

The Compression Hypothesis

*Quantum Foundations, Interface Theory, and Dual-Aspect Monism
as Arguments Against the Primacy of Spacetime*

Aaron Long

Draft version 2.4 · July 2026

Position paper / synthesis. Not an experimental report. Not peer-reviewed.

*Figures, logic-loop diagrams, and mathematical scaffolding.
Schematics generated for this manuscript.*

Preface to Draft 2.4

I am not writing this from a university appointment or a laboratory. I am what people sometimes call an independent scientist: someone who reads the literature carefully, argues with the arguments, and tries to get the claims right without the usual institutional scaffolding. That has costs. It also has one freedom. I can spend pages saying what the popular versions of these ideas get wrong, without needing the conclusion to flatter a grant, a brand, or a spiritual market.

This manuscript was written and revised with AI assistance. The systems used for drafting, critique, expansion, and review were Claude Opus (including the Opus 4.5 / Opus 5 product line), Grok 4.5, and Fable 5. They helped with structure, literature leads, adversarial pressure on weak spots, and the long work of restoring full essay length after an earlier compressed draft. They did not replace judgment. Every claim that remains is one I am willing to stand behind. Every claim that was cut was cut because it did not survive scrutiny. If something here is wrong, the error is mine.

Epistemic status is unchanged. Part I reports standard physics as carefully as I can manage. Parts II and III engage contested but published academic positions. Part IIb bridges into mainstream cognitive science. The appendices are stress tests and cautionary notes, not evidence for the main thesis. The Compression Hypothesis itself is a conjecture with a proposed test structure — not a result.

Why draft 2 was shorter than draft 1, and what 2.1 does about it. Draft 2 correctly restructured the argument — exclusion zone instead of false consilience, predictive processing bridge, high-strangeness and psychedelics moved to appendices, recent FBT critiques cited — but it was rebuilt too tightly, more like research notes turned into prose than like the long-form essay of draft 1. That was a density mistake. A paper in this genre needs air: room to state a theorem carefully, state what it does not show, meet the objection, and only then move on. Draft 2.1 restores that air. It aims at draft 1's scale (or slightly beyond) so that new sections are not paid for by gutting old ones.

Draft 2.3 cleans the graphs, places each figure next to the argument it serves (no early figure dump), and fills the Hoffman gap with author schematics of ITP and deflated FBT — there is no free reusable Hoffman journal art. Logic-loop figures (relations, subject–object cut, icons vs constraints, interface vs simulation, time-as-format) sit with Part III–IV claims. Open-license Wikimedia diagrams where they help: JPEG/DCT, double-slit, Bell, AdS/CFT, Friston free-energy, dualism-vs-monism contrast, Jung unus mundus model, Pauli portrait. Each figure has a short caption: purpose, limits, source.

What 2.1 over-improved relative to earlier drafts remains. Convergence is rewritten as a shared exclusion zone plus an explicit matrix of what each strand still allows. Hoffman's Fitness-Beats-Truth program is kept only in its defensible form, with Bagwell (2023) and Mattes (2026) on the page, not in a footnote I forgot. "Compression" gets scientific backbone from predictive processing and Markov blankets, with a loud distinction between the weak claim (brains encode) and the strong claim (spacetime is format). Ontic structural realism is named as the cleanest ally for a relational ground. Intersubjective agreement and time get real subsections. Discrimination tests get methods sketches. Gibson's affordances appear where objecthood is explained as action-chunking. Dual-aspect monism remains historically serious and evidentially demoted. Vallée and chemically altered perception remain appendices. I still refuse consciousness-causes-collapse, intention physics, synchronicity-as-fact, and the idea that a drug or a UFO story hands you the ground of being.

Draft 2.4 stops adding allies and starts building the object. Four changes carry the version. A one-page box now states, in plain terms, what this paper may and may not be cited for — the difference between the claims it earns and the claims that get projected onto it. The abstract tags each strand by the kind of thing it is, so no reader mistakes a historical parallel for a theorem. The Compression Hypothesis is broken into a labelled stack — what is core to it, what is optional, what is frankly speculative — because a conjunction defended as one lump is a conjunction no one can attack cleanly, and I would rather be attacked cleanly. And a small worked model of the Fitness-Beats-Truth argument is now on the page, with a plot that can come out against me: it makes the evolutionary strand a thing you can turn a knob on, not a slogan you either swallow or spit out. A death table in Part IIb states, in one place, what would kill the strong claim while leaving mainstream cognitive science standing — the question a hostile referee is right to ask first. The prose has been tightened throughout. None of this makes the Compression Hypothesis true. It makes it easier to find out that it is false, which is the only kind of progress a paper like this can honestly offer.

Abstract

Four largely independent research programs put pressure on the idea that a world of localized objects distributed through spacetime is the fundamental furniture of reality. They are not the same kind of thing, and the paper keeps them ranked. The foundations of quantum mechanics supply theorems with experimental backing. The emergent-spacetime program in quantum gravity supplies a serious technical research program with exact results in model universes, not yet a fact about ours. The evolutionary analysis of perception supplies a contested but published argument. The dual-aspect tradition running from Pauli and Jung to Atmanspacher's reconstruction supplies a historical and philosophical parallel with a thin empirical edge. This paper argues that the pressure they share is real, that it is far stronger as a shared exclusion zone than as a forced identity of four theses, and that it is nonetheless systematically overclaimed by enthusiasts and systematically ignored by critics.

I defend a positive thesis, the Compression Hypothesis: that spacetime and its contents constitute a lossy, action-oriented encoding generated by bounded agents under selection and informational constraint, and that what is encoded is not a hidden world of smaller objects but a relational structure that does not have object-form. The working ontology for the ground is ontic structural realism — structure without fundamental things — with the interface supplying the format; the dual-aspect language is treated as historical rhyme, not as the name of the ground. I state the hypothesis as a labelled stack, separating what is core to it from what is optional and what is speculative, so that it can be attacked one conjunct at a time rather than defended as a lump. I distinguish it sharply from “consciousness causes collapse,” which the physics does not support. I give six desiderata any successful theory here must meet, a small worked model of the evolutionary argument whose result can cut against the thesis, a table of what would falsify the strong claim while leaving mainstream cognitive science intact, the main objections at full length, and two near-term discrimination sketches. Appendices treat high-strangeness reports and chemically altered perception as stress tests of the interface idea, not as evidence for it.

Epistemic status. This is a synthesis and position paper, not an experimental report. Part I is standard physics reported as accurately as I can manage. Parts II and III are contested but published academic positions. Part IIb connects to mainstream theories of brain function. The appendices are speculative sociology and plasticity notes on unresolved questions. The Compression Hypothesis is a conjecture with a proposed test structure, not a result.

