapsed or incoherent one should not.

The numbers so obtained are model-confidence-limited for entities we cannot directly interrogate (§15), but the procedure is the same in every case, which is what makes C_{ext} a candidate for genuine comparison across entities.

8 Consciousness is relative to an observable

Because Actual is indexed by an observable, so is consciousness. Write $C_{\text{ext}}(\mathcal{X} | O)$: the apparent consciousness of entity \mathcal{X} with respect to what it is predicting. This is not a hedge; it is a feature. The same bird may be near the peak on $C_{\text{ext}}(\cdot | \text{predator motion})$ and near collapse on $C_{\text{ext}}(\cdot | \text{next winter})$. “Consciousness of what” becomes a first-class question, which is closer to how the word is actually used. A single scalar for a whole entity is then a chosen aggregate over observables, not a primitive.

This relativity is what makes the economic reading precise rather than metaphorical. A product-selling firm has a market-facing prediction machine: it must guess demand, price, quality, inventory, reputation, competition, substitution, timing. It has no guaranteed revenue—only a forecast—so its revenue-facing cloud is wide and it must stay sensitive to Actual, because the customer keeps arriving as numerator and can say yes, no, leave, or switch. Its $C_{\text{ext}}(\text{firm} | \text{revenue})$ is correspondingly high. A rentier firm, by contrast, converts open market uncertainty into lease, title, debt, toll, license, lock-in, or enforceable claim. It does not want to ask each day “will they buy”; it wants to say “they owe.” It narrows its revenue-facing cloud—reduces spread, reduces sensitivity to daily Actual, replaces guess with claim—and so its $C_{\text{ext}}(\text{firm} | \text{revenue})$ is low *by design*. The same firm may still have high C_{ext} with respect to litigation, interest rates, or political risk. When a product firm becomes a rentier, it is not merely changing its business model; it is changing the shape of a prediction machine—collapsing a cloud. Product capitalism stays exposed to Actual; rentier capitalism collateralizes it. None of this asserts corporate sentience; it is a claim about the predictive engagement of a subsystem with a chosen observable.

9 Learning, and why coherent spread is stable

Attention is a conversation between the unconscious prediction machine and conscious felt experience. Actual arrives; Expectation predicted; surprise may appear; the machine asks whether a deviation is noise or a new norm, and the entity answers—primarily through action, not language. Repeated action trains Expectation; repeated Actual reshapes the machine. Learning is the reshaping of the prediction machine after repeated interaction with Actual. The machine is dense and conservative: it does not update after one signal, because it wants low

prediction error and stability, and it changes only when repeated Actuals confirm that a new pattern is real.

This raises an apparent paradox: if the machine minimizes prediction error, does it not narrow its cloud toward collapse, and so toward *lower* consciousness? The resolution is to distinguish minimizing *observed* error from minimizing *expected* error in a rich, nonstationary world. In a genuinely varied environment, a collapsed model is catastrophically wrong the moment the world moves; the expected-error-minimizing model is therefore not the narrowest one but the one whose coherent spread *matches* the structured complexity of its environment, \mathcal{C}_{env} . Too narrow and the model is blindsided; too broad and unstructured and it cannot act. Coherent spread is not fragile—it is the optimum for a complex world. Hence a compact statement:

Consciousness is the internalized shape of the world's own structured uncertainty.

10 A catalogue of entities

The rock. A collapsed prediction machine: narrow cloud, guess nearly equal to Actual, stable Reality, surprise near zero, attention demand near zero. It does not meaningfully guess, worry, learn, or answer prediction with action. It is highly actual and minimally conscious—low consciousness by collapse.

The bird. A less collapsed machine. It guesses food, threat, branch, nest, weather, movement, mate. Its cloud is wider than the rock's, its horizon real but mostly near-field, its action-coupling quick and strong: it flies, hides, hunts, calls, lands, builds, returns. More conscious than the rock, through more coherent spread and stronger coupling.

