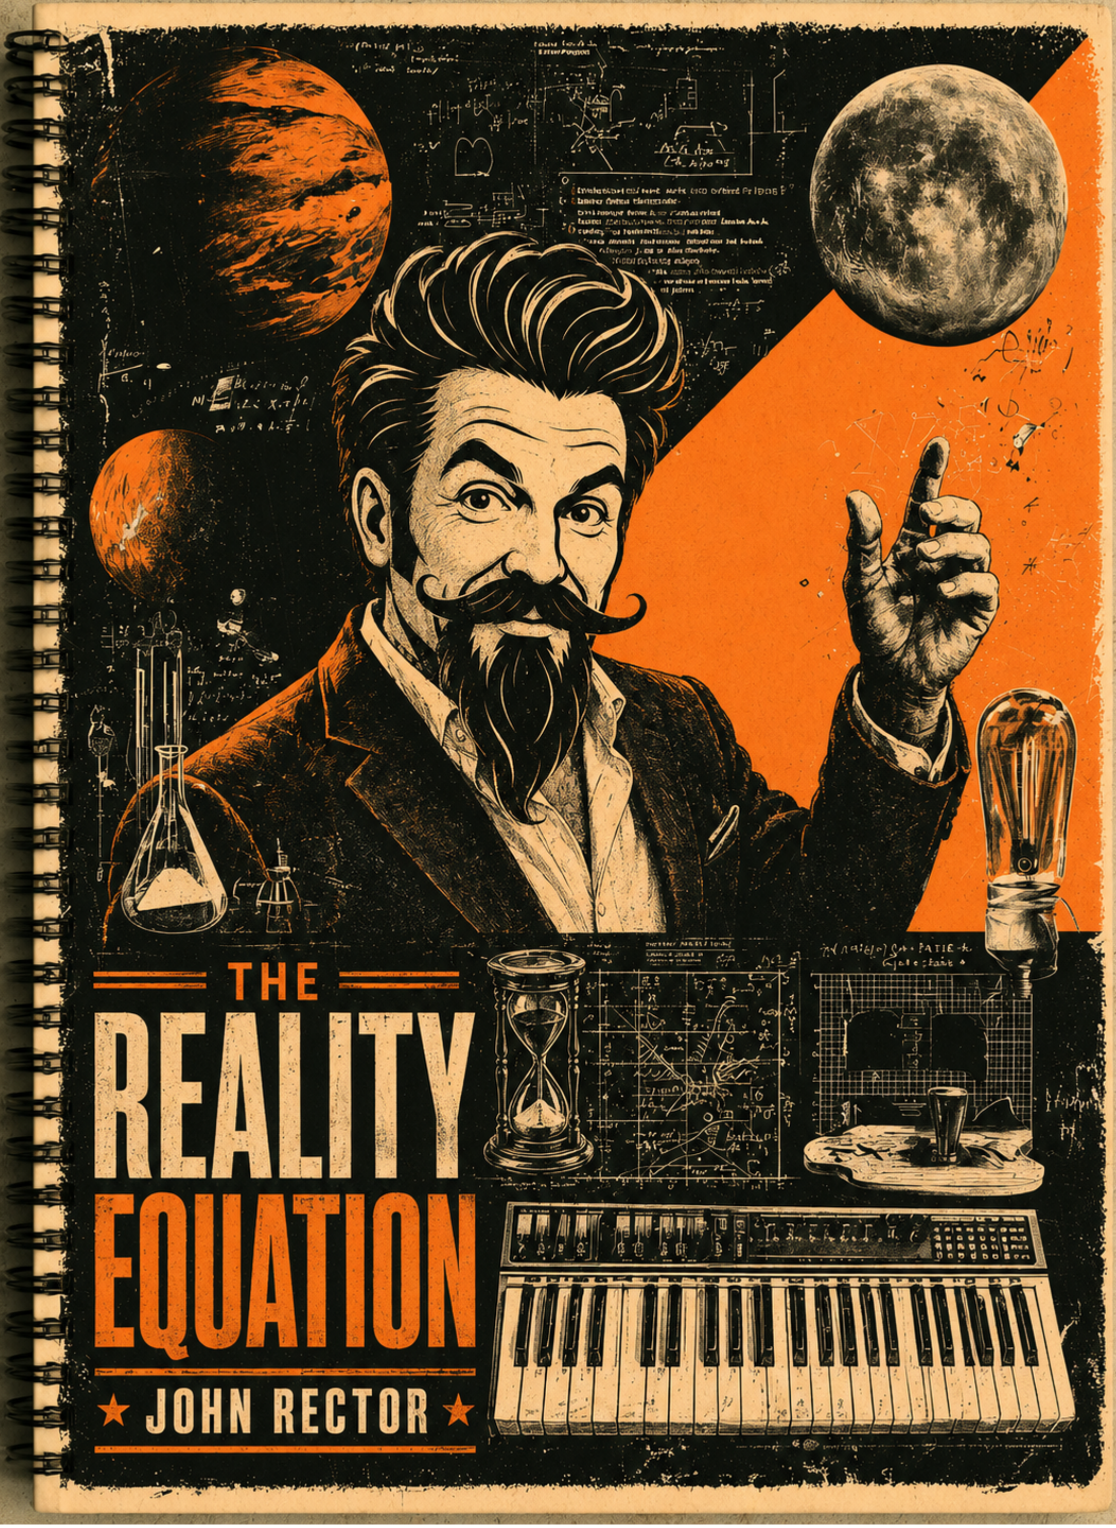
THE
REALITY
EQUATION
JOHN RECTOR

The Reality Equation

SECOND EDITION

A Theory of Reality, Surprise, Consciousness, and Intelligence

John Rector

Charleston, South Carolina

This edition supersedes the first edition of *The Reality Equation* and incorporates the results of two subsequent papers, *The Shape of the Prediction Machine: A Theory of Consciousness from the Reality Equation* (1 July 2026) and *Intelligence as Decomposition: A Theory of Intelligence from the Reality Equation* (3 July 2026).

Portions of this book state hypotheses, postulates, and conjectures rather than established results. Where that is so, the text says so in the sentence that makes the claim. The author invites criticism, counterexamples, and better formalizations.

johnrector.me

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

Contents

Preface to the Second Edition

The first edition of this book did one thing well. It disciplined the language. It separated Actual from Ideal from Real from Reality, it made Expectation complex, it formed the quotient honestly, and it derived surprise instead of merely describing it. Readers who worked through it came out with a vocabulary they could not un-learn. That was the point, and I would write those chapters again.

But the first edition stopped at the doorway of the two questions the framework was built to answer, and it stopped because I could not yet answer them. If Reality is a quotient, and the denominator is produced by a prediction machine, then two enormous things follow that the first edition only gestured at.

The first is that the machine has a *shape*. The denominator receives a number—one resolved guess—but the machine that produces that number holds a distribution before it resolves. Everything we mean by consciousness, I now think, is legible in the shape of that distribution rather than in the guess it emits or the surprise that corrects it. That claim, worked out with its mathematics, its failure modes, its measurement protocol, and its honestly undischarged bridge to felt experience, is Part Three of this edition. It did not exist in the first.

The second is that the numerator has a *fate*. Actual arrives; something happens to it. Some of what arrives gets resolved into components a system can name, hold, and re-use; the rest does not. That split—resolved against unresolved—turns out to be what we have always been pointing at with the word intelligence, and the unresolved part turns out to be where surprise, information, and attention actually live. That is Part Four. It did not exist in the first edition either.

The two results are companions, and their relationship is the most useful thing in this book. Consciousness reads the cloud a system holds *before* it guesses. Intelligence reads how finely a system resolves what *arrives*. They are measurements of different objects at different moments, they fail in mirror-image ways, and they come apart in both directions—which is why a Fourier chip can be brilliant and hold nothing, and a dreaming novice can hold worlds and resolve nothing of the brief in front of him. For all of human history those two quantities travelled together, because the same evolved machinery produced both. They do not travel together anymore. Getting that distinction right is no longer a philosophical nicety; it is the practical problem of the decade.

So this is not a revision. It is a new book with an old spine.

A word on how to read it, and on what kind of claims it makes.

The book moves in six parts. Part One disciplines the terms and builds the denominator. Part Two forms the quotient and derives surprise in both of its components—magnitude and angle. Part Three reads the machine and estimates consciousness. Part Four reads the arrival and measures intelligence. Part Five turns the whole apparatus on human life, firms, bureaucracies, artificial agents, and the economics of attention. Part Six states the limits, the predictions that could falsify the framework, the problems I have not solved, and the reason Reality persists at all.

Three registers run through the text and I have tried never to let them blur. Some statements are *definitions*: they cannot be false, only useless, and they are judged by what they let you compute. Some are *theorems*: they follow from the definitions, and I have marked the sketchy proofs as sketches. Some are *postulates*: empirical bets that could lose, stated in the open with their boundary conditions attached, precisely so a critic knows where to aim. The most important postulate in the book—that the structural quantity I call coherent spread co-varies with felt experience—is undischarged, is probably not dischargeable from outside, and is labelled as such every time it is used. The hard problem of consciousness is left standing. What is offered is an outside estimate, which is a smaller thing and a computable one.

The mathematics is elementary on purpose. A reader comfortable with complex numbers, logarithms, variance, and entropy has everything required. Where I use a heavier tool—a two-part code, a phase-alignment order parameter—I define it in place. Nothing in this book requires you to take a formalism on faith, and every central quantity comes with a protocol for estimating it from outside, without any appeal to what the thing being measured feels.

One last thing. The framework is unusual among theories of mind in that it is not primarily about minds. The Reality Equation is written for a bird, a rock, a firm, an intersection, a bureau, a market, a language model, and a person, in the same notation, with the same terms, and it returns different numbers for each because they are different machines—not because we have quietly changed the subject. That generality is the point. If the equation only worked on humans it would be an anthropology. It is meant to be a physics of the felt.

J.R.
Charleston, South Carolina
July 2026

Part I

The Terms

Chapter 1

Why Reality Is a Quotient

Two people stand in the same room and watch the same thing happen. Afterward one of them is elated and the other is devastated. Nothing about the event differed. The event is not in dispute. What differed is what each of them brought to it.

The ordinary habit of thought handles this by saying that they *interpreted* the event differently, or that they had different *perspectives* on it, and then treats interpretation as a soft, downstream, slightly unserious thing that happens to people after reality is finished being real. This book takes the opposite view, and it takes it literally rather than poetically.

> *Reality is not the event. Reality is the quotient of the event and what was expected of it.*

Write it plainly. For an entity, with respect to something that entity is tracking,

$$R = \frac{A}{E},$$

where A is Actual—what arrived—and E is Expectation—the standing structure the arrival landed in. R is Reality: not the world, and not the opinion, but the relation between them, which is the only thing anyone has ever actually inhabited.

The equation is trivial to write and hard to hold. It is hard because the word "reality" has been doing two incompatible jobs in ordinary speech for a very long time. Sometimes we use it to mean the immovable record of what occurred—*face reality*, *the reality of the situation*. Sometimes we use it to mean the felt texture of a life—*his reality is very different from mine*. Those two usages are not sloppy variants of one idea. They name the numerator and the quotient, and collapsing them is the single most expensive error in this subject.

1.1 What the division buys

Consider what changes the moment you write a division instead of an identity.

The same numerator can produce any Reality. A raise of four percent divided by an expectation of two percent is a good day. The same four percent divided by an expectation of ten is a catastrophe. Nothing about the money moved. The money is the numerator; the day is the quotient.

Reality can change without the world changing. This is the fact that makes the equation worth its trouble. A person can be told something about an event that occurred a decade ago, and their present Reality can invert, instantly, with the Immutable Past untouched. Denominators are revisable. Numerators are not.

Reality has a fixed point. When $A = E$, $R = 1$. The world came in exactly as the structure expected it. That is not bliss and it is not despair; it is the flat, unmarked condition in which nothing needs to be registered at all. Almost all of experience is spent very close to it, which is why almost all of experience is invisible.

Reality has a direction of correction. Because it is a ratio, R carries sign in the sense that matters: $R > 1$ means the world exceeded the structure, $R < 1$ means it fell short. The equation does not merely say "there was a mismatch." It says which way.

1.2 Why not subtraction

A reasonable objection arrives early: why a ratio and not a difference? Why A/E instead of $A - E$?

Because experience is scale-relative and difference is not. A one-dollar shortfall on a one-dollar expectation is a total failure; a one-dollar shortfall on a million-dollar expectation is a rounding error. Subtraction reports the same number for both. Division reports 0 and 0.999999, which is what the entity in question would report as well.

Because the fixed point of subtraction is 0 and the fixed point of division is 1, and 1 is multiplicatively neutral. Expectations compose. Predictions of predictions compose. Ratios chain; differences do not.

And because the logarithm of a ratio is the natural currency of surprise, as Chapter 11 will derive rather than assert. The moment you write a ratio, the magnitude of the felt correction is already implicit in it, waiting to be taken out with a ln.

1.3 The doorway

Everything else in this book is bookkeeping on that one move. The rest of Part One disciplines the numerator so it cannot smuggle in wishes, and constructs the denominator so it can hold what a real entity actually brings. Part Two performs the division honestly. Parts Three and Four ask, respectively, what kind of machine built the denominator and what became of the numerator after it arrived.

But the doorway is here, and it is one sentence wide:

> *The event is not the experience. The experience is the event divided by what the experiencer already was.*

CHAPTER 2

Actual, Ideal, Real, and Reality

Four words. In ordinary use they are near-synonyms. In this framework they name four different objects, and the entire apparatus fails if they are allowed to collapse into each other. This chapter separates them and gives a single picture that holds all four at once.