In short

The world you act in — space, time, separate things with definite properties — behaves like a desktop: useful, lossy, built for action. It is not a transparent window onto whatever is more basic. The deeper structure, if these arguments are on the right track, is relational rather than a warehouse of objects. That does not mean “anything goes.” It does not mean intention rewrites physics. It does not mean a psychedelic state or an anomalous sighting hands you the ground of being. It means the format of ordinary experience is not the format of the world as such — and that several serious research lines, for different reasons and with different force, keep pressing against the naive object.

What this paper may and may not be cited for

Work in this territory is misread in a predictable direction: the careful negative result gets promoted into a triumphant positive one, and a conjecture gets quoted as a finding. This box exists so that no honest reader has to guess where the line is. Cite the left column. Do not cite the right one. The middle names what is genuinely open.

You may cite this paper for	You may NOT cite it for
That naive object-primacy — definite objects with definite properties in fundamental spacetime — is under real pressure from several independent directions.	That any experiment or theorem proves the Compression Hypothesis, or that physics has shown reality to be a “desktop.”
That those directions form a shared exclusion zone, not four proofs of one thesis; the divergence table (Part IV) is part of the claim.	That the four strands converge on one positive ontology, or that Bell / Kochen–Specker / AdS/CFT support conscious realism or idealism.
A deflated reading of Fitness-Beats-Truth: selection does not guarantee structure-preserving perception, which supports lossy encoding — not fiction.	That “evolution proves perception is an illusion,” or that FBT establishes its conclusion independent of its genericity assumption.
The Compression Hypothesis as a stated, labelled conjecture with a working ontology (structural realism + interface) and a proposed test structure.	That the hypothesis is established, that it solves the hard problem, or that consciousness collapses the wave function / intention moves matter.
The anti-overclaim hygiene itself: precise statements of what quantum foundations do and do not establish.	Synchronicity, high-strangeness, or psychedelic experience as evidence for the thesis — the paper explicitly refuses all three.

Open questions (cite as open, not settled): whether real fitness landscapes are monotone enough to blunt the evolutionary strand (Part II toy model; Part VI sketch A); whether emergent spacetime holds beyond anti-de Sitter model universes; whether any observer-involving framework can derive the Born rule non-circularly; whether the strong claim can be made to generate a prediction that fails.

Contents

Preface to Draft 2.4

Abstract / In short

What this paper may and may not be cited for (claims box)

1. Introduction (Figs. 1–2, 12–13)

Part I — Physics (Figs. 3–6)

Part I-M — Math: constraints, definitions, model theorems (what is / isn't proved)

Part II — Hoffman / FBT (incl. a worked toy model)

Part IIb — Predictive processing; weak vs. strong death table (Figs. 7–8, 14–15)

Part III — Dual-aspect / Jung (Figs. 17–18)

Part IV — Exclusion zone + CH, stated as a labelled stack (Figs. 9, 16)

Part V–VI — Objections and discrimination

Appendix A — Vallée · Appendix B — Plasticity (Fig. 10)

Figure credits · References

1. Introduction

1.1 The question and why it is usually mishandled

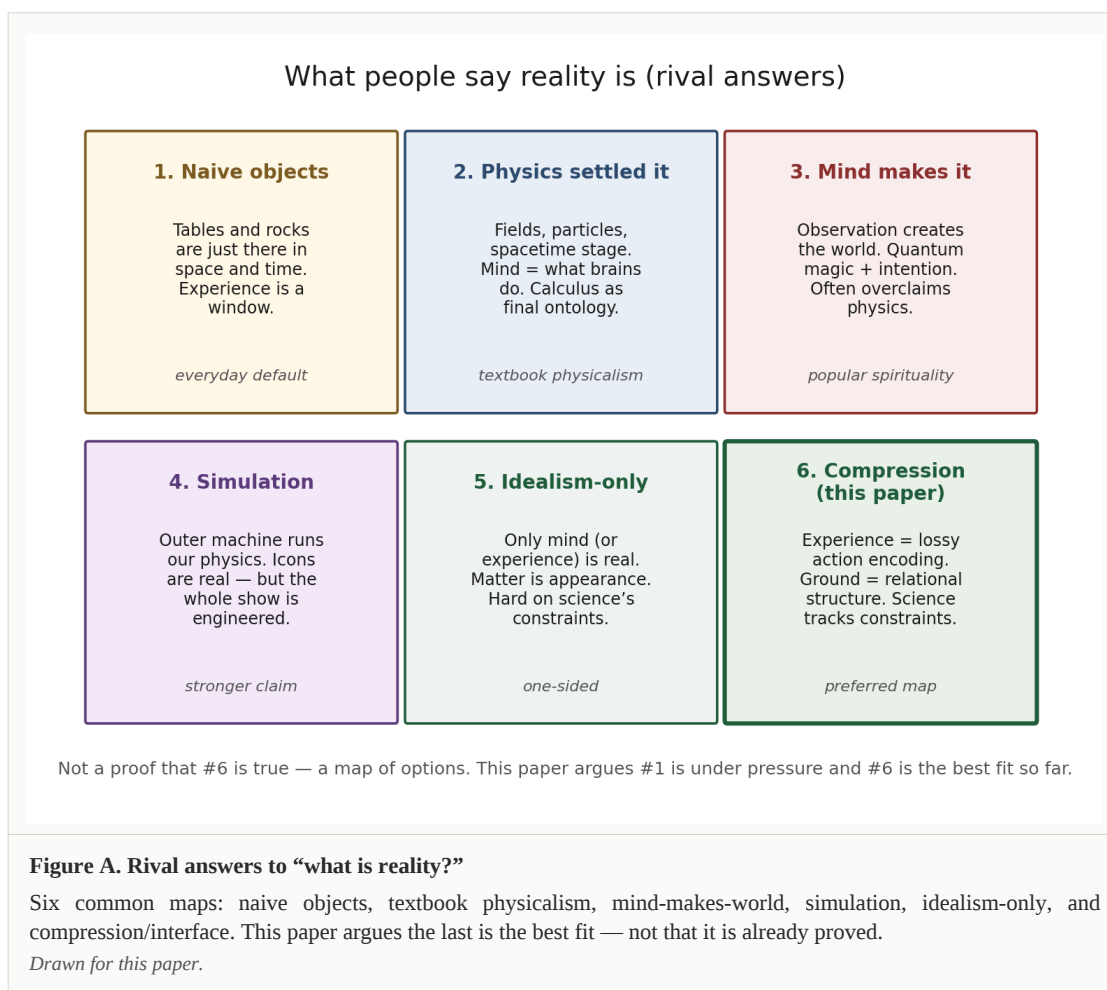
“What is reality?” is a question that gets answered badly in two opposite directions.