The human being. A highly spread-out and coherent machine. It guesses across seconds, years, decades—identity, love, death, money, reputation, meaning, children, work, legacy, history. Its horizon is vast, its sensitivity high: a sentence can redirect a life, a silence can become memory, a look can become history. It can feel many this-and-thats and still resolve them into one guess, one action, one direction. That resolution is choice: not necessarily metaphysical free will, but the felt selection of this over that, here over there, approach over avoid. The human is highly conscious because the human machine is spread, coherent, sensitive, future-reaching, and action-coupled.

The sun. A great maker of history—heat, light, orbit, season, fusion, growth—yet from our frame a highly stable, comparatively collapsed machine with a predictable pattern. Power is not consciousness; scale is not consciousness; energy is not consciousness. Coherent spread is.

The intersection. An engineered entity whose prediction machine is built into lights, lanes, timing, sensors, and rules. It anticipates collision, congestion, delay, error, flow. Its cloud is engineered rather than biological; its apparent consciousness is the degree to which that cloud carries spread, stays coherent, responds to Actual, and shapes action through signals and design.

The incoherent entity. The other way to be low in consciousness: high spread, low coherence. Not collapse but noise—many unresolved signals that cannot be named, differentiated, or resolved, a machine that cannot perform the transform that turns spread into music. Low consciousness by incoherence.

10.1 Worked scorings

Table 1 scores this cast on $[0, 1]$ scales. These are *illustrative* placements meant to display the structure of the theory, not measurements; the protocol of §7 is how real numbers would be earned. Note the diagnostic pairs: rock versus white noise (equally low, opposite modes); anesthesia versus delirium (the two ways off the human peak); and the generalized seizure—massive phase-synchrony (κ high) but collapsed differentiation (σ low)—which on this theory is

why hypersynchrony coincides with the *loss* of consciousness rather than its intensification.

Entity (observable)	σ	κ	C_{ext}	Regime
Sun (radiative state)	0.05	0.96	0.05	collapse
Rock (mechanical state)	0.02	0.90	0.02	collapse
Generalized seizure (cortex)	0.14	0.94	0.13	collapse (hypersynchrony)
Anesthesia (cortex)	0.12	0.78	0.09	collapse
Thermostat / intersection	0.18	0.85	0.15	engineered, narrow-coherent
Rentier firm (revenue)	0.25	0.80	0.20	collapsed by design
Bird (near-field survival)	0.45	0.70	0.32	coherent spread
Product firm (market)	0.70	0.75	0.53	coherent spread
Human, awake (open world)	0.80	0.86	0.69	coherent spread (peak)
Human, delirium (open world)	0.85	0.26	0.22	incoherence
White noise / static	0.95	0.06	0.06	incoherence

Table 1: Illustrative scorings, $C_{\text{ext}} = \sigma \times \kappa$. The two low- C_{ext} clusters sit at opposite corners of Figure 2: collapse (low σ) and incoherence (low κ).

11 The two failure modes and their signatures

The theory earns its keep by predicting *two distinct* ways to lose consciousness, which single-axis theories tend to blur. The following are offered as resonances to be tested, not as validated results:

- **Collapse (low σ).** Deep dreamless sleep and general anesthesia reduce the differentiation of cortical states. A generalized seizure is the extreme: enormous phase-synchrony (κ high) with almost no differentiation (σ collapsed)—on this theory, precisely why hypersynchrony abolishes consciousness rather than heightening it.
- **Incoherence (low κ).** Delirium and some acute psychotic states show abundant, shifting activity that fails to resolve into stable, distinguishable content.
- **Coherent spread (the peak).** Ordinary waking awareness: broadband activity with organized long-range phase relations.

This is a genuine differential prediction: two states can have opposite (σ, κ) signatures yet similar behavioral unconsciousness, distinguishable by whether spread or coherence collapsed.