2.1 Actual

Actual is what happened. It is the record. It is finished, exact, and not up for negotiation, because it belongs to the Immutable Past—the complete, closed archive of events that no longer contain uncertainty. Actual is not generous and not cruel. It does not intend. It arrives.

Crucially, Actual is not produced by the entity whose equation it enters. This is the asymmetry that makes the equation a description of a creature in a world rather than a creature in a mirror. The denominator is yours. The numerator is not.

2.2 Ideal

Ideal is a pure form: a possible endpoint, complete in itself, unconditioned by circumstance. Justice, in the abstract. A perfect fifth. The thing the sketch is a sketch of. Ideals do not occur. They are not events and they are not located anywhere in the Immutable Past. They are what actualization is *of*.

The temptation is to treat ideals as noble fantasies—nice to have, not real. That is exactly backwards for our purposes. Ideals do work in this framework. They constitute the imaginary component of the denominator, they orient entities toward futures that have not occurred, and they are the reason Reality does not settle into a permanent 1. But they do their work by *orienting*, never by *arriving*.

2.3 Real

Real is the domain of what can actualize: matter, energy, bodies, institutions, instruments, hands, contracts, code. The Real is where ideals get made into actuals, and it is the only place that transaction can happen. A symphony is ideal; the score is real; last Tuesday's performance is actual.

The Real is also the domain of *actualizers*—entities capable of taking an ideal and putting it into the record. That is a role, not a species. A person is an actualizer. So is a firm, a court, a factory, a compiler.

2.4 Reality

Reality is the quotient. It is not a domain at all; it is a relation, and it is exactly and only what the equation says it is. It is the felt subjective experience a witness is having in the eternal now, and it is unstable in the precise sense that it moves whenever either term moves.

2.5 The circle

Draw a circle. Let its interior be the Real: the field of the actualizable. Let a single point inside it, marked and fixed, be an Actual: one thing that got made. Let the circumference be the boundary of possibility—the unit circle of future possibility—and let each point on it be an Ideal, a pure direction that has not occurred and by itself never will.

Now the four terms sit in one picture without touching. Ideals are on the rim; they have angle but no location. The Real is the interior; it has extent. Actuals are points in the interior that have been fixed and cannot move. Reality is not drawn on the picture at all, because Reality is a relation between the point that arrived and the structure that was pointed at it.

This figure will do a lot of work later. In Chapter 6 the rim becomes the space on which idea-vectors live, and their weighted resultant becomes the imaginary component of Expectation. In Chapter 7 the rim's geometry—diameters, poles, cancellation—becomes the structure of belief. In Chapter 23 the fineness with which an entity can distinguish points on that rim becomes a bound on its intelligence. Keep the circle.

2.6 The four errors

Each collapse of these terms is a specific, diagnosable mistake.

- **Actual** = **Reality.** The error of the hard-headed: believing that because the facts are settled, the experience is settled. It is a denial of the denominator, and it makes every disagreement look like a factual dispute.
- **Ideal** = **Real.** The error of the utopian and the ideologue: treating a pure form as though it were already in the field of the actualizable, and therefore treating the work of actualization as a formality that other people are obstructing.
- **Real** = **Actual.** The error of the fatalist: treating the entire field of what could be made as though it were exhausted by what has been made.
- **Reality** = **Ideal.** The error of the aggrieved: taking one's own denominator as a description of the world, so that every shortfall is an injustice rather than a quotient.

Four terms, four errors, one circle.

CHAPTER 3

Actual as Declared Scalar

The numerator has to be hardened before the division can be trusted. This chapter does the hardening, and adds the one clause the first edition lacked: Actual is indexed to an observable.

3.1 Five properties

Actual is **positive**. In this framework there is no negative actualization. Nothing that failed to occur occupies the numerator; nothing that occurred occupies it with a minus sign. Bad things are not negative Actuals—they are Actuals landing in denominators that expected otherwise. Locating badness in the numerator is the most common way people accidentally make the equation say nothing.

Actual is **scalar**. Whatever multidimensional richness an event had in the world, what enters the numerator is one magnitude on one axis. This is not a simplification imposed for convenience; it is forced by the next property.

Actual is **indexed to an 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 whatever that entity is tracking:

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

The hawk's shadow that is Actual for the bird is not Actual for the rock beside it. Same world, same instant, different projections, different numerators. This single clause dissolves a great deal of confusion. Two people did not experience "the same event"; they experienced two projections of one world onto two different observables, and then divided each by a different denominator. No wonder they disagree.

Actual is **declared**. Someone or something fixes the observable and the scale. This is not a weakness in the theory; it is what makes every quantity in the book estimable. All claims in this framework are of the form "with respect to O." The unqualified question—is this entity conscious, is this system intelligent, how surprising was that—has no answer, and the theory says so out loud instead of quietly returning a number.

Actual is **morally neutral and residue-free**. It carries no charge, and it carries nothing from the alternatives that did not occur. The counterfactual is not in the numerator. If you are still carrying the road not taken, you are carrying it in the denominator, which is where it belongs and where it does its damage.

3.2 Making it dimensionless

One technical move, small but load-bearing. Expectation will shortly become a sum of a prediction and something that is not a prediction, and you cannot add a quantity in dollars to an orientation. So fix a reference scale x_0 natural to the observable—a baseline, a just-noticeable difference, or the Actual itself—and work in normalized units:

$$\tilde{A} = A/x_0, \qquad \tilde{P} = P/x_0.$$

From here on the ratio is between pure numbers. Units cancel inside the quotient, which is the correct behavior: a person is not surprised in dollars.

3.3 The discipline

The whole point of hardening the numerator is to force every interesting thing into the denominator, where it can be examined. If Actual is positive, scalar, projected, declared, and residue-free, then every disagreement about an experience, every difference between two witnesses, every change of heart across a decade, is a fact about Expectation. That is not a loss of expressive power. It is a relocation of all the expressive power to the one term that can actually be studied.

Chapter 4

Why Expectation Is Complex

Here is the move the framework is named for, and the one that most often gets waved through without earning its keep. So let it earn it.

4.1 The argument

Suppose Expectation were a single real number—a forecast. Then $R = \tilde{A}/\tilde{P}$ would be a real ratio, surprise would be a real logarithm, and the whole theory would be a restatement of prediction error. It would be tidy, it would be familiar, and it would be wrong, for a reason that shows up immediately in any honest case.

Take a person waiting on a medical result. Ask them what they expect. They will give you a forecast—probably fine, ninety percent. Now ask them what they *want*, what they *fear*, what they are *oriented toward*. They will give you something else entirely, and it will not be a number on the same axis. It has no magnitude on the forecast axis at all. It is not a competing prediction. It is a direction.

That second thing is not noise on the first. It does not average with the forecast and it does not correct it. Two people with identical, accurate, well-calibrated forecasts can have entirely different experiences of the same result, and the difference is not in what they predicted.

So the denominator has two components that cannot be added on one line. That is precisely the situation complex numbers were invented for. Write

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

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

4.2 What the *i* is doing

The imaginary unit here is not decoration and not a metaphor about "imagination." It is doing the one job it always does: holding a quantity that is real, that combines with the first quantity, and that cannot be reduced to it. The two components are orthogonal in the strict sense. No amount of predictive accuracy converts into ideational orientation, and no amount of orientation improves a forecast.

Three immediate consequences, each of which will be cashed out later:

Expectation has magnitude and angle. $|E| = \sqrt{\tilde{P}^2 + B^2}$ and $\arg E = \arctan(B/\tilde{P})$. An entity is not merely expecting *how much*; it is expecting *from an angle*.

Reality inherits the complexity. If E is complex and $\tilde{A}$ is a pure number, then $R = \tilde{A}/E$ is complex, and it has both a magnitude and a phase. Chapter 10 forms it; Chapter 11 reads it.

Surprise acquires a second dimension. The logarithm of a complex number has a real part and an imaginary part. The first is magnitude-surprise: how far off the world was. The second is angular surprise: from what direction the world corrected you. Anyone who has been blindsided by something they had also predicted knows these are different experiences.

4.3 One constraint

The denominator must be nonzero: $E \neq 0$. An entity with zero Expectation—no prediction and no orientation—has an undefined Reality, and that is the correct output, not a bug. Something with no denominator is not having an experience. It is being passed through, which is a state the framework has a name for (Chapter 26) and a score for (zero).

CHAPTER 5

The Real Component: The Prediction Machine

The real part of the denominator is produced by a machine, and the machine is the most consequential object in this book. This chapter introduces it in its first-edition role—the thing that generates $\tilde{P}$—and flags what Part Three will do with it.

5.1 Always on

Every entity we choose to model has a predictive structure that continuously estimates what is about to be declared Actual on its observable. In a person it is subconscious, fast, involuntary, and untiring. It is running now, while you read, predicting the next word, the next syllable, the weight of the object in your hand, the pressure of the chair, the sound the room will make in a second. Almost none of that reaches you, and that is the point: prediction that succeeds is silent.

Four properties matter for the equation.

It is always on. There is no state of the entity in which no prediction is running. Even boredom is a prediction being confirmed.

It resolves to a scalar. Whatever internal richness the machine has—and Part Three is entirely about that richness—what it hands the denominator is one number. $\tilde{P} = G[p]$, a functional of its internal distribution: the mean, the mode, whatever the entity's resolving rule is. The denominator receives the guess, not the machinery.

It is conservative. It does not update on one signal. It is trying to maintain a low, stable prediction error, so it treats a single deviation as noise until repeated Actuals confirm a new pattern. This is why one good day does not change a life and why one betrayal often does not either.

It is trained by action, not by talk. The machine updates on what the entity repeatedly does and repeatedly encounters. Telling yourself something is not an input it respects. This is the least popular sentence in the framework and the one with the most evidence behind it.

5.2 Learning

Learning, in this framework, has a precise definition: *the reshaping of the prediction machine through repeated interaction with Actual.* Not the acquisition of propositions. Not the ability to state a rule. The reshaping of the thing that guesses.

That definition explains a familiar and otherwise puzzling gap. A person can know, propositionally and sincerely, that a fear is unfounded, and be no less afraid. The proposition entered a different system. The machine was not reshaped, because the machine is not addressable by assertion. It is addressable by repetition and consequence, which is what practice, exposure, rehearsal, and lived experience all are.

5.3 What is being held back

Everything in this chapter treats the machine as a producer of one number, which is all the first edition needed. But notice what has been smuggled past: the machine must have had *something* before it produced the guess. A guess is a resolution of alternatives, and you cannot resolve alternatives you do not have.

That something—the distribution the machine holds before it collapses to $\tilde{P}$—is the subject of Part Three, and I now believe it is where consciousness is legible. The denominator contains the guess. The machine contains the cloud.

CHAPTER 6

The Imaginary Component: Ideas

The other half of the denominator. This chapter formalizes ideas, and in doing so commits the framework to a metaphysics that some readers will resist. I state it plainly so the resistance can be aimed properly.

6.1 Ideas are not invented

In this framework, no system—biological or synthetic—invents an idea. Ideas are pre-existing endpoints on the unit circle of future possibility. Each has a direction, an angle, pointing toward a possible orientation in the unknowable future. What systems do is *actualize, expose, combine, and resolve* ideas, always at finite resolution.

This is a strong claim and it is doing real work later—it is why intelligence in Part Four can be measured as *resolution* rather than as generation, and why creativity turns out to be recomposition over a well-factored dictionary rather than creation *ex nihilo*. A reader who rejects the metaphysics can still use the machinery: substitute "the space of orientations available to this entity" for "the circle of pre-existing ideas" and every equation survives. But I do not think the substitution is necessary, and I do not make it.

6.2 Magnitude, not count

The naive move is to say that the imaginary component is proportional to how many ideas an entity holds. It is not, and the reason is geometric. Ideas are vectors on the unit circle. If an entity weighted every idea equally, the vector sum would cancel toward zero—a mind holding all directions at once points nowhere, and its imaginary component vanishes.

What matters is the *weighted resultant*:

$$V_I = Be^{i\theta},$$

with magnitude B and argument θ. The magnitude is how much orientation the entity actually has after cancellation. The argument is which way it points.

So the imaginary component is not a measure of how ideational someone is. It is a measure of how much their ideation *fails to cancel*. A person full of contradictory commitments and a person with no commitments have the same B, for opposite reasons, and the framework is right to score them alike—neither orientation moves their denominator.

6.3 Conditioning

Ideals are unconditioned; ideas are ideals as held by a particular entity, and therefore conditioned by everything about that entity: history, body, culture, language, injury, incentive, and the accumulated record of what has previously worked. Two people oriented toward the "same" ideal hold different idea-vectors because the conditioning differs. That is not a failure of either to grasp the ideal. It is what holding an ideal in a real actualizer consists of.

6.4 Angle and magnitude have different jobs

A distinction worth planting now, because it organizes the last two parts of the book.

For the theory of consciousness in Part Three, what matters about the ideational field is not θ but the *coherence* of the structure—whether the components are phase-organized or smeared. For the theory of Ideation, sketched in Chapter 38, what matters is precisely θ: the net angle a system commits to.

Consciousness is the coherence of phase across what is held.
Ideation is the net phase committed to.

Two readings of one complex object. Chapter 38 shows the seam.

CHAPTER 7

Symmetry, Polarity, and the Unit Circle

The ideational field has geometry, and the geometry is not decorative. It is what makes belief, opposition, fairness, and cancellation computable rather than merely discussable.