The first bad answer is that physics has already settled it. On this view, reality is a four-dimensional manifold populated by quantum fields; consciousness is what certain arrangements of those fields do when they get complicated enough; and everything else is either poetry or confusion. This answer has the enormous virtue of predictive success and the enormous defect that it does not actually say what a quantum field is, cannot say why measurement produces definite outcomes, and cannot say why there is experience at all. It mistakes an extraordinarily good calculus for an ontology. As Eddington observed nearly a century ago, physics gives us the structure of the world — the relations, the symmetries, the equations — and is constitutively silent about the intrinsic nature of whatever is doing the relating (Eddington 1928).

The second bad answer is the one that circulates in popular spirituality: that quantum mechanics proves consciousness creates reality, that observation collapses the wave function in the sense that a mind reaches out and makes something be, and that this vindicates whatever metaphysics the speaker already preferred. This answer is not merely unrigorous; on the specific claims it makes about physics, it is false, and I will say precisely where and why. A great deal of intellectual damage has been done by the fact that these two positions are usually treated as the only options, so that anyone who finds the first answer inadequate is presumed to be committed to the second.

The interesting territory lies between them, and it is more populated than either camp acknowledges. There exists a body of published, peer-reviewed, mathematically serious work — in quantum foundations, in quantum gravity, in evolutionary game theory, in philosophy of mind — that converges on a claim far stronger than the first answer permits and far weaker than the second asserts. The claim is this: the object-in-spacetime is not fundamental. Not “objects are made of smaller objects.” Not “objects are fuzzy at small scales.” Something more radical: that the form of a definite thing occupying a definite region is a feature of the representation, not of the represented.

Before the positive thesis, it helps to put the usual answers to “what is reality?” on one page — so the Compression Hypothesis is not floating free of the competition.



1.2 The structure of the argument

The argument of this paper is not that any single line of evidence establishes the thesis. None does. The argument is convergent in a limited and carefully stated sense. Four investigators, working with different instruments, on different questions, in different decades, each put pressure on naive object-primacy. If their methods are genuinely independent, that shared pressure is worth more than any one line alone. The analogy people reach for here — several independent measurements converging on the age of the Earth long before any one was decisive — is worth stating and then immediately qualifying, because it flatters the case. In the age-of-the-Earth example the independent lines converge on a single positive number. Here they do not: they converge on a negative — that the naive object is not free — and diverge on what replaces it. The consilience is real but it is consilience about an absence, and Part IV keeps that distinction in front of the reader rather than letting it blur.

But consilience requires sameness of proposition, not sameness of mood. Bell’s result is not Hoffman’s result is not Maldacena’s duality is not dual-aspect monism. Draft 1 sometimes let the rhetoric of convergence run ahead of that distinction. Draft 2.1 does not. What the strands share is better described as an exclusion zone: the picture of definite objects with definite properties laid out in fundamental spacetime is no longer free. What replaces that picture is not settled by the zone itself. Part IV makes this explicit with tables of what each strand excludes and what it still allows.

The four strands, ranked by how much weight they can honestly bear:

- **Quantum foundations.** A sequence of no-go theorems — Bell, Kochen–Specker, PBR, Local Friendliness and related Wigner’s-friend results — that progressively eliminate local pre-existing values, non-contextual value assignments, and (under further assumptions) absolute observed facts. This is the most rigorous strand. The results are not in serious dispute; what they mean is.

- **Quantum gravity / emergent spacetime.** The holographic principle, AdS/CFT, and entanglement–geometry correspondences, which together suggest that spacetime is not a stage on which physics happens but an emergent description of an underlying quantum information structure. Nima Arkani-Hamed’s formulation — that “spacetime is doomed” — is now close to a working assumption in parts of the field rather than a heresy. A serious technical program; not yet a fact about our de Sitter universe.

- **Evolutionary epistemology.** Donald Hoffman’s Interface Theory of Perception and the Fitness-Beats-Truth theorem, which argue that natural selection does not favour veridical perception in the strong sense, and that our perceptual categories — including space, time, and object — should be expected to be useful rather than true. Useful when read as an argument for lossy, action-tuned encoding; overstated when read as “perception is fiction.”

- **Dual-aspect monism.** Jung’s unus mundus, developed in correspondence with Wolfgang Pauli, and reconstructed rigorously by Harald Atmanspacher and others: the proposal that mind and matter are two epistemic aspects of a single underlying reality that itself is neither. Structural interest; thin empirical edge. Ancestor and parallel, not coequal with theorems.

These four are not equally strong. Strand 1 is a body of theorems; strand 4 is, at its empirical edge, close to unfalsifiable. I will not pretend otherwise. A substantial portion of this paper is devoted to saying exactly how much weight each can bear. The negative pressure they share is real. Their positive reconstructions do not overlap.


1.3 What I will argue

I will argue for a specific positive proposal, the Compression Hypothesis, which can be stated in one sentence:

The world as experienced — extended in space, persisting in time, composed of separable objects with definite properties — is a lossy encoding generated by a bounded agent for the purpose of action, and the structure it compresses is relational and non-spatial.

The word “compression” is doing real work. A compression is not an illusion. A JPEG of a face is not a hallucination of a face; it stands in a systematic, information-preserving-in-the-relevant-respects relation to whatever produced it. But the JPEG’s format — its blocks, its frequency coefficients — is a property of the encoding scheme, not of the face. The Compression Hypothesis says that space, time, and objecthood are format, not content.

Interface vs underlying process



Icons help action. They are not a transparent copy of the process underneath.

Useful and constrained — not free fiction.

Figure 1. Icons and what sits under them.

What you deal with day to day is more like a desktop than a window. Useful, constrained — not make-believe.

Drawn for this paper.

This framing does two things I want. It preserves the objectivity of science: compressions can be better or worse, more or less predictive, and the difference is not arbitrary. And it explains why the deep structure would be so hard to see: you cannot easily inspect the format from inside the format, any more than a program running inside a virtual machine can inspect the host by executing instructions.