12 Relation to existing frameworks

The theory is not free-floating; its quantities have recognizable relatives, which is some evidence it carves at real joints.

- **Predictive processing and the free energy principle** (in the spirit of Clark, Hohwy, and Friston). $R = A/E$ and $S = \ln R$ are a prediction-error and surprisal calculus, and the expected surprise \bar{S} of §5 is a free-energy-like quantity. The distinctive move here is to refuse to equate consciousness with prediction error, locating it in the *shape* of the cloud instead.
- **Integrated Information Theory** (Tononi). The product $\sigma \times \kappa$ (differentiation times organization) is structurally parallel to the “differentiated yet integrated” character of Φ . This theory gains a cleaner two-failure-mode geometry and an explicit route to measurement; it shares the burden of a correspondence postulate (§13).

- **Criticality and the edge of chaos** (in the tradition of self-organized criticality). The result $\sigma^* = \sigma_c$ says consciousness peaks at the boundary where coherence begins to fail—a critical point—connecting the inverted-U to a body of work on neural systems operating near criticality.

These are positioning notes, not endorsements or derivations.

13 The correspondence postulate

Nothing in the mathematics *forces* a structural quantity to be about felt experience. This is the explanatory gap, and the theory does not claim to cross it. It states the bridge honestly, as a postulate on a par with the central identity of other structural theories.

Postulate 1 (Correspondence). *Wherever an entity has felt experience at all, the richness of that experience co-varies with C_{ext} , the coherent spread of its prediction machine.*

This is a hypothesis about co-variation, testable in the one system we can interrogate from the inside—ourselves—not a proof that coherent spread *is* experience. Declaring its status is a strength: it tells a critic exactly where to aim.

14 Testable predictions

The formalization yields commitments that could turn out false, which is what makes this a theory rather than a picture.

Prediction 1 (Inverted-U over spread). *Within a fixed machine, driving spread past capacity σ_c lowers C_{ext} . Over-broadened, phase-disorganized states should show reduced markers of awareness, not increased.*

Prediction 2 (Two dissociable collapses). *Loss of consciousness comes in a low- σ variety and a low- κ variety, with opposite (σ, κ) signatures despite similar behavioral unconsciousness.*

Prediction 3 (Complexity-matching). *C_{ext} tracks the structured complexity of the niche an entity must model, up to its capacity σ_c ; enrich the niche and C_{ext} rises to meet it, until capacity is exceeded and the system tips into incoherence.*

Prediction 4 (Observable-relativity). *The same entity yields different C_{ext} for different observables; a firm's revenue- C_{ext} falls as it rentierizes even when its total activity is unchanged.*

Prediction 5 (Shape, not size, of surprise). *High- C_{ext} systems are marked by structured counterfactual surprise landscapes even when their realized surprise is low; accurate prediction does not by itself lower consciousness.*

15 Open problems

Honest residue for future work, stated so others can attack it.

- **The correspondence postulate** is undischarged and probably cannot be discharged from outside. The best available move is to make it precise and test its co-variation where we have inside access.
- **Choosing the observable.** Because C_{ext} is observable-relative, a principled way to select or aggregate observables is needed before single-number comparisons across very different entities are fully meaningful.

- **Estimating the cloud** for entities we cannot interrogate—a rock, the sun—means estimating a distribution we can only model, not sample; those numbers are limited by model-confidence by construction.
- **The resolving functional** $\mathcal{G}[p]$ —mean, mode, or sample—is left open and may itself be entity-specific.
- **Uniqueness of the measure.** Whether $\sigma \times \kappa$ is the right combiner, or merely a good first one, is open; alternative combiners making different predictions would be welcome.