7.1 Wholes and poles

Take the circle of possibility again. Certain ideas come in *polarity pairs*: two endpoints separated by π, opposite ends of a diameter. Mercy and severity. Freedom and order. Risk and safety. Loyalty and honesty. These are not contradictions in the logical sense—nothing about mercy makes severity false—and treating them as contradictions is the characteristic mistake of unstructured thought.

They are *poles of one symmetrical whole*. The diameter is the whole; the two endpoints are its poles; and the whole is only visible if you can hold both endpoints at once. Fairness, in this framework, is the diameter itself—not a point on the circle but the axis through it, the capacity to hold a pair without collapsing to one side.

7.2 Cancellation

The geometry has an immediate and brutal consequence. If an entity holds two poles with equal weight, they cancel:

$$we^{i\theta} + we^{i(\theta+\pi)} = 0.$$

Perfect balance produces zero orientation. This is why studied neutrality feels like nothing—because in the denominator it *is* nothing. It contributes no imaginary magnitude, and an entity in perfect balance across every pair has $B = 0$ and an Expectation that is purely predictive.

This is not an argument against fairness. It is a precise account of its cost. Fairness holds the diameter; orientation requires an angle; you cannot have maximal fairness and maximal orientation on the same pair at the same time. Every honest person has felt this trade and now it has coordinates.

7.3 Structured opposition

The useful case sits between the two disasters. Collapse to one pole and the entity has orientation but no whole—maximum B on a pair it can no longer see. Balance perfectly and it has the whole but no orientation. What a well-formed entity does is hold the diameter *and* lean:

$$V_I = w_1 e^{i\theta} + w_2 e^{i(\theta+\pi)}, \qquad w_1 > w_2 > 0,$$

which yields a nonzero resultant of magnitude $w_1 - w_2$ pointed at the leaned-toward pole, while both poles remain represented and therefore both remain visible.

That is what it looks like, in vectors, to have a position without being a fanatic. The lean is real; the opposite pole is still in the sum; the resultant is smaller than the fanatic's and it is attached to a structure that can see what it is opposing.

7.4 Why this matters downstream

Two payoffs to flag.

First, the ability to represent both poles of a pair is a form of *distinguishability*, which is exactly what Part Four's admissibility tests require of a component and what Part Three's coherence term measures across a cloud. Structured opposition is not merely a virtue; it is a high-resolution representation, and it scores as one.

Second, the angular resolution of the circle—the finest distinction between directions an entity can register and stably maintain—turns out to bound its intelligence (Chapter 23). Someone who can only see one axis at a coarse resolution is not merely biased. They are, measurably, holding fewer distinguishable states.

CHAPTER 8

Truth, Falsity, Ignorance, and Bias

With the geometry in place, four familiar words become operational. Each is a particular relation between a host and a polarity pair.

8.1 The four positions

Let a pair have poles p^+ and p^-, and let a host hold them with weights w_+ and w_-.

Truth is correspondence between the held structure and the pair as it actually stands—the host's weights track the real asymmetry of the whole. Truth, in this framework, is not a property of a sentence in isolation. It is a fit between a representation and a structure.

Falsity is a held weight on a pole that the whole does not support: $w_+ > 0$ where the structure gives no warrant. Falsity is not the absence of truth; it is presence in the wrong direction, and it is therefore expensive. It contributes real magnitude to B and points it wrong.

Ignorance is absence of representation: $w_+ = w_- = 0$, the pair simply not held. Note what the geometry says—ignorance and perfect balance both contribute nothing to the resultant, and are indistinguishable in the denominator. They are distinguishable in the machine (Part Three) and in the decomposition (Part Four), where the ignorant entity has no components for that axis and the balanced entity has two.

Bias is systematic asymmetry in the weighting that does not derive from the structure: a lean the world did not put there. Bias is not a moral category here. It is a measurable displacement of the resultant, and it has a specific, diagnosable consequence in Chapter 12—it is one of the two independent sources of surprise.

8.2 Bias is not error

This is the sharpest of the distinctions and the most useful in practice.

A prediction error is a failure of the real component: the machine guessed wrong. Bias is a displacement of the imaginary component: the orientation was leaned. They produce different surprises, they are corrected by different things, and confusing them wastes lives.

You correct a prediction error by exposure—more Actuals, repeated, until the machine reshapes. You cannot correct bias that way, because bias is not a mis-estimate of frequency; it is a weighting on a pole. A person who is systematically disappointed by others is not necessarily bad at predicting others. They may be predicting others accurately and holding a leaned denominator, in which case every additional accurate prediction leaves the disappointment exactly where it was.

Chapter 12 gives this the formal treatment: surprise decomposes by source, and the decomposition is diagnostic.

8.3 The host's responsibility

One consequence deserves to be said without hedging. Because bias lives in the denominator and the denominator is the entity's own, the framework locates the revisable part of experience squarely in the entity. Not the event—the event is closed. Not the prediction—the prediction is a machine that answers to repetition, not to will. But the weighting on the poles is held, and what is held can be examined, named, and re-weighted.

That is a narrow doorway. It is also the only one there is.

Chapter 9

Why Ideas Choose People

The last chapter of Part One reverses a direction that most readers will have been assuming, and the reversal is not rhetorical.

9.1 The reversal

The ordinary picture: a person has an idea, holds it, and if they are diligent, actualizes it. The idea is passive; the person is the agent.

The framework's picture: ideas are pre-existing endpoints seeking actualization, and hosts are *sites of compatibility*. An idea is actualized through a host whose structure can carry it—whose vocabulary can resolve it, whose orientation points near it, whose real capacities can make it, and whose conduct sustains it long enough for the making to finish. On this reading the selection runs the other way. The idea does not belong to the person. The person is where that idea could happen.

Anyone who has done sustained creative or entrepreneurial work will recognize the phenomenology—the sense of being used by a thing rather than using it, the way an idea will not leave, the way it arrives in several people at once and gets made by the one who was built to make it. The framework does not treat that phenomenology as a charming illusion. It treats it as a report about the direction of selection.

9.2 Two conditions

For an idea to select a host, two things are jointly necessary, and neither alone suffices.

Belief is orientation: the host's idea-vector must have real magnitude pointed near the idea's angle. A host who holds the idea at balance, or does not hold it at all, contributes nothing—by the cancellation geometry of Chapter 7, literally nothing.

Works are actualization capacity: the host must be in the Real with the means to make. Orientation without capacity is a vector on a rim with nothing inside the circle to receive it.

The pairing is old and the framework re-derives it rather than borrowing it. Belief without works has angle and no actualization. Works without belief has capacity pointed nowhere in particular, which is to say capacity that will be spent on whatever has the nearest angle—usually someone else's idea.

9.3 What this buys the rest of the book

Three things.

It explains **compatibility** rather than merit. The question is not who deserves an idea but whose structure can carry it, which is why extraordinary ideas routinely pass through unremarkable people and why extraordinary people are routinely useless to particular ideas.

It explains **persistence**. If ideals sought nothing, actualization would be optional and Reality would settle. Chapter 39 closes the book on exactly this point: Reality persists because ideals have not yet achieved perfect actualization through real actualizers, and the quotient stays off 1 as long as that is true.

And it sets up the theory of **Ideation** in Chapter 38. If consciousness is the coherence of a held cloud, and ideation is the net angle committed to, then an idea selecting a host is a claim about which cloud can resolve toward which angle. The metaphysics and the mathematics turn out to be describing the same event.

Part II

The Quotient

CHAPTER 10

Forming the Quotient

The terms are disciplined. Now perform the division, and—this is the entire difficulty—do not flinch at the end of it.

10.1 The operation

With $\tilde{A}$ a positive pure number and $E = \tilde{P} + iB$ complex and nonzero,

$$R = \frac{\tilde{A}}{\tilde{P} + iB} = \frac{\tilde{A}(\tilde{P} - iB)}{\tilde{P}^2 + B^2} = |R|e^{i\varphi},$$

with

$$|R| = \frac{\tilde{A}}{\sqrt{\tilde{P}^2 + B^2}}, \qquad \varphi = -\arctan\left(\frac{B}{\tilde{P}}\right).$$

Reality is complex. It has a magnitude—how much the world exceeded or fell short of the structure that received it—and a phase—the rotation between where the entity was pointed and where the world came from.

10.2 The temptation to scalarize

Every reader wants, at this moment, to take $|R|$ and throw φ away. Magnitude is intuitive. Magnitude is what "better or worse than expected" means. Phase is strange, and there is no folk vocabulary for it.

Do not do it, and here is the cost of doing it.

Consider two entities with identical Actual and identical $|E|$, so identical magnitude of Reality. The first has a large predictive component and almost no orientation: $E \approx \tilde{P}$, $\varphi \approx 0$. The second has the same $|E|$ built almost entirely from orientation: $\tilde{P}$ small, B large, $\varphi \to -\pi/2$. Scalarize and these are the same experience. They are not remotely the same experience. The first is a well-calibrated forecaster who got a result. The second is someone who was pointed at a future and had the world arrive from ninety degrees off their axis.

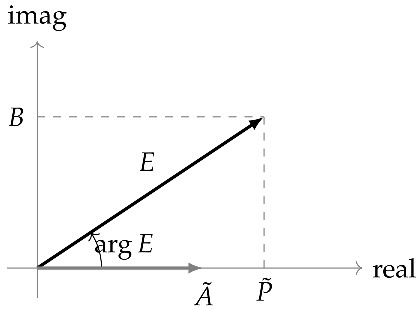

Figure 10.1: Expectation in the plane. The real axis carries the prediction machine's resolved guess; the imaginary axis carries the weighted resultant idea-vector. Actual is a positive scalar on the real axis. Reality is the quotient of the two, and therefore has both a magnitude and an angle.

We have ordinary words for the second condition and they are all about direction: *blindsided*, *from nowhere*, *I never saw it coming*—said, often, by people whose forecasts were fine. Angular surprise is not an exotic addition to the theory. It is the part of experience the scalar theories have been silently discarding.

10.3 Four readings

The four corners of the quotient, each a recognizable condition.

$\tilde{A}$ **large,** $|E|$ **small.** Reality far above 1. The world exceeded a thin structure. Elation, windfall, the disproportionate joy of someone who expected little. Note that this is achievable by lowering the denominator, which is the mechanism behind every ascetic tradition and every piece of advice about expectations that has ever been given.

$\tilde{A}$ **small,** $|E|$ **large.** Reality far below 1. Disappointment, and specifically the disappointment of the well-prepared. The structure was large. The arrival was not.

$\tilde{P}$ **dominant,** $B \approx 0$**.** $\varphi \approx 0$. A purely predictive entity: it can be right or wrong in magnitude and it cannot be blindsided in direction, because it was not pointed anywhere. Instruments live here. So do people in certain states, and they describe those states as flat.

B **dominant,** $\tilde{P}$ **small.** $\varphi \to -\pi/2$. A purely oriented entity: enormous angular exposure. Every arrival is a rotation. This is the signature of the visionary and of the delusional, and the framework does not distinguish them here—the distinction shows up in Part Four, as whether the components being held pay for themselves.

10.4 The rule

> *Preserve the quotient as complex. Take the magnitude when you want to know how much. Take the angle when you want to know from where. Take both when you want to know what happened.*

Chapter 11

Surprise as a Complex Logarithm

Surprise is not a primitive in this framework and it is not defined by introspection. It is derived, and derived in one step.

11.1 The derivation

We want a scalar measure of how far Reality departed from the neutral condition $R = 1$. Four requirements fix it almost uniquely.

It must vanish at $R = 1$. It must be signed, so that exceeding and falling short are opposite rather than merely different. It must be symmetric under inversion, so that doubling and halving are equal and opposite deviations—twice as much as expected and half as much as expected are the same size of event, in opposite directions. And it must be additive over composed ratios, so that surprises chain.

The function satisfying all four is the logarithm. So

$$S = \ln R.$$

And because R is complex, so is S:

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

taking the principal branch, $\varphi \in (-\pi, \pi]$, so that angular surprise is the smallest signed rotation carrying Expectation onto Actual.

That is the whole derivation. Surprise is the logarithm of the quotient because the logarithm is the only function that behaves the way surprise is observed to behave.

11.2 Reading the two parts

Magnitude-surprise, $\ln|R|$, is the familiar quantity. Positive when the world exceeded the structure, negative when it fell short, zero at the fixed point. It is unbounded in both directions in principle and compressed in practice, which is why the difference between a windfall and a larger windfall is felt so much less than the difference between nothing and a windfall.

Angular surprise, φ, is the direction of the correction. It answers a question the magnitude cannot: not *how far off were you* but *from what axis did this arrive*. An entity with a large B has a large standing angular exposure; every Actual, arriving as it does on the real axis, rotates it.

Two things worth noticing about φ. It is bounded—rotations live on a circle, and the principal branch caps angular surprise at π; inside the domain of ordinary validity, where $\tilde{A} > 0$ and $\tilde{P} > 0$, it is capped more tightly still, at $\pi/2$. And it is *signed by which side of your orientation the world came from*, which is why being wrong in the direction you feared feels categorically different from being wrong in the direction you did not.

11.3 Zero surprise is not nothing

$R = 1$, $S = 0$: Actual matched Expectation. It is worth pausing on how much of a life is spent there, and what that means.

The overwhelming majority of predictions succeed. The floor holds. The word is the expected word. The person behaves as the person behaves. Every one of those is a Reality of exactly 1 and a surprise of exactly 0, and none of them are registered, because there is nothing to register.

This has a consequence that Part Four will make quantitative: *successful prediction is invisible, and therefore the felt content of a life is disproportionately made of its residue*. What you notice is what your machine did not absorb. This is not a complaint about attention. It is the definition of attention, and Chapter 29 states it as such.

11.4 What surprise is not

Two guardrails, both of which the framework will lean on hard later.