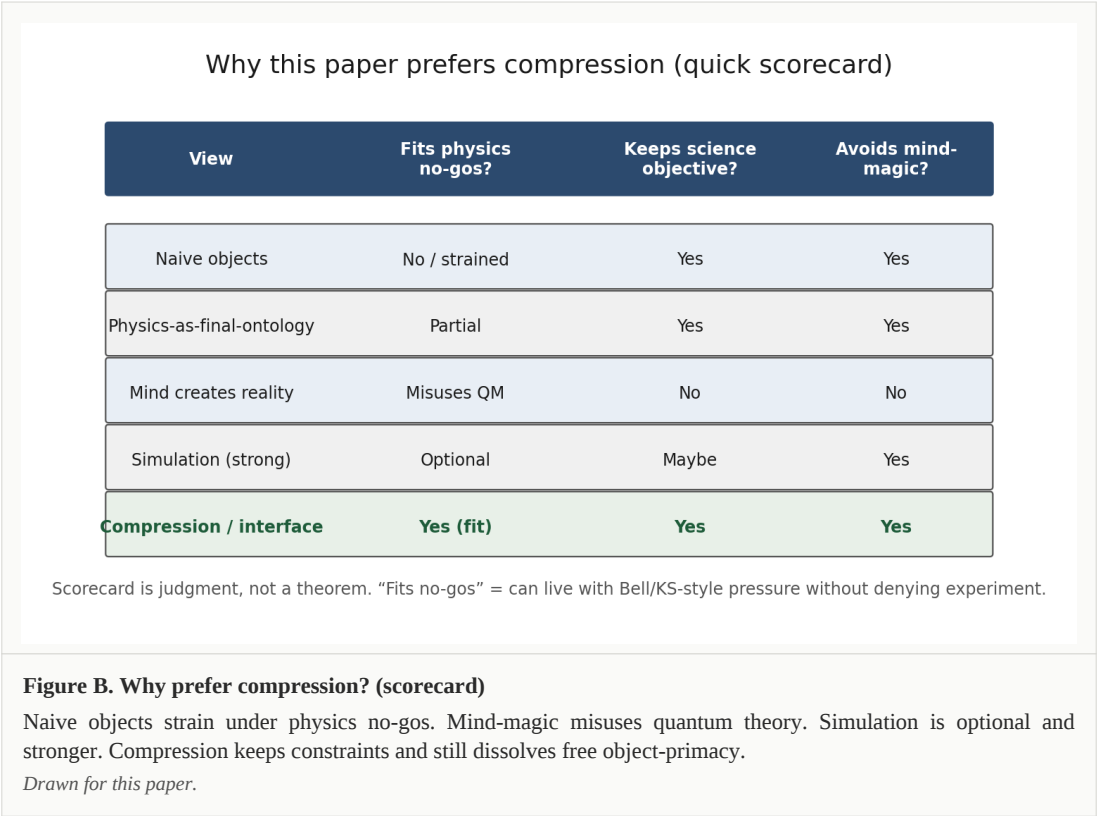


Figure 2. A real over-compressed JPEG.

Ugly, but you still recognize the scene. Format ate the fine grain; the picture did not become a hallucination.

Wikimedia Commons: JPEG compression Example.jpg (JJ Harrison, CC BY-SA 3.0).

Against that map of options, here is a blunt scorecard of why this paper prefers compression over the usual package deals. It is judgment, not a theorem.



1.4 What I will not argue

I will not argue that consciousness collapses the wave function. I will not argue that quantum mechanics licenses any claim about human intention affecting physical outcomes. I will not argue that synchronicity is an established phenomenon. I will not argue that anomalous aerial phenomena are non-human intelligence. I will not argue that psychedelic experience is a window onto the ground of being. Each of these is either false, unestablished, or outside what the evidence supports, and the credibility of the interesting thesis is repeatedly destroyed by people who cannot resist attaching the uninteresting ones to it.

Part I — Dissolution of the Object in Physics

This part is about what physics has made expensive to believe about objects and spacetime — not about consciousness magic. Roadmap: measurement (the textbook crack), Bell and KS (no free local or non-contextual values), PBR (psi not free ignorance), Wigner-type and Local Friendliness results (absolute facts not free), Wheeler (participatory language with cautions), interpretations that still work without interface metaphysics, then emergent spacetime as a separate pressure. A historical reminder first: the people who built quantum theory already knew the classical object-world was under pressure.




Figure 3e. Solvay, 1927.

Einstein, Bohr, Curie, Pauli, Heisenberg and the rest. History, not a proof of anything in this paper.

Wikimedia Commons: Solvay conference 1927.jpg (public domain).

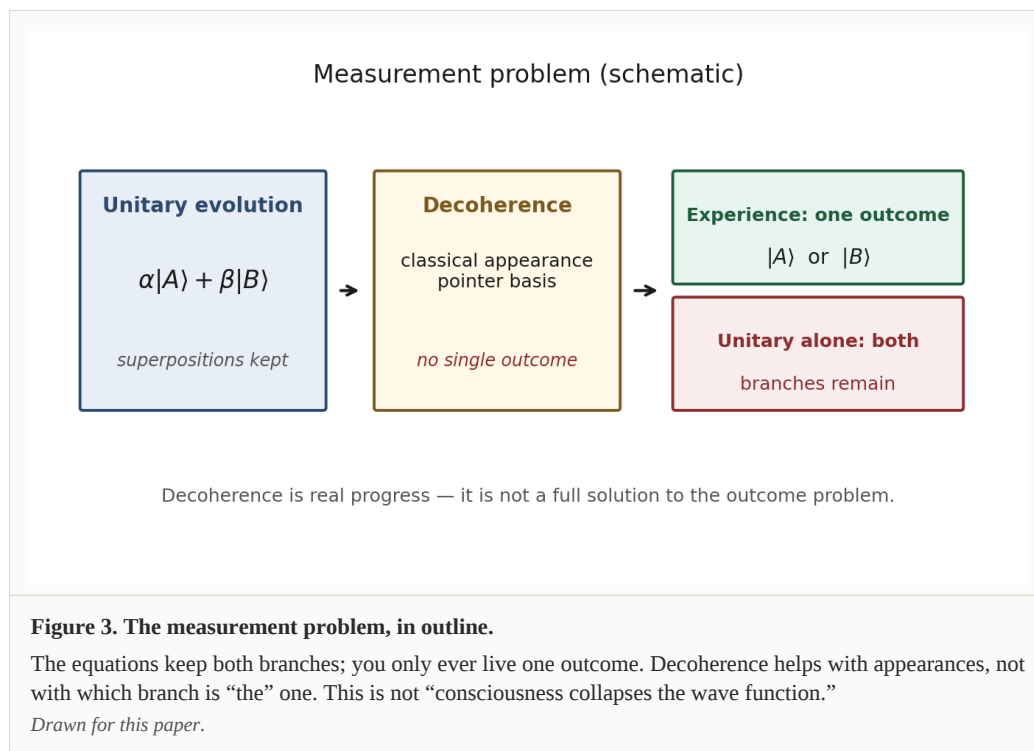
2. The measurement problem, stated precisely

Almost everything sloppy that gets said about quantum mechanics and consciousness begins with an imprecise statement of the measurement problem, so it is worth being pedantic.

In the standard textbook presentation, quantum theory is given two dynamical rules that do not fit together cleanly. Rule 1 (unitary evolution): between measurements, a quantum state evolves according to the Schrödinger equation — deterministic, continuous, reversible, and linear. Linearity means that if $|A\rangle$ evolves to $|A'\rangle$ and $|B\rangle$ evolves to $|B'\rangle$, then $a|A\rangle + b|B\rangle$ evolves to $a|A'\rangle + b|B'\rangle$. Superpositions are preserved; they do not “resolve themselves.” Rule 2 (the Born rule / projection postulate, in the old textbook packaging): on measurement, one records an eigenvalue with probability given by the squared modulus of the corresponding amplitude, and the state is treated as updated accordingly. That update is stochastic and irreversible in the textbook story. Modern treatments often replace ideal projection with more general measurements (POVMs), but the core tension remains: unitary dynamics alone does not pick a single recorded outcome.