16 Toward Ideation

The phase-coherence reading of κ (§4.2(ii)) leaves a gift for the companion theory. If coherence is phase coherence, then the imaginary component of Expectation—idea-orientation $B e^{i\theta}$ —is not a separate department but the same angular structure seen at the level of commitments rather than components. Consciousness and Ideation become two readings of one complex object:

$$\begin{aligned} \text{Consciousness} &\sim \text{the coherence of phase across the cloud,} \\ \text{Ideation} &\sim \text{the net phase } \theta \text{ the cloud commits to.} \end{aligned}$$

A conscious entity holds many phase-organized possibilities (high κ); an ideating entity collapses them toward a dominant angle θ . Ideation is the directional resolution of a coherent cloud. Reality is complex; Ideation is its angle; Consciousness is its coherence.

17 Definitions

Entity that which a Reality Equation represents, relative to an observer’s model.

Observable (O) the quantity, indexed to the entity, that Actual measures.

Actual (A) the exact, given projection of the Immutable Past onto the observable.

Expectation (E) the complex denominator, $E = \tilde{P} + iB$, in dimensionless units.

Real component (\tilde{P}) the prediction machine’s single resolved guess, $\mathcal{G}[p]$.

Imaginary component ($B e^{i\theta}$) the weighted resultant idea-vector; magnitude B , angle θ .

Reality (R) the complex quotient \tilde{A}/E .

Surprise (S) $\ln R = \ln |R| + i\varphi$, principal branch; magnitude- and angular-surprise.

Prediction machine the structure in the model that generates \tilde{P} .

Probability cloud (p) the machine’s distribution over outcomes before it resolves into \tilde{P} .

Spread (σ) the normalized width of the cloud.

Coherence (κ) the cloud’s organization into distinct, stable modes; equivalently, its phase alignment.

Apparent consciousness (C_{ext}) the outside estimate, $C_{\text{ext}} \approx \sigma \times \kappa$.

Sensitivity (G) the smallest guess–Actual deviation the entity registers.

Probability horizon (Λ) how far ahead the cloud reaches before resolving.

Action-coupling (K) the degree to which the cloud shapes behavior.

Choice the felt resolution of coherent spread into one guess, one action, one direction.

Learning the reshaping of the machine through repeated interaction with Actual.

18 Axioms

1. Every entity we choose to model can be represented by a Reality Equation, relative to an observer and an observable.
2. Actual is the exact, given projection of the Immutable Past onto the observable; it is common in source, specific in projection.
3. Expectation is the complex, dimensionless denominator $E = \tilde{P} + iB$, with $E \neq 0$.
4. The real component \tilde{P} is a single number: the prediction machine's guess, a functional $\mathcal{G}[p]$ of its cloud.
5. The cloud belongs to the machine, not to the denominator; the denominator receives the resolved guess.
6. Spread is the width of the cloud; coherence is its organization into distinct, stable modes (equivalently, its phase alignment).
7. Apparent consciousness is coherent spread, $C_{\text{ext}} \approx \sigma \times \kappa$; it vanishes if either factor vanishes.
8. Coherence degrades once spread outruns capacity σ_c , so C_{ext} traces an inverted-U, peaking at the edge of coherence.
9. Collapsed certainty is not consciousness; incoherent uncertainty is not consciousness; coherent uncertainty is consciousness.
10. Surprise is the complex correction of the guess by Actual; consciousness is read from the machine, not from the correction.
11. Consciousness is relative to an observable.
12. Angle matters for Ideation; shape matters for Consciousness.
13. The identification of C_{ext} with felt experience is a postulate, not a theorem.

19 Conclusion

The hard problem remains. But the outside estimate is now something one can compute rather than merely assert. Do not ask only what an entity feels; ask what the shape of the prediction machine is that generates its real component. The denominator contains the guess; the machine contains the cloud; the cloud resolves into a number. Spread is the width of the cloud; coherence is its organization; consciousness is their product. Too little spread is collapse; too little coherence is noise; coherent spread—held right up to the edge where coherence would fail—is consciousness. The cloud reaches the equation through the shape of the surprise it is prepared to feel. Reality is complex; Ideation is its angle; Consciousness is its coherence.