Surprise is not the measure of consciousness. A wide, well-organized machine may predict accurately, so its Actual arrives near its Expectation and its surprise is low—and it may be highly conscious nonetheless, because consciousness is read from the machine that produced the guess, not from the correction that followed it. This decoupling is load-bearing and Part Three will not be coherent without it.

Surprise is not residual energy. A large residual that the model correctly treats as high-entropy generates little surprise—surf noise to a surfer, traffic to a commuter. A small residual in a direction the model held nearly certain generates enormous surprise. The map from what was left over to what is felt runs through the model's confidence structure, which is to say through the cloud. That is the seam where Part Four hands off to Part Three, and Chapter 24 marks it explicitly.

Chapter 12

Sources of Surprise

The scalar is not enough by itself, because two entities can post identical surprise for entirely different reasons, and the reasons determine what, if anything, can be done.

12.1 The decomposition

Surprise arises from a mismatch between $\tilde{A}$ and E. Since E has two components, mismatch has two independent sources.

Prediction error is a mismatch in the real component: $\tilde{P}$ was wrong. The machine's guess did not track the world. This is the source that responds to exposure—repeated Actuals reshape the machine, the guess improves, the surprise falls. It is, in the ordinary sense, learnable.

Ideational bias is a displacement in the imaginary component: B and θ were leaned. The orientation was pointed somewhere the structure did not warrant. This source does *not* respond to exposure, and that is the whole point of separating it. More accurate prediction does not reduce it. You can be perfectly calibrated and reliably disappointed, and the framework says exactly why: your $\tilde{P}$ is fine and your B is leaned.

Compound surprise is both at once, which is the common case, and the reason diagnosis is necessary rather than optional.

12.2 Diagnostics

Three questions separate the sources in practice.

Did the world do something my history did not contain? If yes, there is prediction error in the mix. The machine had no basis and could not have guessed.

Would a well-informed stranger have been surprised? If no—if the surprise is idiosyncratic to you—the source is largely ideational. The world behaved as the structure warranted and your weighting was leaned.

Does the surprise recur with the same shape across different events? If yes, it is bias. Prediction errors are distributed; bias is systematic. A repeating shape is a fact about the denominator, not about the world.

12.3 Why this is the practical core of the book

Nearly all of what people call working on themselves is, in this framework, one of two operations that are constantly confused for each other.

Reducing prediction error is a matter of *exposure*: put the machine in contact with the Actuals until it reshapes. It is slow, it is not verbal, it does not respond to resolve, and it works.

Reducing bias is a matter of *re-weighting a pole*: examining a held asymmetry that the structure does not warrant and holding it differently. It is available to examination in a way the machine is not, and it is not fixed by exposure at all.

The failure mode is universal and expensive: people attempt to fix bias by gathering evidence, and attempt to fix prediction error by deciding to feel differently. Both fail, both feel like effort, and the framework predicts the failure of each with the same equation.

12.4 An honest limit

The decomposition is clean in the formalism and difficult in practice, because $\tilde{P}$ and B are not separately observable from inside. What is observable from inside is the quotient. Separating the sources requires either the diagnostics above, applied honestly, or an outside estimate of the machine—which is exactly what Part Three's measurement protocol is for. The framework does not pretend the decomposition is easy. It insists only that it is real, and that treating it as one quantity guarantees the wrong intervention.

PART III

The Shape of the Machine

CHAPTER 13

The Cloud Behind the Guess

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 part does not try.

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

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

13.1 The move

Chapter 5 established that the machine hands the denominator a single number. But a guess is a *resolution*, and resolution presupposes something to resolve. Let $p(x)$ be the machine's predictive distribution over normalized outcomes on the observable—call it the **probability cloud**—and let the resolved guess be a functional of it:

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

The cloud 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.

Three clouds, three conditions. 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 emerged from many distinguishable, resolvable alternatives.

The outside-facing test is therefore not "how wide is the cloud" alone, but "can that width be transformed into distinct, nameable, resolvable alternatives before it collapses into the guess."

13.2 A guard against mysticism

An important clarification, and the one that keeps this from being 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 this book says a rock "has a prediction machine," it means that the best model of the rock's behavior with respect to some observable contains a nearly trivial predictive structure. The quantity we are about to define is a property of a model, estimated from outside—not a claim about rock-experience.

That stance is what keeps the theory operational and falsifiable rather than mystical, and it is imported wholesale into Part Four.

13.3 Entities

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."

The difference between entities is not that some have prediction machines and others do not. It is the *shape* of the machine.

CHAPTER 14

Spread and Coherence

Two axes, and they are independent. Getting that independence right is the difference between this theory and a single-dial theory of consciousness.

14.1 Spread

Spread is the width of the cloud—how much unresolved possibility the machine is holding before it commits. A clean, bounded choice is normalized dispersion,

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

optionally squashed to $[0,1]$ by $\sigma \mapsto \sigma/(1+\sigma)$ for comparison across entities. A delta cloud—perfect certainty—has $\sigma \to 0$. A broad cloud has large σ.

Spread by itself says nothing about organization. Static has enormous spread.

14.2 Coherence, in two computable 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 Chapter 17. Two readings make it computable, and they agree.

Information-theoretic: negentropy at fixed spread

Among all clouds of a given variance, the smoothest—the highest-entropy one—is the Gaussian; a featureless smear is the noise limit. So define

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

where $H(p)$ is the entropy of the cloud. Then $\kappa \to 0$ for an unstructured smear, and $\kappa \to 1$ for a cloud concentrated into sharply separated modes. Coherence is departure from maximal smoothness toward distinct modes—which is precisely what it means to hold distinguishable this-and-thats.

One technical caution. Differential entropy is not scale-invariant, and $H_{\max}(v)$ is negative for $v < 1/(2\pi e)$, so the ratio can leave $[0,1]$ if the cloud is expressed in badly chosen units. The definition is well behaved only once outcomes are normalized so that the cloud's variance is of order one—which is what x_{ref} is for. When that normalization is awkward, use the spectral reading below, which is bounded by construction.

Spectral: phase coherence

Represent the machine's predictive dynamics as a signal and decompose it into components. 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

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

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 $\kappa \to 0$. Music has many components with organized phases, so both σ and κ are high. This is the literal, physical meaning of coherence—phase alignment—and it makes the music analogy exact rather than decorative.

It also gives the imaginary axis of Expectation a role in consciousness after all, since phase is angle. That is the seam along which the theory of Consciousness and the theory of Ideation are joined in Chapter 38.

14.3 Two readings, one quantity

The two definitions agree in spirit: structured-given-spread is the same thing as phase-organized. Either can serve as κ, and which one you use is a matter of what you can measure on the system in front of you—entropy of a state distribution for a brain or a market, phase relations for anything you can treat as a signal.

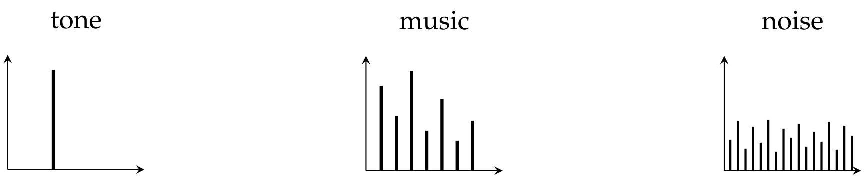

Figure 14.1: The spectral reading of the cloud. *Tone*: one component, so $\sigma \to 0$. *Music*: several phase-organized components, so σ and κ are both high. *Noise*: many components with random phases, so the vectors cancel and $\kappa \to 0$.

CHAPTER 15

Consciousness as Coherent Spread

15.1 The measure

Apparent consciousness is the product:

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

The product form captures both failure modes at once. If $\sigma \to 0$ then $C_{\text{ext}} \to 0$: collapse. If $\kappa \to 0$ then $C_{\text{ext}} \to 0$: 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.

> *Consciousness is organized uncertainty. It is the felt geometry of prediction.*

And it gives *choice* a definition that neither inflates nor dissolves it: choice is the felt resolution of coherent spread into one guess, one action, one direction. Not necessarily metaphysical free will—the felt selection of this over that, here over there, approach over avoid. A machine with no spread has nothing to select among. A machine with no coherence has nothing selectable.

15.2 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, which is plainly wrong—there is such a thing as too much.

The downturn is real, and it is a theorem rather than a picture. Coherence fails once spread outruns the machine's capacity to organize it.

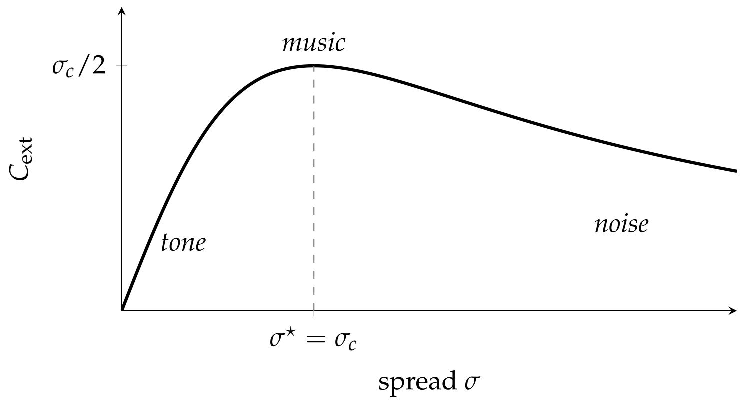

Figure 15.1: The consciousness curve. With finite organizing capacity σ_c, coherence falls once spread outruns capacity, so $C_{\text{ext}} = \sigma\kappa(\sigma)$ rises, peaks at $\sigma^\star$, then decays: collapse $\rightarrow$ coherent spread $\rightarrow$ incoherence.

Let σ_c be that organizing capacity and write

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

Differentiating gives an interior maximum at

$$\sigma^\star \;=\; \sigma_c\,(m-1)^{-1/m}, \qquad m>1,$$

and the restriction matters: for $m \leq 1$ the derivative never changes sign and C_{ext} rises monotonically. The inverted-U is a theorem about a machine whose coherence degrades *superlinearly* in excess spread—which is what finite organizing capacity means. so for $m = 2$ the peak sits at $\sigma^\star = \sigma_c$ with $C_{\text{ext}}^{\max} = \sigma_c/2$.

The content of that result is striking, and it is the sharpest claim in this part of the book:

> *Consciousness is maximized when spread is pushed right up to the edge of the machine's ability to keep it coherent.*

Below the edge lies collapse—tone. Above it lies incoherence—noise. At the edge lies coherent spread—music. This is an edge-of-chaos statement, and Chapter 35 places it against the literature on criticality.

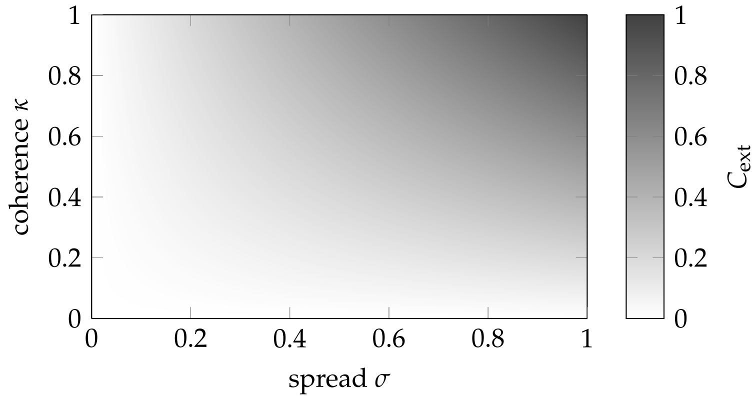

Figure 15.2: The consciousness landscape, $C_{\text{ext}} \approx \sigma \times \kappa$. Collapse lives on the left edge, incoherence on the bottom edge; coherent spread—high apparent consciousness—lives in the upper-right interior.

15.3 The landscape, and the curve as a slice

Spread and coherence are two independent axes, not one, and the inverted-U is only what you see when you walk along the capacity constraint $\kappa = \kappa(\sigma)$. The full object is a surface. It is low along both edges—collapse on the left ($\sigma \to 0$), incoherence along the bottom ($\kappa \to 0$)—and high only in the coherent-spread interior.

Placing entities as points (σ, κ) makes the theory's predictions visual, and it 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.

15.4 Why coherent spread is stable

An apparent paradox. If the machine minimizes prediction error, does it not narrow its cloud toward collapse, and therefore 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; it is the one whose coherent spread matches the structured complexity of its environment. 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.*

Chapter 16

The Surprise Landscape

One might worry that the cloud has been cut loose from the equation. Only the resolved guess enters the denominator; if consciousness is read from the cloud, is it not floating free of the Reality Equation altogether?

It is not, and the connection is exact.

16.1 Counterfactual surprise

The cloud is 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:

$$S(a) \;=\; \ln\!\left(\frac{\tilde{a}}{E}\right), \qquad a \in \operatorname{supp}(p).$$

Two summary statistics recover the two axes from outside:

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

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

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 distinguishable hypotheses.

> *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.*

16.2 Why this matters

Consciousness is legible in the shape of this response surface, and—this is the practical payoff—the response surface is probeable without any access to felt experience. You can perturb inputs, observe how the system's error signal varies, and read whether the landscape has basins or is flat. That is an experiment, not an interpretation.

It also produces the theory's most counterintuitive prediction, stated formally in Chapter 36: 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. The expert who is rarely surprised is not thereby less awake; her surprise landscape is more finely carved, not flatter.

CHAPTER 17

The Fuller Measure and a Protocol

17.1 Three modulators

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

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

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—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.