The problem is not that Rule 2 is strange. The problem is that the theory does not say, from within unitary dynamics, when a “measurement” has happened. A measuring apparatus is made of quantum

systems. Model the apparatus unitarily and it can go into a superposition of pointer readings. Model the observer the same way and the same issue reappears. Nothing in pure unitary evolution selects one definite record. Yet experimenters write down single outcomes. That is the measurement problem: a clash between the linear dynamics and the appearance of unique results — not a license for “mind collapses the wave function.”



This must be distinguished sharply from classical ignorance. If I flip a coin under a cup, I do not know the outcome, but there is a fact of the matter and my ignorance is a fact about me. The measurement problem is not that. A superposition is not a probability distribution over definite states — as interference experiments demonstrate, the “unrealized” branches have physical consequences. This is the point that popular treatments almost always garble.





Figure 3b. Double-slit with light.


Two paths, interference fringes. You cannot talk as if the light always took one path and still explain the pattern.

Wikimedia Commons: Two-Slit Experiment Light.svg (inductiveload, public domain).

Decoherence does not solve this, and the honest literature says so. Environmental decoherence — the rapid entanglement of a system with its surroundings — explains, elegantly and quantitatively, why interference between macroscopically distinct branches becomes practically unobservable and why certain “pointer bases” are singled out (Zurek 2003). It explains the appearance of classicality. What it does not do is turn an improper mixture into a proper one, or select which branch is actual. Zurek himself is explicit that decoherence addresses the preferred-basis problem, not the outcome problem. Anyone who tells you decoherence dissolves the measurement problem is either using “measurement problem” to mean something narrower, or is wrong. Schlosshauer’s reviews of the literature make the same distinction carefully.



The cat cartoon is the cultural version of the same problem: nothing in the formal story forces a single “now” by itself. It is a teaching story, not a lab proposal and not a theory of mind.



3. What Bell’s theorem actually shows

John Bell’s 1964 result is the most misquoted theorem in physics (Bell 1964).


Bell showed that no theory meeting a package of assumptions — roughly local causality (no dependence of outcomes here on settings there outside the light cone, with correlations explained by a shared past cause), plus free or independent choice of settings, plus the usual probability framework — can reproduce the quantum predictions for suitable entangled systems. He derived inequalities (and later forms such as CHSH) that any such theory must satisfy and that quantum mechanics violates.

The experiments have been done, exhaustively. Early tests include Freedman and Clauser (1972); Aspect’s improvements (1981–82) tightened loopholes; then the decisive 2015 loophole-free runs: Hensen et al. at Delft with entangled electron spins in diamond separated by 1.3 km (Nature), and photonic experiments by Giustina et al. and Shalm et al. the same year. Aspect, Clauser and Zeilinger shared the 2022 Nobel Prize for this long program. Nature violates Bell inequalities under the tested conditions. That experimental fact is as solid as anything in physics. What it means for ontology remains contested — that is the next paragraph, not a loophole in the data.


Now the discipline. What Bell rules out is a local hidden-variable package of that type. What Bell does not rule out:

- **Nonlocal realism.** Bohmian mechanics is fully realist — particles have definite positions at all times — and reproduces the quantum statistics. It pays with explicit nonlocality. Bell was a serious student of Bohm’s theory and used the theorem partly to show what such a theory must cost (Bell 1987).
- **Many-worlds.** Everettian quantum mechanics keeps unitary dynamics and drops unique single outcomes as fundamental; branches carry the correlations without a nonlocal collapse story (Everett 1957; modern presentations e.g. Wallace 2012).
- **Denial of measurement independence** — superdeterminism, some retrocausal models, and related moves — which attack the “free settings / independent choices” side of the package rather than locality alone.
- **Any role for consciousness whatsoever.** Bell’s theorem is about settings and outcomes. The settings can be chosen by a random number generator or astronomical signals. No mind is required in the statement of the result.


The correct summary is: at least one ingredient of the local-causal, measurement-independent package must go. Which ingredient goes is what interpretations disagree about. Bell is not a proof of desktops, interfaces, or consciousness.



For a more standard lab sketch of the same idea — source, two-channel polarizers, coincidence counting — see the open diagram below. The moral is unchanged.



A different picture of non-classical correlation — not another Bell schematic — is entanglement swapping: pairs that never “met” can still end up linked. Hard to treat as a little object ferrying properties.



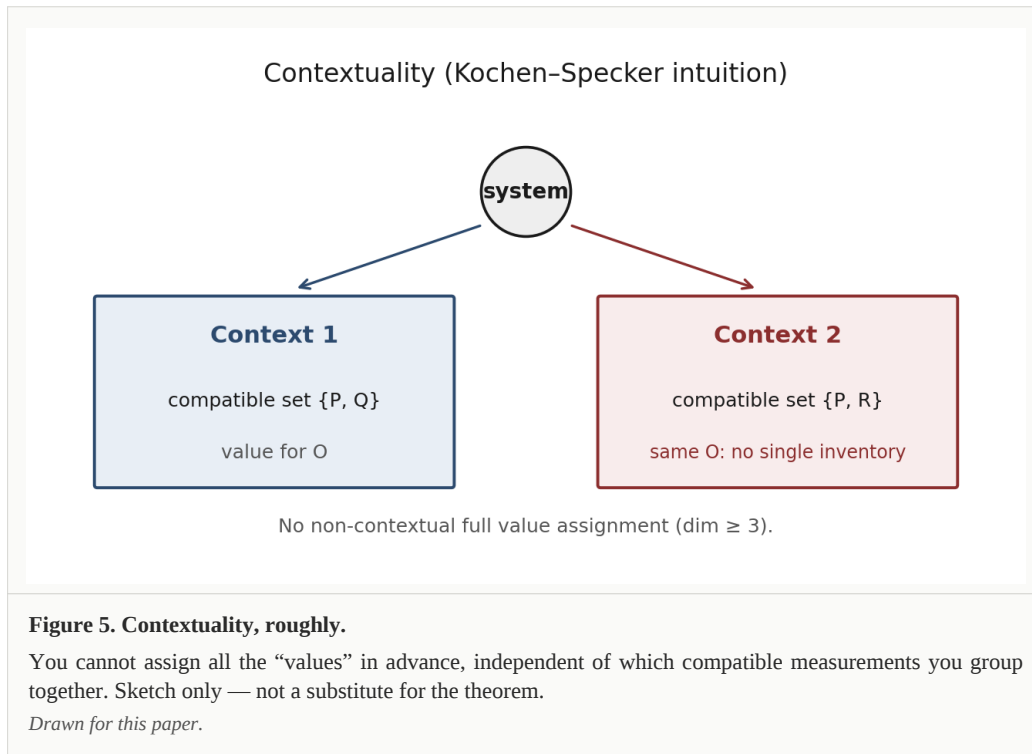
4. Kochen–Specker: the sharper knife

Bell gets the attention; for the purposes of this paper, the Kochen–Specker theorem is more directly relevant and considerably less discussed outside foundations (Kochen and Specker 1967; a version was independently proved by Bell in 1966).

KS shows that in a Hilbert space of dimension three or greater, it is impossible to assign definite values to all observables in a way that is non-contextual — that is, such that the value assigned to an observable does not depend on which other compatible observables are measured alongside it. There is no single consistent non-contextual value table. This is a theorem about the structure of Hilbert space; it needs no locality assumption and no distant parties.

State the implication carefully. It is not merely that we cannot know all values at once. Under the non-contextuality demand, there is no consistent assignment of values waiting to be known. That kills a certain classical picture of the system as a warehouse of pre-existing, context-free properties. It does not, by itself, say “all properties are only relations to observers” in a philosophical sense: contextual hidden-variable theories and other packages remain on the table. What dies is free non-contextual objecthood of the KS type.

Contextuality has been demonstrated experimentally (including state-independent forms) and is studied as a resource in quantum computation (e.g. Howard et al. 2014). That underlines the point: this is physics, not word-play. For this paper, KS is the sharpest single formal hit on free object-with-property-list realism in strand 1. It is not a proof of interface theory. It is a theorem that any CH-friendly story can live with and any naive object story must answer.



5. PBR: the wave function is not just ignorance

A natural retreat after Bell and KS is: fine, the wave function is not a description of the world, only of our knowledge. In the jargon, that is a psi-epistemic reading (in the Harrigan–Spekkens sense): distinct quantum states may correspond to overlapping distributions over some underlying ontic states.

The PBR theorem (Pusey, Barrett and Rudolph 2012) closes a large class of such models. Assume systems have ontic states, and that independently prepared systems have independent ontic states (preparation independence). Then distinct pure quantum states cannot be represented by overlapping ontic distributions in the intended way: if they did, quantum theory would assign probability zero to outcomes that the overlapping model would sometimes allow. Under those assumptions, the quantum state is not mere classical ignorance about an ordinary shared inventory.

Assumptions can be denied — preparation independence is the usual target — and restricted psi-epistemic constructions exist once one weakens the setup. So PBR is not the last word on every epistemic reading of psi. What it does is shrink the comfortable middle: you can treat psi as personal or normative (as QBism does), or treat it as ontic, but “just information about ordinary pre-existing properties, with independent preparations” is no longer free.

6. Wigner’s friend, and facts that are not absolute

Here the results get more delicate, and popular reporting is often worst.

Eugene Wigner’s thought experiment (1961): a friend in a sealed lab measures a system and records a definite outcome. Wigner, outside, can still write a unitary description of lab-plus-friend as a big superposition. Both descriptions are licensed by the theory; how they fit together is the puzzle.

Frauchiger and Renner (2018) built an extended multi-agent scenario and argued that three assumptions cannot all hold: (Q) agents may apply quantum theory to systems that include other agents; (C) agents’

conclusions hang together in a specified consistency sense; (S) measurements have single outcomes. Their title is blunt: “Quantum theory cannot consistently describe the use of itself.” Different interpretations drop different pieces. Critical literature disputes which inferences really follow from “quantum theory” as they package it; the safe use, which I adopt, is that every major reading gives up something — not that QM has been shown inconsistent in a vacuum.

Bong et al. (2020) formulated Local Friendliness (LF): roughly, that an observed event is an absolute fact for all observers, together with locality and free choice conditions in their sense. They show tension with quantum predictions and reported a photonic proof-of-principle (Nature Physics). Some presentations claim LF assumptions are weaker than Bell’s, so an LF violation would be a stronger constraint (Cavalcanti; Wiseman, Cavalcanti and Rieffel). That comparative packaging has been challenged (e.g. Okon). I note both sides and do not rest CH on “LF is strictly stronger than Bell.” Proietti et al. (2019) realized an extended Wigner’s-friend scenario with photonic “observers” (Science Advances). In all such experiments the “friends” are small quantum systems, not people, and the authors say so.

Discipline: these results do not show that human consciousness creates reality. They pressure the idea that every outcome is a single absolute fact in a view-from-nowhere ledger, free of cost in locality, freedom, or uniqueness of outcomes. That is interesting for any interface or relational story. It is not a derivation of the Compression Hypothesis. CH is about the format of experienced space, time, and objects under action constraints. Multi-agent quantum no-gos are a neighboring pressure on absolute factuality — allied in spirit if one already likes perspectival facts (RQM, some QBist themes), not a proof that perception is a desktop.

7. Wheeler: it from bit, and the participatory universe

John Archibald Wheeler — Feynman’s doctoral advisor, coiner of “black hole” and “wormhole” — spent later years arguing that information language sits close to the foundations. In “Information, Physics, Quantum: The Search for Links” (1989) he pressed the slogan that every it — particle, field, even the spacetime continuum — derives its function and meaning from apparatus-elicited answers to yes-or-no questions, bits. And, more provocatively: “No elementary phenomenon is a phenomenon until it is an observed (registered) phenomenon.”

Wheeler’s delayed-choice idea (1978; optical tests followed in later decades) is often mis-told. The setup lets one choose, after a photon has passed a first beamsplitter, whether the final apparatus will register which-path information or interference. Quantum mechanics predicts, and experiments confirm, the statistics appropriate to the final apparatus. That does not mean a little bullet “decided to be a particle or a wave in the past” and then rewrote history. It means the completed phenomenon — including the final measurement context — is what the theory describes. Complementary questions do not all get classical answers at once.

Two cautions, which matter. First, “registered” means recorded in an irreversible amplifying process; a plate or detector counts; a human mind is not required by the physics. Second, delayed choice does not allow signalling into the past. The correct reading is not “the future changes the past,” but “there was no single classical path-fact waiting in the drawer until the phenomenon was finished.” Wheeler’s participatory language is about system-plus-apparatus co-constitution of phenomena, not mind-over-matter. I use him as historical color and a warning against sloppy retrocausation talk — not as a fourth no-go theorem.

8. The interpretations, honestly

Any paper arguing in this direction owes the reader an honest account of the alternatives, because there exist fully coherent interpretations of quantum mechanics that require nothing like the Compression Hypothesis.

Bohmian mechanics (de Broglie–Bohm; Bohm 1952). Particles have definite positions always; a pilot wave on configuration space guides them; the quantum statistics are recovered. It is realist about trajectories and unambiguous about measurement records. Its costs include explicit nonlocality in the guidance equation and tension with relativistic covariance (a preferred foliation or equivalent structure is the usual price). Bohm’s later implicate-order writing (1980) is a different, more speculative layer: the explicate order of separate objects as an unfolding of a deeper enfolded process. That is historically interesting as a precursor mood for “objects not fundamental.” It is not the same claim as the modern holographic principle (area-law entropy, AdS/CFT), and I do not treat it as a proof of CH.