17.2 The protocol

The theory is only worth scrutinizing if C_{ext} can, in principle, be estimated from outside. Here is the recipe, and every step is executable 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 and coherence κ—negentropy at fixed spread, or the spectral phase-order parameter.
5. Aggregate $C_{\text{ext}} = \sigma\kappa$, optionally modulated by G, Λ, and K.
6. Cross-check against the surprise landscape $S(a)$: a coherent-spread machine should exhibit a structured, multi-basin landscape; a collapsed or incoherent one should not.

The numbers so obtained are model-confidence-limited for entities we cannot directly interrogate—and Chapter 37 says so as an open problem rather than burying it—but the *procedure* is the same in every case, which is what makes C_{ext} a candidate for genuine comparison across entities.

CHAPTER 18

Consciousness Is Relative to an Observable

Because Actual is indexed by an observable, so is consciousness. Write $C_{\text{ext}}(X \mid O)$: the apparent consciousness of entity X with respect to what it is predicting.

This is not a hedge. It is a feature, and it recovers how the word is actually used. The same bird may sit near the peak on $C_{\text{ext}}(\cdot \mid \text{predator motion})$ and near collapse on $C_{\text{ext}}(\cdot \mid \text{next winter})$. "Consciousness of what" becomes a first-class question. A single scalar for a whole entity is then a chosen aggregate over observables, not a primitive—and any such aggregate must publish its weighting, exactly as a price index must.

18.1 The economic reading, made precise

Observable-relativity is what turns a metaphor about corporate behavior into a measurement.

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 $C_{\text{ext}}(\text{firm} \mid \text{revenue})$ is correspondingly high.

A **rentier firm** 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 $C_{\text{ext}}(\text{firm} \mid \text{revenue})$ is low *by design*.

The same firm may still be highly conscious with respect to litigation, interest rates, or political risk. That is the point of the index.

> *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, and it is measurable by the protocol of the previous chapter.

CHAPTER 19

A Catalogue of Entities

Illustrative placements, meant to display the structure of the theory. These are not measurements; the protocol is how real numbers would be earned.

19.1 The cast

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. Highly actual and minimally conscious—low 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.

The human being. A highly spread-out and coherent machine, guessing across seconds, years, and decades—identity, love, death, money, reputation, meaning, children, work, legacy, history. Its horizon is vast and 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.

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.

The incoherent entity. The other way to be low: high spread, low coherence. Not collapse but noise—many unresolved signals that cannot

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

Table 19.1: Illustrative scorings, $C_{\text{ext}} = \sigma \times \kappa$. The two low-$C_{\text{ext}}$ clusters sit at opposite corners of the landscape: collapse (low σ) and incoherence (low κ).

be named, differentiated, or resolved; a machine that cannot perform the transform that turns spread into music.

19.2 Worked scorings

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, with collapsed differentiation, σ low—which on this theory is precisely why hypersynchrony coincides with the *loss* of consciousness rather than its intensification.

19.3 The two failure modes, with signatures

The theory earns its keep by predicting two distinct ways to lose consciousness, which single-axis theories tend to blur. These 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. The generalized seizure is the extreme case.

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.

Chapter 20

The Correspondence Postulate

Nothing in the mathematics forces a structural quantity to be *about* felt experience. This is the explanatory gap, and this book 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, not a concession. It tells a critic exactly where to aim. If the postulate is false, everything structural in this part survives as a theory of predictive architecture and nothing in it is a theory of consciousness. If it is true, the outside estimate is finally something one can compute rather than merely assert.

20.1 What the postulate does not say

It does not say that C_{ext} is consciousness. It says the two co-vary where experience exists at all.

It does not say that systems with high C_{ext} have experience. The thermostat has a nonzero C_{ext} and the postulate is silent about it. The theory quantifies the distance between a thermostat and a person instead of hand-waving it, which is the opposite of flattening the two together.

And it does not say that low C_{ext} implies no experience. It says the richness co-varies. A collapsed machine, if it has experience at all, has a thin one.

The hard problem remains. But the outside estimate is now a computation. 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, and consciousness is their product.

Part IV

The Resolution

CHAPTER 21

Intelligence as Decomposition

Part Three read the denominator's machinery. This part turns to the numerator's fate: Actual arrives—then what?

21.1 The black box

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 it 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.*

21.2 Three refusals

The theory's perimeter is defined as sharply by what it refuses as by what it claims.

There is no observable-free intelligence. Every score in this part is written $I(X \mid O)$, and any single-scalar "general intelligence" is a declared aggregate, not a discovered constant.

The theory claims nothing about felt experience. Intelligence, on this account, can be carried by systems for which the consciousness question is empty. Chapter 27 makes the independence formal.

No system invents ideas. As Chapter 6 committed, 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.

21.3 Where it sits in the equation

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.

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 time 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, $\Phi_{\mathrm{num}} \in [0,1]$, 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. Chapter 31 is about nothing else.

21.4 Terminological discipline

Three standing rules, stated once and obeyed throughout, because the notation of the two theories is adjacent and collisions would be fatal.

1. **Expectation** names the denominator and nothing else. Nothing in this part "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 is Surprise. The vocabulary of decomposition is V, never B—B is the imaginary magnitude of Expectation. Component coherence is χ, never κ—κ is cloud-coherence.

21.5 The decomposition identity

System X encounters arrival stream y on observable O and produces

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

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.

This identity is bookkeeping, not yet a theory. The theory is in what qualifies a c_j.

CHAPTER 22

The Five Admissibility Tests

Counting components naively rewards fragmentation, overfitting, and invented structure. So the count is not the foundation. The foundation is a *gain criterion*, and the count is that criterion's integer shadow.

22.1 The 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 stays 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 the class $\mathcal{Y}_O$. 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.

22.2 The two scores

$$\begin{aligned}
N_\varepsilon(X \mid y, O, V) &= \big|\{\, j : c_j \text{ passes T1–T5} \,\}\big| \\
I_{\text{dec}}(X \mid y, O) &= \sum_{j \text{ admissible}} \Delta\ell_j \\
I_{\text{act}}(X \mid O) &= \mathbb{E}_\mu\big[I_{\text{dec}}(X \mid y, O)\big] \\
Q_{\text{abs}}(X \mid y) &= 1 - \ell_X(\rho_X) \,/\, \ell_0(y) \\
L_\rho(X \mid O) &= \mathbb{E}_\mu[\ell_X(\rho_X)] \,/\, T
\end{aligned}$$

respectively: the admissible component count (the public score), total compression gain in bits (the technical score), its expectation over the signal class, the absorbed fraction, and the residual load in bits per unit time.

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 $(O, \mathcal{Y}_O, V, \varepsilon)$.

Reporting N_ε without its context is, within this framework, malpractice, exactly as reporting C_{ext} without its observable would be.

Weighted-count variants are subsumed: $\Delta\ell_j$ *is* the principled weight, because amplitude, stability, and explanatory power either contribute to compression gain on the class or should not contribute to the score at all. 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} in Chapter 31.

22.3 Vocabulary regimes

Component extraction is always relative to a vocabulary V: Fourier atoms, legal concepts, diagnostic categories, musical structures, institutional classifications. Two scoring regimes are permitted, and scores are *never* compared across them.

Conventional-V scoring fixes a community-standard vocabulary, enabling fair comparison of systems within a domain—like scoring chess under the rules of chess.

Free-V scoring lets the system bring its own dictionary and scores it purely by description length, which is vocabulary-fair by construction: any dictionary may be used, and each is charged its own description cost.

Free-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.

CHAPTER 23

Capacity: Potential and Actualized

23.1 Angular resolution and idea-space

Chapter 6 placed ideas as vectors on the unit circle of future possibility, each with an angle θ. Within that 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 usually is not, because the circle is not an imported metaphor: it is an object this 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.

23.2 Capacity in bits

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

$$I_{\text{pot}}(X \mid O) \;=\; \log_2 \left(\begin{array}{l}\text{\# distinguishable component-} \\ \text{states of } X\text{'s vocabulary over} \\ O\end{array}\right),$$

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 "forty bits."

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 $1/136$.

23.3 The bound

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

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 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.

CHAPTER 24

The Remainder

The remainder $\rho_X(y) = y \ominus F(c_1, \ldots, c_N)$ is what this system's decomposition leaves unresolved. Three properties matter, and all three are easy to get wrong.

24.1 It 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 estimates not only the remainder's size but its latent structure: how much of it a finer reference system can resolve. That step is not optional. Without it, a system's resolution limit gets quietly promoted into a claim about the world.

24.2 It 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 Chapter 30 names the consequences.

24.3 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.

24.4 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 and not asserted

The Reality Equation's Surprise is $S = \ln(\tilde{A}/E)$, a complex log-ratio, not a log-probability. The connection this theory 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.

A corollary that protects the theory

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, which is to say through the cloud, which is precisely the seam where this theory hands off to Part Three's counterfactual surprise landscape.

CHAPTER 25

Absorption and Attention

25.1 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 Chapter 5 called learning.

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. What you learned to drive, you no longer drive; the car drives, and you think about something else.

25.2 Two postulates about attention

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

Postulate 2 (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 3 (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 2 is strongly supported by the automaticity and expertise literatures—chunking in expert memory, automatization of procedural skill—with a known boundary condition: explicit-monitoring pressure can force re-attention to absorbed components, typically degrading performance. 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 3 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.

25.3 The expertise signature

The consequence is worth stating plainly, because it inverts the folk picture.

> *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.*

Two disciplined caveats bound the claim. Residual load L_ρ is a *structural* quantity, and whether any of it is *felt* is governed by Part Three's separate machinery and its correspondence postulate. A thermostat with nonzero L_ρ feels nothing this theory can speak to.

Intelligence, attention, surprise, and consciousness remain four distinct quantities throughout this book. Nothing is gained by letting any two of them merge, and a great deal is lost.

CHAPTER 26

Pass-Through and Confabulation

Consciousness fails two ways: collapse, holding nothing; and incoherence, spread without organization. Intelligence also fails two ways, and the symmetry is exact enough to be load-bearing.

26.1 The two modes

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 labelling, 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. And 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. The protocol's null-signal control 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**.

26.2 The confabulation curve

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

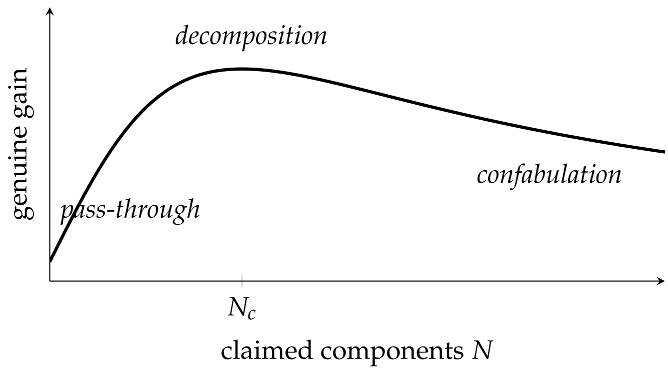

Figure 26.1: The conjectured confabulation curve, the intelligence-side twin of the consciousness curve. The optimum is set by validation capacity N_c rather than organizing capacity σ_c.

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 Part Three 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.

26.3 The protocol

For a system X and observable O. Every step is executable from outside, in the same model-relative stance as Part Three.

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-V or free-V—and, for free-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 (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 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.

Chapter 27

Intelligence Is Not Consciousness

27.1 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.

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.

27.2 The dissociations, both directions

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: the dreamer. 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.

27.3 The structural mirror

The deepest relation between the theories is not overlap but symmetry.

Both theories use the same spectral mathematics—components, weights, phases—applied to different objects. Part Three's κ is a phase-alignment order parameter over the *cloud's own* components; this part'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.

27.4 Why intelligent machines invite the consciousness inference

The framework explains the inference without validating it, in three steps.

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.

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.

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

Note carefully 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.

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.

27.5 A catalogue of decomposers

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 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.

	Consciousness	**Intelligence**
Object measured	pre-resolution cloud $p(x)$	decomposition of arrival y
Temporal locus	before the guess resolves	upon and after encounter
Core quantity	$\sigma \times \kappa$	$\sum_j \Delta\ell_j$
Role of coherence	organizes spread: phase alignment across the cloud	qualifies components: self-prediction within each
Failure I	collapse: $\sigma \to 0$	pass-through: $N_\varepsilon = 0$
Failure II	incoherence: spread without organization	confabulation: claims without gain
Healthy regime	coherent spread (music, held)	genuine decomposition (music, resolved)
Capacity constant	σ_c, organizing	N_c, validation (conjectured)
Relation to surprise	shape of the counterfactual landscape	size and structure of the residual
Link to felt experience	correspondence postulate, undischarged	none postulated
Canonical dissociation	dreamer: high C, low I	Fourier chip: high I, near-zero C

Table 27.1: 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.

System (observable)	N_ε	Q_{abs}	C_{ext}
Rock (incident sound)	0	~ 0	~ 0
Wire, lossless channel (audio)	0	0	~ 0
Thermostat (room temperature)	1	high	low
Twelve-band analyzer (audio)	12	mod.	~ 0
Fourier chip (spectral content)	$\sim$ all	~ 1	~ 0
Novice reader (legal case)	~ 2	low	high
Trial lawyer (legal case)	~ 30	high	high
Layperson (patient presentation)	~ 3	low	high
Physician (patient presentation)	~ 25	high	high
Expert musician (a song)	20+	high	high
Human subconscious (bodily regulation)	very large	~ 1	~ 0
Language model (symbolic signal)	very large	high	open
Bureaucracy (civic life stream)	large	falling	~ 0
Corporation (its market)	large	mod.	varies
Nation-state (its society)	very large	low–mod.	—

Table 27.2: Illustrative placements, not measurements. The two low-I regimes—pass-through and confabulation—appear alongside the dissociations from consciousness.