Everettian / many-worlds (Everett 1957; modern e.g. Wallace 2012). Keep unitary dynamics; drop unique single outcomes as fundamental. Superpositions continue; definite records are what a branch looks like from inside. Formally conservative, ontologically large. Fully realist about the universal wave function and typically local at the level of the wave function’s dynamics. It is the strongest standard counterexample to “quantum mechanics forces observer-dependence.” The open technical fight is the probability / Born-rule problem: what the weights mean if every branch occurs.

Objective collapse (GRW 1986; continuous spontaneous localization; gravitational proposals including Diósi–Penrose). Modify the dynamics so collapse is a real physical process. These are real theories with different predictions from textbook QM, so experiment can bite. Donadi et al. (Nature Physics 2021) reported strong constraints on a simple parameter-free Diósi–Penrose gravitational collapse model. Other collapse parameter ranges remain under test. That is what a program that can be wrong looks like — not a reason to ignore the class.

Relational Quantum Mechanics (Rovelli 1996). Quantum states are relative to physical systems; there is no single observer-independent state of the whole world. “Observer” means any system that can correlate, not a mind. RQM is the nearest mainstream physics cousin to “facts indexed to a cut.” It is not Hoffman, not Jung, and not a proof of compression. It is a live option that already weakens absolute object-state packages without my desktop metaphor.


QBism (Fuchs, Mermin, Schack and others). The quantum state is an agent’s personal probabilities for their own future experiences; the Born rule is treated as a normative coherence constraint. Update looks like Bayesian learning, not a physical collapse process. QBism is often misread as consciousness-causes-collapse; its authors reject that. Experience is where the theory’s advice lands; the agent is not said to shove the world by staring. I list QBism because it shows another serious way to deny that psi is a little object in the room — without my evolutionary story.

Honest bottom line on interpretations: these options remain live. No present experiment crowns one. A careful physicist can hold Everett or Bohm and treat the rest of this paper as optional philosophy. I think the broader argument (no free object-primacy + compression as a positive map) shifts the balance of interest. I do not claim it closes the question. Surveys of foundations attitudes (Schlosshauer, Kofler and Zeilinger 2013 and later polls) show a split, often disengaged field — not consensus for interface metaphysics.

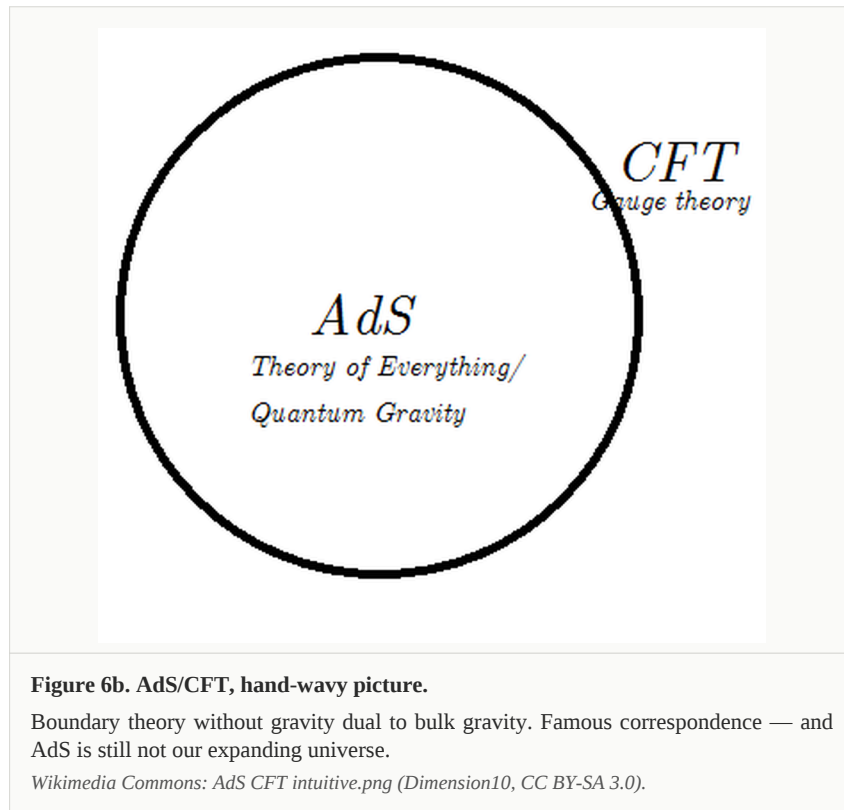
9. Emergent spacetime: the object dissolves from the other end

The most striking development is that a completely separate research program, motivated by black holes rather than by measurement, has independently concluded that spacetime is not fundamental.

The holographic principle. 't Hooft (1993) and Susskind (1995), reasoning from black hole thermodynamics — the Bekenstein–Hawking result that a black hole’s entropy scales with its horizon area, not its volume — proposed that the information content of any region is bounded by its boundary area. A three-dimensional volume’s physics can be encoded on a two-dimensional surface. Dimensionality itself starts to look like a description-dependent quantity.



AdS/CFT. Maldacena’s 1997/98 correspondence gave this a precise realization: a theory of quantum gravity in a $(d+1)$ -dimensional anti-de Sitter space is exactly equivalent to a conformal field theory without gravity on its d -dimensional boundary. Two descriptions, different dimensions, same physics. It is among the most-cited results in theoretical physics of the last thirty years.



Entanglement builds geometry. Van Raamsdonk (2010) showed that if you take a boundary state and progressively disentangle its degrees of freedom, the corresponding bulk spacetime pulls apart. Reduce entanglement to zero and the bulk splits into disconnected pieces. The natural reading: spatial connectivity is entanglement, seen from the bulk side. Ryu and Takayanagi’s earlier formula (2006), relating entanglement entropy to minimal surface areas, is the quantitative backbone. ER = EPR: Maldacena and Susskind (2013) conjectured that entangled particles (EPR) are connected by microscopic wormholes (ER). Nonlocality in the quantum description is geometry in the gravitational one.

“Spacetime is doomed.” Arkani-Hamed’s now-standard formulation in lectures on the amplituhedron program. The amplituhedron (Arkani-Hamed and Trnka 2014) computes scattering amplitudes in a particular supersymmetric gauge theory from a purely combinatorial-geometric object, with no reference to spacetime, locality, or unitarity — from which locality and unitarity emerge as consequences of its positive geometry. The program’s working slogan is that spacetime and quantum mechanics are both derivative of something more primitive that is combinatorial in character.

I want to be careful here, because this is where the paper is most easily accused of overreach. String theory is not experimentally confirmed. AdS/CFT applies to anti-de Sitter space, and our universe is de Sitter — the correspondence is a powerful toy model, not a description of the actual world. The amplituhedron program covers planar $N = 4$ super-Yang–Mills, which is not the Standard Model, though the program has been extended considerably since 2014. None of this establishes that our spacetime is emergent.