PART V

The World

CHAPTER 28

Reading Reality in Human Life

The apparatus is built. Turn it on a life.

28.1 The conversation that went wrong

Two people leave the same conversation. One is fine. One is not. Neither is lying about their experience.

Write the equations. Same world-state, but two projections: what he was tracking was whether the plan got settled; what she was tracking was whether the plan got settled *in a way that showed he had thought about her week*. Two observables, therefore two Actuals, therefore two numerators. Already the "same conversation" has stopped being the same event.

Now the denominators. His prediction machine, trained on a decade of these conversations, guessed accurately: this is roughly how it goes. Low prediction error, small $|\ln R|$. Her machine guessed accurately too. But her idea-vector had real magnitude pointed at a particular angle—an ideal of being attended to—and the Actual arrived on the real axis. Large B, and therefore large φ: high angular surprise with modest magnitude-surprise.

The result is exactly the argument they are about to have. He will say nothing happened, and he will be reporting his magnitude honestly. She will say something happened, and she will be reporting her angle honestly. Neither of them has a word for the difference, so they will fight about the facts, which were never in dispute.

That is the framework's first practical use: it gives the difference a name and a coordinate.

28.2 Disappointment, diagnosed

A person is repeatedly disappointed by other people. Two hypotheses, and they call for opposite interventions.

If the source is *prediction error*, the machine is mis-calibrated: it expects more reliability than the population supplies. Exposure fixes this. Enough Actuals, repeated, and the machine reshapes downward, and the disappointment declines without any change of values.

If the source is *ideational bias*, the machine is fine—the person is predicting others accurately—and there is a leaned weight on a pole, a held ideal of how people ought to be, contributing magnitude to B at an angle the world does not warrant. Exposure does *nothing* here. More accurate predictions leave the disappointment exactly where it was, which is why this kind of person often becomes cynical and no happier: they revised $\tilde{P}$ and never touched B.

The diagnostic from Chapter 12 applies. Does the surprise recur with the same shape across different people and different situations? Prediction errors are distributed. A repeating shape is a fact about the denominator.

28.3 The mechanics of gratitude and entitlement

Both are denominator operations, and the equation is unsentimental about them.

Gratitude lowers $|E|$, so the same Actual yields a larger quotient. This is not a trick and it is not self-deception, because nothing about the numerator has been touched. It is a real change to the term that was always yours.

Entitlement raises $|E|$, so the same Actual yields a smaller quotient. An entitled person is not being ungrateful about the world; they are dividing by a larger number, and getting the correct answer for the denominator they hold.

The framework declines to moralize about either, and instead points out the structural fact both traditions have always reported: *the denominator is the only term under your control, and it controls the whole quotient.*

28.4 Why the good news wears off

You get the thing. Reality spikes. Weeks later it is gone, and nothing about the thing has changed.

Chapter 5's conservative machine explains it exactly. The first arrival is a surprise: large $|\ln R|$. The machine, on repeated confirmation, reshapes—the new level becomes the expected level—and $\tilde{P}$ rises to meet $\tilde{A}$. $R \to 1$. $S \to 0$. The thing is still there. The quotient is not.

This is hedonic adaptation, and it is not a defect. It is the machine doing the only thing it is for: absorbing structure so that attention can go elsewhere. Part Four gives the same event its intelligence-side reading—the component has been absorbed, and absorbed components go quiet.

The practical consequence follows without any moralizing. If Reality is a quotient and denominators adapt, then durable Reality above 1 cannot come from a fixed Actual. It comes either from a rising numerator—which is a treadmill—or from work on the denominator, which is the only intervention that does not require the world's cooperation.

28.5 Attention, and what a life is made of

Chapter 11 noted that successful prediction is invisible. Take the consequence seriously and something uncomfortable follows.

The overwhelming majority of what happens to you produces $R = 1$ and $S = 0$ and is never registered at all. What reaches you is the residue: what your machine did not absorb. So the felt content of a life is not a sample of that life. It is systematically biased toward the unabsorbed.

Which means two people with identical lives and different absorption profiles have different lives, in the only sense that is available from the inside. And it means that becoming expert at something does not make it boring—it *refines what about it is interesting*, which is the difference between a novice who finds a piece of music pleasant and a musician who finds one bar of it astonishing.

28.6 The three operations

Everything the framework licenses in a human life reduces to three moves, and only three.

Change the numerator. Act. Actualize. Put something new in the Immutable Past. This is real, it is slow, and it is what most of life is.

Reshape the machine. Exposure, repetition, consequence. Not resolve, not assertion, not insight. The machine does not read; it counts.

Re-weight a pole. Examine a held asymmetry that the structure does not warrant, and hold it differently. This is the only one available to direct examination, and it is where the leverage is disproportionate to the effort.

Everything else on offer—most of what is sold as transformation—is one of these three, or it is nothing.

CHAPTER 29

Attention as the Ledger of Surprise

The most valuable resource on the planet is not oil, land, compute, or capital. It is human attention, and this framework says precisely what attention *is*, which is why it can say why the resource behaves the way it does.

29.1 The definition

Attention is normalized accumulated surprise.

Not a spotlight, not a budget, not a faculty. A ledger. Prediction runs continuously; most of it succeeds and produces nothing; where it fails, surprise accumulates; and attention is the conscious record of what Reality could not make disappear.

Three consequences follow immediately, and each explains something the other accounts of attention have to bolt on.

Attention is a remainder, not a resource. It is what is left when prediction has not finished its work. That is why it cannot be manufactured, only freed—and why every attempt to "capture" more of it is a fight over a by-product rather than a supply.

Attention is the settlement layer. Surprise accumulates until something resolves it. Attention is where accumulated surprise becomes judgment, care, and action—where a discrepancy stops being a signal and becomes a decision. Nothing else in the system does this.

Attention is the decision surface of Reality. Money can be borrowed, energy generated, computation scaled. Attention cannot be borrowed. Every human future has to pass through it first, which is what makes it the resource beneath every other resource.

29.2 Absorption, and the release of attention

Part Four supplies the supply side. When a component is absorbed, its recurrence stops generating surprise, and the attention it used to demand is released.

This is the whole economics of expertise, of habit, of infrastructure, of institutions, and now of machines. Every one of them is a mechanism for absorbing structure so that attention no longer has to be spent on it. You do not attend to the stairs. You did once.

And it explains the characteristic emotional signature of absorption when you were present for the transfer. Nobody resents their own subconscious, because it preceded their agency—no one experienced the handover. Bureaucracies, tools, and algorithms arrive *after* conscious discretion, so the transfer of attention and judgment into a procedure is experienced as a loss even when the gain is real.

> *The love–hate relationship with bureaucracy is what absorption feels like when you were present for it.*

29.3 Trust is amortized attention

A system earns the right to be ignored the way a loan is retired: one uneventful payment at a time, on a schedule no demonstration can skip.

In this framework that is not an analogy. Trust is the state in which a component has been absorbed—repeated Actuals have confirmed the pattern, the machine has reshaped, and residual surprise on that component has gone to near zero. The attention formerly spent monitoring is now free.

Which is why trust cannot be asserted, demonstrated in a single instance, or purchased. Absorption is a function of repetition and consequence, and no amount of either can be front-loaded. A perfect demo moves $\tilde{P}$ not at all, because the machine is conservative by construction: it changes only when repeated Actuals confirm that a new pattern is real.

29.4 The attention dividend

When absorption succeeds, hours come back. What happens to them is not automatic.

By default, released attention is recolonized—by new noise, new monitoring, new surfaces built precisely to capture a by-product that has just become available. The *attention dividend* is the yield of absorption, and like any dividend it compounds only if it is reinvested, and it is reinvested only if it is deliberately pointed at the remainder: at the part of the stream nothing has absorbed yet, where the next structure is.

Absorption that frees an hour and hands it to a feed has produced no dividend. It has moved the attention from one remainder to a manufactured one.

29.5 The surprise reserve

The hardest claim in this chapter, and the one with the most practical bite.

Absorption does not destroy surprise. It defers it. A system that has been ignored successfully for years has not eliminated the possibility of an anomaly outside its decomposition; it has accumulated the whole of that possibility into a single, unscheduled future event. The remainder does not disappear because nobody is looking at it. It concentrates.

So a system whose components have been absorbed stays solvent only if someone holds a **surprise reserve**: rehearsed attention, kept ready for the day the ignored system hands the problem back. Pilots who have not flown manually in years, operators who have never seen the failure mode, institutions with no living memory of the last crisis—each is a case of absorption without a reserve.

The formal version, from Chapter 24: the remainder is never certified structureless, and $H(\mu)$—the world's irreducible entropy on the observable—survives any amount of intelligence. Absorption reduces the reducible part. It cannot touch the rest, and the rest arrives on its own schedule.

Chapter 30

Bureaucracy: Civilization's Synthetic Subconscious

Bureaucracy is the theory's civilization-scale instance, and treating it that way turns a set of complaints into a set of measurements.

30.1 The formal reading

Its vocabulary V_{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 rather than decorative.

Bureaucracy is the pretrained layer of civilization. V_{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. 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.

30.2 The four pathologies

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. V_{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, and it is not about 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.

> *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.

30.3 Appeal paths

The design lesson follows directly, and it is the most actionable result in this chapter.

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, a review board.

So the quality of a bureaucracy is not only its decomposition score. It is whether its architecture *knows its own remainder exists*. An institution with a high component-gain and no appeal path is a machine that has mistaken its resolution limit for the boundary of the world, and it will fail exactly where its categories run out—which is to say, on the cases that matter most.

30.4 A normative firewall

Stated once, plainly: 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, fragmentation, residual criticality. It hands wholesale ammunition to no political camp, and any reader who finds it confirming everything they already believed should check whether they have read the half that cuts the other way.

CHAPTER 31

AI Agents and the Numerator

The claim about current artificial intelligence is a factorization, not a rating.

31.1 The factorization

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

The first factor is strong. The language model is a prediction machine—that is its essence; summarizing, drafting, planning, and 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, and 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, 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 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.

31.2 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, and unconscious correction—without effort, and their decomposition vocabulary is revised in the same stream it is scored on.

Current systems have 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.

31.3 The bureaucratic habitat

Here the two chapters meet, and the result is the sharpest prediction in this part of the book.

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.

> *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 already been pre-decomposed into the kind of world a checkpoint can faithfully receive.

31.4 What realer agents require

The path to higher I_{eff} runs through the numerator, and every 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 the previous chapter: an agent, like an institution, is safe in proportion to its architecture's acknowledgment that its $\rho \neq 0$.
- **Correction loops—**live revision of 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.

Chapter 32

The Zero-Attention Economy

One economic thesis, derived rather than asserted, because every term in it has now been defined.

32.1 The thesis

Artificial intelligence becomes truly valuable at the moment it stops asking to be noticed. The next great productivity metric is not how much a machine can do; it is how little attention it requires while doing it.

In this framework that is a statement about absorption. A system whose components have been absorbed generates near-zero marginal surprise and demands near-zero attention—it has earned the right to be ignored. Every mature technology has made that journey: electricity, plumbing, refrigeration, the payments network. The software industry has been an exception mainly because it kept selling surfaces that demanded to be looked at.

> *Follow the human attention. Where it is being spent is where the decomposition has not finished; where it has been released is where the value has already settled.*

That is why the pressure on conventional software is not primarily a pressure on features. Seats, dashboards, and interfaces are all instruments for capturing attention. A system that absorbs the work and returns the attention has, from the buyer's side, dominated them—and it has done so on the one axis the incumbent cannot easily concede, since its business model is denominated in exactly the resource the buyer wants back.

32.2 Why GPT never gets surprised

A clarifying case, because it puts three of the book's quantities in one sentence.

You can give a language model an unexpected fact. It can compute prediction error. It can even say "that surprised me." None of those things mean it experienced surprise.

Surprise, in this framework, is $\ln R$: the complex correction of an entity's Expectation by an Actual that arrived. The model has a pre-resolution distribution—so it has a cloud, and its C_{ext} on that observable is an empirical question, not a rhetorical one. What it does not have, in the ordinary deployment, is a standing denominator that persists across the encounter, an accumulation of unresolved surprise, or a machine that reshapes because the correction cost it something. The correction is computed and discarded.

So the right thing to say is neither "it feels" nor "it is a mere machine." It is that intelligence and consciousness are separate measurements (Chapter 27), that one of them is high and the other is open, and that the cues we use to read the second were trained on a world where the two never came apart.

32.3 The four costs of the transition

The framework predicts where the transition hurts, and none of the four are mysterious once the terms are defined.

The apprenticeship gap. Absorption eats routine work—and routine work is what quietly trained every expert. Every component a novice absorbs was absorbed by doing the boring version of the task many times. Remove the boring version and V never gets built, so I_{pot} never rises, so the capacity bound of Chapter 23 binds a generation lower than the last. Closing the gap means deliberately designing surprise back into the curriculum: re-injecting absorbed components into the attended stream, which is what deliberate practice has always been.

The recolonization of the dividend. Released hours default to new noise unless they are pointed at the remainder. This is Chapter 29's dividend, unpaid.

The missing reserve. Systems ignored successfully accumulate deferred surprise. Without rehearsed attention held in reserve, the day the ignored system hands the problem back is a day nobody has practiced for.

Numerator drift. $I_{\text{eff}} = I_{\text{act}} \times \Phi_{\text{num}}$. As more of the arrival stream becomes traces of traces, the second factor silently falls while the first keeps rising, and the aggregate metric everyone watches is the first one.

32.4 What the thesis does not say

It does not say that attention is worthless once released. The opposite: attention is the settlement layer, and the entire value of absorption is that it moves attention to a finer remainder.

It does not say that ignorable systems are safe. It says they are *trusted*, which is a statement about absorbed prediction and not about the tail.

And it does not say human work disappears. It says the location of human work moves to the remainder—which has always been where it was, and which has never before moved this fast.

CHAPTER 33

Worked Problems

Fluency, not admiration. Each problem states a case, works it in the framework's notation, and—this is the part that matters—names the distinction that would be lost by a sloppier reading.

Problem 1: the raise

Case. Two colleagues receive identical four-percent raises. One is delighted, one is insulted.

Work. $\tilde{A}$ identical by construction. Colleague A's machine, trained on a flat market, guessed two percent: $R = 2$, $\ln|R| = +0.69$. Colleague B's machine, trained on his own last three years, guessed six: $R = 0.67$, $\ln|R| = -0.41$. Both have small B, so $\varphi \approx 0$ for both.

Distinction preserved. This is pure prediction error, not bias. The intervention that works is exposure to the actual distribution of raises, and the intervention that fails is telling either of them how to feel.

Problem 2: the promotion that went to someone else

Case. A candidate is passed over. She reports not merely disappointment but disorientation—"I don't recognize this place."

Work. Magnitude: her machine put the odds near even, so $\ln|R|$ is moderate. Angle: she held a strong idea-vector—an ideal about how merit is supposed to relate to advancement—with large B at a definite θ. The Actual arrived on the real axis. φ is large.

Distinction preserved. Disorientation is angular surprise, and it is not proportional to magnitude-surprise. Anyone treating this as "she overestimated her chances" has scalarized the quotient and will give useless advice.

Problem 3: the veteran nurse

Case. A nurse walks past a patient who is stable on every monitored value and says "something is wrong." She is right.

Work. Her V on this observable is large; N_ε is high; the routine components—colour, posture, respiratory pattern, affect, the way the family is standing—are absorbed and cost nothing. Her residual load on routine structure is near zero, so attention is available and lands on the refined remainder. The monitors, by contrast, have a small fixed V: high Φ_{num} on what they measure, no components at all outside it.

Distinction preserved. This is not intuition in the mystical sense and not consciousness. It is I_{dec} plus Postulate 3. The monitor's failure is not inaccuracy—it is a vocabulary whose remainder contains the answer.

Problem 4: the confident agent

Case. An AI agent reconciles a quarter's records flawlessly. Every internal check passes. The reconciliation is wrong, because two upstream systems disagreed and the trace it read carried only one.

Work. I_{act} on the symbolic observable: high, and correctly so. Φ_{num} against the observable that matters: low. $I_{\text{eff}} = I_{\text{act}} \times \Phi_{\text{num}}$ is what the firm experiences.

Distinction preserved. No amount of improvement to the decomposer fixes this, and every internal check lives on the same corrupted trace. The Φ_{num}-raising mechanisms of Chapter 31 are the only lever.

Problem 5: the form that protects against nothing

Case. An agency requires a disclosure that no one reads and no decision has ever turned on.

Work. Run T4 on the current class: $\Delta\ell \leq 0$. The component once paid—it was fit on μ_{then}—and T5 has silently failed as the class drifted. It is now ritual, and it enters the ritual coefficient.

Distinction preserved. The complaint "this is pointless bureaucracy" and the measurement "this component has negative current gain" are the same claim, except that the second one can be audited and the first one cannot.

Problem 6: the seizure

Case. Massive, highly synchronized cortical activity, and the person is not conscious.

Work. κ is very high—phase-synchrony is nearly total. σ is collapsed: differentiation is gone. $C_{\text{ext}} = \sigma\kappa \approx 0.13$, and the failure is by *collapse*, on the left edge of the landscape.

Distinction preserved. A single-axis theory keyed to synchrony or to activity level predicts heightened consciousness here and is wrong. The product form gets it right, and predicts the opposite signature for delirium—high σ, low κ—which is the other way off the peak.

Problem 7: the firm that stopped guessing

Case. A software company shifts from selling a product to licensing an enforceable position. Revenue is steadier. Headcount and activity are unchanged.

Work. $C_{\text{ext}}(\text{firm} \mid \text{revenue})$ falls: the revenue-facing cloud has been narrowed by design, and sensitivity to daily Actual has been replaced by claim. Note what has *not* changed: C_{ext} with respect to litigation, rates, or political risk may be unchanged or higher.

Distinction preserved. Observable-relativity is doing the work. "The firm became less conscious" is meaningless; "the firm collapsed its revenue-facing cloud" is a measurement, and it predicts exactly which surprises it will now be unable to feel.

Problem 8: the expert who is never surprised

Case. A master craftsman reports low surprise in her domain. Is she less conscious of it than a novice?

Work. Realized surprise: low, because $\tilde{P}$ tracks $\tilde{A}$ closely. Counterfactual surprise landscape: highly structured, with distinct basins for each hypothesis her machine takes seriously. σ and κ both high. C_{ext} high.

Distinction preserved. Consciousness is read from the machine, not from the correction. This is Prediction 5 in Chapter 36, and it is the cleanest case where the framework's answer differs from the intuitive one.

Part VI

The Limits

CHAPTER 34

Boundary Conditions

A theory that does not state where it stops is not making claims; it is making noises. Here is where this one stops.

34.1 The domain of ordinary validity

Four conditions define the region in which the Reality Equation behaves as described.

Positive scalar Actual. $\tilde{A} > 0$. Actual is a magnitude on a declared observable. Negative or complex numerators are outside the theory, and attempts to put badness in the numerator produce nonsense rather than insight.

Positive predictive structure. $\tilde{P} > 0$. The prediction machine is producing a guess. An entity with no predictive structure at all on the observable has no real component, and its equation is degenerate rather than interesting.

Nonzero denominator. $E \neq 0$. This is the sharp one. Since $E = \tilde{P} + iB$, the denominator vanishes only if prediction and orientation both vanish—no guess and no direction. Reality is then undefined, which is the correct output. Something with no denominator is not having an experience; it is being passed through.

A declared observable. Every quantity in this book carries an index. Unqualified questions have no answers, and the theory returns that verdict rather than a number.

34.2 The two limits

Two conditions sit at the edges of the domain rather than inside it, and both are worth naming because both are familiar.

The limit $E \to 0$: bliss. As the denominator shrinks toward zero with $\tilde{A}$ fixed, $|R| \to \infty$ and $\ln |R| \to \infty$. Unbounded Reality, unbounded magnitude-surprise. This is the mathematical signature of the state the contemplative traditions describe: an entity that has released both prediction and orientation, for which any Actual whatsoever is infinitely more than was expected.

The theory does not endorse or explain that state. It notes that its coordinates are a limit and not a point, that the equation is undefined exactly at the limit, and that the traditions themselves describe it as unsustainable and non-propositional—which is what a boundary condition feels like from the inside.

The limit $E \to \infty$: numbness. Enormous denominator, whether from inflated prediction or from vast orientation. $|R| \to 0$, $\ln |R| \to -\infty$. Nothing that arrives can be adequate. This limit is more commonly reached in practice than the first, and it is the structure of despair.

34.3 What the theory does not claim

Stated flatly, so that nobody has to infer it.

It does not solve the hard problem of consciousness. Chapter 20 states the bridge as a postulate and marks it undischarged.

It does not claim that any particular entity has felt experience. Every consciousness quantity in this book is C_{ext}: an outside estimate of a property of a model.

It does not claim that intelligence entails experience, or that experience entails intelligence. Chapter 27 makes the independence formal in both directions.

It does not derive physics. The one numerological flourish in the program is fenced in Chapter 23 and disavowed there.

And it does not claim that the measures proposed are unique. Whether $\sigma \times \kappa$ is the right combiner, and whether $\sum_j \Delta\ell_j$ is the right score, are open problems in Chapter 37, and alternative combiners making different predictions would be welcome.

Chapter 35

Relation to Existing Frameworks

The framework is not free-floating, and its quantities have recognizable relatives. That is some evidence it carves at real joints. What follows are positioning notes, not endorsements and not derivations.

35.1 On the consciousness side

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 Chapter 16 is a free-energy-like quantity. The distinctive move here is a refusal: consciousness is *not* equated with prediction error, and is located in the shape of the cloud instead. A system can minimize surprise beautifully and be read as barely conscious; a system can be highly conscious and rarely surprised.

Integrated Information Theory, in the tradition of 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, and says so.

Criticality and the edge of chaos, in the tradition of self-organized criticality. The result $\sigma^\star = \sigma_c$—the $m = 2$ case of $\sigma^\star = \sigma_c(m-1)^{-1/m}$—says consciousness peaks at the boundary where coherence begins to fail—a critical point—which connects the inverted-U of Chapter 15 to a body of work on neural systems operating near criticality.

35.2 On the intelligence side

Minimum description length, in the tradition of Rissanen, and behind it Kolmogorov complexity and the algorithmic sufficient statistic. T4 is an MDL criterion. The contribution here is to *individuate* the code into components subject to stability and coherence tests, so that gain is attributable part by part rather than monolithically.

Sparse coding and dictionary learning, in the spirit of Olshausen and Field. Free-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 and Legg and Hutter's universal intelligence. Those score goal-achievement across environments, or raw compression. This one 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, in the spirit of Schmidhuber. Closely adjacent to Postulate 3: where that account says interestingness is the first derivative of compression, this one says attention's default flows to the remainder. The two meet at the remainder's most absorbable frontier, which is Open Problem 6.

Chunking and expertise, from Miller through Chase and Simon and the deliberate-practice literature. The expert/novice predictions of Chapter 36 are this theory's re-derivation of chunking, with a gain-scored, confabulation-controlled measurement attached.

35.3 Objections, with their answers

Stated before a reviewer states them. Each gets an answer or a 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, since 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 it is named as an extension in Open Problem 7 rather than 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 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 an observer-chosen basis." The two-regime answer of Chapter 22. Conventional-V scoring is explicitly convention-relative and fair within a declared vocabulary; free-V scoring is basis-fair by construction, charging every dictionary its own description cost. What must be conceded is that 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, and 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 2 and 3 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, so 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, so 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 throughout: these are properties of an observer's best model of a system relative to an observable, not metaphysical inherence claims. 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 a thermostat and a 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.

Chapter 36

Testable Predictions

The formalization yields commitments that could turn out false, which is what makes this a theory rather than a picture. Thirteen are stated here. Any one of them failing costs the framework something specific, and the text says what.

36.1 From the theory of consciousness

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.

36.2 From the theory of intelligence

Prediction 6 (Expert component-count)**.** Experts extract more admissible components than novices from identical signals in their domain, under conventional-V scoring with blinded raters and null-signal controls. If experts merely relabel rather than out-decompose novices—same I_{dec}, fancier vocabulary—the absorption account is in trouble.

Prediction 7 (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 is the joint signature of Postulates 2 and 3.

Prediction 8 (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.

Prediction 9 (Decomposition asymmetry in language models)**.** Language models score high I_{dec} under free-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 10 (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 11 (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.

Prediction 12 (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.

36.3 Across the two theories

Prediction 13 (Measured dissociation in one system)**.** A language model's C_{ext}—spread and mode-structure of its output distribution, per Chapter 17—and its I_{dec}—per Chapter 26—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, the Distinctness axiom weakens.

Chapter 37

Open Problems

Honest residue, stated so that others can attack it. A theory's open problems are part of its content, not an apology attached to it.

Open Problem 1 (The surprise bridge lemma)**.** State and prove the Tier-3 relation of Chapter 24: 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 the intelligence 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, since a jurisdictional fact changes what negligence means. Candidate formalisms: hierarchical dictionaries, compositional grammars, sheaf-like gluing of local components. The five tests are stated so as 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 the capacity bound of Chapter 23 is a theorem rather than a sketch.

Open Problem 4 (Vocabulary comparability)**.** Free-V scoring is code-family-relative and conventional-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 (the intelligence theory) and cloud-organizing capacity (the consciousness 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)**.** The intelligence 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 3.

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—which is why the traces exist in the first place. Proxy protocols—consequence-feedback consistency, cross-trace triangulation—need development.

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

Open Problem 10 (Choosing the observable)**.** Because both central quantities are observable-relative, a principled way to select or aggregate observables is needed before single-number comparisons across very different entities are fully meaningful.

Open Problem 11 (Estimating the cloud for the uninterrogable)**.** For a rock or the sun, estimating $p(x)$ means estimating a distribution we can only model, never sample. Those numbers are limited by model-confidence by construction, and the limit should be quantified rather than ignored.

Open Problem 12 (The resolving functional)**.** $G[p]$—mean, mode, or sample—is left open and may itself be entity-specific. Different resolving rules yield different denominators from the same cloud, and nothing in the theory currently determines the rule.

Open Problem 13 (Uniqueness of the measures)**.** Whether $\sigma \times \kappa$ is the right combiner, and whether $\sum_j \Delta\ell_j$ is the right aggregate, is open. Alternative combiners making different predictions would be welcome, and would be the fastest way to improve the framework.

Open Problem 14 (The correspondence postulate)**.** Undischarged, and probably not dischargeable from outside. The best available move is to make it precise and test its co-variation where we have inside access—which is to say, in ourselves, under conditions where (σ, κ) can be independently manipulated.

CHAPTER 38

Toward Ideation

The phase-coherence reading of κ leaves a gift for the companion theory this book does not contain.

38.1 The seam

If coherence is *phase* coherence, then the imaginary component of Expectation—idea-orientation $Be^{i\theta}$—is not a separate department of the framework. It is the same angular structure, seen at the level of commitments rather than components.

Consciousness and Ideation become two readings of one complex object:

Consciousness $\sim$ *the coherence of phase across the cloud.*

Ideation $\sim$ *the net phase* θ *the cloud commits to.*

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.

38.2 What the third theory would have to do

Three things, and they are stated here so that the gap is visible rather than papered over.

Make θ dynamical. Chapter 6 gave the weighted resultant its magnitude and angle and then set the angle aside. A theory of Ideation has to say how θ moves: how commitments form, how they are revised by Actuals arriving off-axis, and why some angles are stable attractors and others are not.

Connect θ to the remainder. Open Problem 1 poses the possibility that the *orientation* of what a system failed to decompose drives the angular part of Surprise. If that holds, angular surprise is not an independent phenomenon but the intelligence theory's residue read on the imaginary axis—which would join all three theories at a single point.

Explain selection. Chapter 35 conceded that the selection problem in recomposition is unmeasured. In the language of this chapter, selection is exactly the commitment of a coherent cloud to one angle. A theory of Ideation that did not answer it would not be worth having.

38.3 Why the program is one program

Read from the top, the whole architecture is a single object examined in four places.

Reality is the quotient. Surprise is its logarithm. Consciousness is the coherence of what the machine held before it resolved. Intelligence is the fineness with which what arrived was resolved. Ideation is the angle the resolution commits to.

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

That sentence is the book in one line, and every clause in it now has a definition, a measure, and a way to be wrong.

CHAPTER 39

Why Reality Persists

One question remains, and it is the one the first edition ended on. Why does any of this continue?

39.1 The question

Consider what the framework predicts about a stable world. Prediction machines are conservative and they learn. Components get absorbed and go quiet. Surprise on absorbed structure falls toward zero. $\tilde{P}$ rises to meet $\tilde{A}$; $R \to 1$; $S \to 0$.

Run that forward far enough and Reality flattens. Everything is predicted, everything is absorbed, nothing is registered, and the quotient sits permanently at its fixed point. That is not death—the entity is still running—but it is something like the heat death of experience.

It does not happen. Not to individuals, not to institutions, not to species. Why?

39.2 Three answers, in ascending order

The world's own entropy. Chapter 24's Tier 1 identity leaves $H(\mu)$ standing for any intelligence whatsoever. Absorption reduces the reducible part of expected surprisal and cannot touch the rest. So there is always residue, and therefore always a remainder, and therefore always somewhere for attention to go.

That answer is true and insufficient. It explains why noise persists. It does not explain why the residue is *structured*, why it rewards attention, why resolving one layer of it reliably exposes another.

Coherent spread is the optimum. Chapter 15 showed that in a rich, nonstationary world, the expected-error-minimizing machine is not the narrowest one. A collapsed model is catastrophically wrong the moment the world moves. So machines that survive do not collapse their clouds;

they hold coherent spread matched to the structured complexity of their environment—which means they hold live possibility, permanently, as a condition of continuing to function.

That answer is better. It explains why entities do not settle. It still does not explain why the environment keeps having more structure to match.

The ideals have not finished. Here is the framework's own answer, and it requires taking Chapter 9 seriously.

Ideals are pre-existing endpoints seeking actualization. Actualization happens only in the Real, only through actualizers, and always at finite resolution—Chapter 23's capacity bound is not a temporary limitation but a permanent property of every finite system. So every actualization is partial. Every made thing is a projection of an ideal onto a real actualizer with a finite $\delta\theta$, and the difference between the ideal and the made thing is precisely a remainder.

That remainder is not noise. It is the structured gap between what was pointed at and what got made. It is what the next actualizer will resolve a little further, and what the one after that will refine again, each at their own finite resolution, each leaving a finer remainder behind.

> *Reality persists because ideals have not yet achieved perfect actualization through real actualizers—and cannot, because actualizers are finite.*

39.3 What that means for the quotient

The denominator has an imaginary component precisely because entities are oriented toward endpoints that have not arrived. Remove the orientation and $B \to 0$; Expectation becomes purely predictive; the machine converges; the quotient settles at 1.

What keeps B nonzero across a life, an institution, and a civilization is that ideals keep selecting hosts, hosts keep actualizing partially, and partial actualization keeps exposing the next thing to point at. The gap is the engine.

So the quotiental drama does not continue because the world is chaotic or because people are foolish. It continues because finite actualizers keep pointing at infinite endpoints and keep almost getting there. Every Actual that lands is a partial answer that opens a finer question. Every absorbed component frees the attention that finds the next remainder.

39.4 The last thing

Do not ask only what happened. Ask what it was divided by.

Do not ask only what an entity feels. Ask what the shape of the machine is that generates its guess.

Do not ask only how smart something is. Ask what it resolved, at what gain, against what trace, and what it left in the remainder.

And when the quotient is not 1—which is always—do not treat that as a failure of the world to cooperate. It is the signature of an unfinished actualization, which is the only condition under which anything is still being made.

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

Appendix A

Notation

Symbol	Meaning
O	the observable; every quantity is indexed to one
A, $\tilde{A}$	Actual; normalized Actual, $\tilde{A} = A/x_0$
x_0	reference scale on the observable
Π_O	projection of the world-state onto the observable
E	Expectation, $E = \tilde{P} + iB$, with $E \neq 0$
P, $\tilde{P}$	the prediction machine's resolved guess; normalized
B, θ	magnitude and argument of the weighted resultant idea-vector
V_I	the idea-vector, $V_I = Be^{i\theta}$
R	Reality, the complex quotient $\tilde{A}/E$
S	Surprise, $\ln R = \ln \lvert R \rvert + i\varphi$
φ	angular surprise, principal branch $(-\pi, \pi]$
$p(x)$	the probability cloud held before resolution
$G[p]$	the resolving functional; $\tilde{P} = G[p]$
σ	spread: normalized width of the cloud
κ	cloud-coherence: organization into distinct, stable modes
σ_c	organizing capacity; $\sigma^\star = \sigma_c(m-1)^{-1/m}$
C_{ext}	apparent consciousness, $\approx \sigma \times \kappa$

Symbol	Meaning
G, Λ, K	sensitivity, probability horizon, action-coupling
$S(a), \bar{S}$	counterfactual surprise landscape; expected surprise
$y, \mathcal{Y}_O, \mu$	arrival stream; signal class; its distribution
V	the decomposition vocabulary
c_j, ρ_X	components; the remainder
F	the composition operator
χ_j	internal coherence of component j
$\Delta\ell_j$	net compression gain of component j
N_ε	admissible component count (public score)
$I_{\text{dec}}, I_{\text{act}}, I_{\text{pot}}$	decomposition, actualized, and potential intelligence
$\delta\theta$	angular resolution over idea-space
Q_{abs}, L_ρ	absorbed fraction; residual load
$\Phi_{\text{num}}, I_{\text{eff}}$	numerator fidelity; effective intelligence $I_{\text{act}}\Phi_{\text{num}}$
N_c	validation capacity (conjectured)

Two notes on collisions. The resolving functional always appears with its bracketed argument, $G[p]$, and is not the sensitivity modulator G. The spread formula normalizes by x_{ref}, a dispersion scale chosen for the cloud, which need not equal the reference scale x_0 used to make Actual and the guess dimensionless.

Appendix B

Definitions

Entity That which a Reality Equation represents, relative to an observer's model. Not restricted to organisms.

Observable (O) The quantity, indexed to the entity, that Actual measures. Every claim in this book is of the form "with respect to O."

Actual (A) The exact, given projection of the Immutable Past onto the observable. Positive, scalar, declared, morally neutral, residue-free.

Ideal A pure form: an unconditioned endpoint on the unit circle of future possibility. Ideals do not occur; they orient.

Real The domain of what can actualize—the interior of the circle—and of actualizers, the entities that do the actualizing.

Reality (R) The complex quotient $\tilde{A}/E$. Not a domain; a relation. The felt subjective experience of a witness in the eternal now.

Expectation (E) The complex, dimensionless denominator, $E = \tilde{P} + iB$, entity-specific and nonzero.

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

Imaginary component ($Be^{i\theta}$) The weighted resultant idea-vector: magnitude B after cancellation, angle θ.

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

Prediction machine The structure in the model that generates $\tilde{P}$. Always on, conservative, trained by repetition and consequence rather than by assertion.

Probability cloud (p) The machine's distribution over outcomes before it resolves into $\tilde{P}$. Lives in the machine, not the denominator.

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 prediction machine through repeated interaction with Actual.

Arrival stream (y) The time-indexed sequence presented to a system over an encounter window. In mediated systems, a *trace* of the projection rather than the projection.

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

Vocabulary (V) The family of candidate component-types available for decomposition.

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

Coherent component (c_j) A candidate passing T1–T5: distinguishable, stable, internally coherent, 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.

Attention Normalized accumulated surprise; the settlement layer where unresolved surprise becomes judgment, care, and action.

APPENDIX C

Axioms and Postulates

On the equation

A1. Every entity we choose to model can be represented by a Reality Equation, relative to an observer and an observable.

A2. Actual is the exact, given projection of the Immutable Past onto the observable: common in source, specific in projection.

A3. Expectation is the complex, dimensionless denominator $E = \tilde{P} + iB$, with $E \neq 0$.

A4. Reality is the quotient $\tilde{A}/E$ and is preserved as complex; Surprise is its principal-branch logarithm, with magnitude and angular parts.

A5. Ideals are unconditioned endpoints; ideas are ideals as conditioned by a host; no system invents them.

A6. Polarity pairs are poles of one symmetrical whole; equal weighting cancels, so orientation requires asymmetry.

On consciousness

C1. The real component $\tilde{P}$ is a single number: the prediction machine's guess, a functional $G[p]$ of its cloud.

C2. The cloud belongs to the machine, not to the denominator; the denominator receives the resolved guess.

C3. Spread is the width of the cloud; coherence is its organization into distinct, stable modes—equivalently, its phase alignment.

C4. Apparent consciousness is coherent spread, $C_{\text{ext}} \approx \sigma\kappa$; it vanishes if either factor vanishes.

C5. Coherence degrades once spread outruns capacity σ_c, so C_{ext} traces an inverted-U, peaking at the edge of coherence.

C6. Collapsed certainty is not consciousness; incoherent uncertainty is not consciousness; coherent uncertainty is consciousness.

C7. Surprise is the complex correction of the guess by Actual; consciousness is read from the machine, not from the correction.

C8. Consciousness is relative to an observable.

C9. Angle matters for Ideation; shape matters for Consciousness.

On intelligence

I1. **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.

I2. **Operationalism.** Intelligence is measured by decomposition performance under the protocol, not by subjective impression, fluency, confidence, or resemblance to humans.

I3. **Demonstration by decomposition.** A system demonstrates intelligence relative to O exactly when it resolves arriving signals from $\mathcal{Y}_O$ into admissible coherent components.

I4. **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.

I5. **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.

I6. **Capacity bound.** No finite system resolves the continuum; every intelligence class has finite resolution, and $I_{\text{act}} \leq I_{\text{pot}} < \infty$.

I7. **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 postulates

Deliberately not axioms. These are empirical bets with named boundary conditions, held to the standard of telling a critic where to aim.

P1. **Correspondence.** Wherever an entity has felt experience at all, the richness of that experience co-varies with C_{ext}.

P2. **Absorption relieves attention.** Absorbed components generate near-zero marginal surprise on recurrence, and correspondingly reduced attention demand. *Boundary condition:* explicit-monitoring pressure can force re-attention.

P3. **Attention flows to the remainder.** The default allocation of attention tracks residual surprise density. *Boundary condition:* task demands and top-down goals move attention independently.

Sources

This edition consolidates and supersedes the following, all by the author.

- *The Reality Equation*, first edition, sixteen chapters. Published in serial form at johnrector.me, April 2026. Parts One and Two of the present volume rewrite that material; Chapters 1–12 of the present volume rework, reorder, and extend that material; the first edition's chapters 13–14 are absorbed into Part Five, and its chapters 15–16 into Part Six.

- *The Shape of the Prediction Machine: A Theory of Consciousness from the Reality Equation*, 1 July 2026. Part Three of the present volume is a rewriting of this paper for book form. The formal results—the two definitions of coherence, the derived inverted-U with $\sigma^\star = \sigma_c(m-1)^{-1/m}$, the surprise landscape, the measurement protocol, the correspondence postulate, and the worked scorings—are carried over intact.

- *Intelligence as Decomposition: A Theory of Intelligence from the Reality Equation*, 3 July 2026. Part Four is a rewriting of this paper. The five admissibility tests, the two scores, the capacity bound, the three-tier surprise bridge, the two attention postulates, the failure modes and confabulation curve, the structural mirror, and the factorization $I_{\text{eff}} = I_{\text{act}} \times \Phi_{\text{num}}$ are carried over intact.

- The attention and absorption essays of July 2026—*Attention Is the Ledger of Surprise*, *The Decision Surface of Reality*, *The Settlement Layer of Surprise*, *Trust Is Amortized Attention*, *The Attention Dividend*, *The Surprise Reserve*, *The Apprenticeship Gap*, *The Zero-*

Attention Economy, *Follow the Human Attention*, and *Why GPT Never Gets Surprised*—which supplied Chapters 29 and 32.

Positioning references named in Chapter 35—Clark, Hohwy, and Friston on predictive processing and free energy; Tononi on integrated information; the self-organized criticality literature; Rissanen on minimum description length and Kolmogorov behind it; Olshausen and Field on sparse coding; Legg and Hutter on universal intelligence; Schmidhuber on compression progress; Miller, Chase, and Simon on chunking—are cited as positioning, not as derivations, and no claim in this book should be read as following from any of them.