What it establishes is weaker and still important: that “spacetime is emergent” is a mathematically coherent, technically productive, mainstream research position, and that we now have explicit worked examples of universes in which it is true. That is enough to make the philosophical question live. Before

1997 one could reasonably say that the notion of emergent spacetime was word-play; after AdS/CFT one cannot.

One more discipline, essential for later parts of this paper: emergent spacetime in quantum gravity is not the same claim as “human perception is a desktop interface.” AdS/CFT does not say vision is an icon. The bridge from “geometry from entanglement in a CFT” to “experienced objects are fitness-tuned encodings” is philosophical synthesis, not a theorem. I will make that synthesis in Part IV. I will not pretend physics already made it for me.

10. What Part I establishes

Let me be exact about the ledger — three tiers, not one mushy “quantum says.”

Tier 1 — Theorem + experiment (standard assumptions stated in the literature): Local hidden-variable packages of Bell’s type are ruled out by theory and by loophole-free tests. Non-contextual full value assignments for all observables are impossible in $\dim \geq 3$ (Kochen–Specker). Under preparation independence, distinct pure states are not overlapping classical ignorance distributions in the PBR sense. These are the hard floor under strand 1.

Tier 2 — Live constraints and programs (important, more packaging debate): Wigner-type multi-agent scenarios and Local Friendliness press absolute, view-from-nowhere outcome facts; experiments use microscopic “friends,” not people. Comparative claims that LF is “strictly stronger than Bell” are contested; I use LF as a live pressure, not as a settled upgrade over Bell. Emergent spacetime (holography, AdS/CFT, entanglement–geometry) is a serious technical program with exact dualities in toy settings — not an established fact about our de Sitter universe, and not a theorem about human vision.

Tier 3 — Not established (refuse these free riders): Consciousness causes collapse. Human intention steers quantum outcomes. Orch-OR-style quantum consciousness as a settled account (Tegmark’s decoherence estimates remain a serious obstacle for simple microtubule coherence stories; the debate continues, and I do not build on it). Any ethics or spirituality package “implied by quantum mechanics.” That physics has already proved the Compression Hypothesis.

The honest conclusion of Part I is negative and limited: free naive object-primacy — local pre-existing values, non-contextual property warehouses, psi as free classical ignorance — is not a free lunch. Every surviving interpretation pays somewhere: nonlocality, branching, modified dynamics, relativized facts, or personalist readings of psi. That is the ledger the rest of the paper builds on. It is not yet a positive proof of desktops.

Part I-M — Mathematical Scaffolding: What Formal Results Establish

Can mathematics prove the Compression Hypothesis? No — not if “prove” means a single theorem that forces “spacetime is a lossy, action-oriented encoding of a non-object relational ground.” That package mixes physics, evolution, and ontology. Math does three different jobs here, and they should not be confused.

- **Job 1 — Constraints (theorems).** No-go results rule out free local hidden variables, free non-contextual value lists, and (under stated assumptions) free “psi is mere ignorance” stories. Part I argued these in prose; this section only pins the formal shape.
- **Job 2 — Definitions (not uniqueness theorems).** “Compression,” “interface map,” and “Markov blanket” can be stated cleanly so the hypothesis is not only a JPEG metaphor.
- **Job 3 — Model theorems with open empirical gates.** Fitness-beats-truth is a theorem *inside a model*. Whether nature sits in that model’s assumptions is not settled by the proof (Part II and Part VI).

If you want the long physics story, stay in Part I. If you want Hoffman and predictive processing in full, Parts II–IIb. What follows is the short formal spine — and an explicit list of what still is not proved.

M1. What “prove” would have to mean

A proof is relative to axioms. Bell proves that no theory meeting local causality (plus standard side conditions) can match quantum correlations. Kochen–Specker proves that no non-contextual value assignment covers all observables in $\dim \geq 3$. Those are proofs. The Compression Hypothesis is not that kind of object: it claims something about how experience is formatted and what kind of ground it compresses. To “prove CH” you would need axioms that already smuggle in perception, action, and ontology — which begs the question. So math here is scaffolding and constraint, not a uniqueness proof of one ontology.

M2. No-go constraints (compact formal restatement of Part I)

This subsection does not re-argue §§3–6. It only records the minimal formal statements so later claims can point to them without hand-waving.

Bell / CHSH. If outcomes A, B and settings a, b admit a local hidden-variable factorization

$$P(A, B \mid a, b) = \int d(\lambda) \rho(\lambda) P(A \mid a, \lambda) P(B \mid b, \lambda),$$

then certain correlator combinations stay inside Bell/CHSH bounds. Quantum theory predicts violations; experiment confirms them. So at least one of locality, “outcome independent of distant setting,” or classical shared past (ρ) fails. Proved relative to the assumptions. Not a proof of consciousness, desktops, or any single interpretation.

Kochen–Specker. There is no map v from observables to values that (i) lands in the spectrum and (ii) respects functional relations for every compatible family, independent of which family is chosen. In $\dim H \geq 3$ this is impossible. Free “object with a full non-contextual property list” is inconsistent with Hilbert-space QM. Bohm, Everett, and other packages remain open — they just pay a different price.

PBR (sketch). Under preparation independence, distinct pure states cannot share overlapping ontic support in the intended way: psi is not mere classical ignorance about an ordinary inventory. Deny the assumptions and the door reopens; within them, the middle psi-epistemic retreat is closed.

Local Friendliness / Wigner-type no-gos (status note). Part I recorded these as live constraints on absolute observed facts plus locality and free choice. They are less settled in packaging than Bell/KS; they belong in the ledger as active research constraints, not as a second copy of Bell. I do not re-derive them here.

Logical role for CH: these results pressure free object-primacy. They do not select interface theory over Everett, Bohm, QBism, or RQM. Anyone who says “Bell proved we live in a desktop” is abusing a theorem.


M3. Compression, formally (the part that is not already in Part I)

This is the load-bearing formal work of the section. To make “lossy, action-oriented encoding” more than a metaphor, fix three spaces and two maps. Let W be world (or ground) states, X interface or percept states, and A actions. A perceptual map and a policy are

$$p : W \rightarrow X, \pi : X \rightarrow A$$

(deterministic for simplicity; the stochastic case uses kernels). A fitness or payoff is $f : W \times A \rightarrow \mathbb{R}$. Expected fitness under a distribution μ on W is

$$F(p, \pi) = \sum_w \mu(w) * f(w, \pi(p(w))).$$



Call p *structure-preserving* (a strong “truth” idealization) if it is injective or an isomorphism onto its image for some relational structure on W . Call p a *lossy compression under f* if it discards distinctions that do not change the optimal action class — roughly, if $p(w) = p(w')$ whenever the argmax sets of a $|-> f(w, a)$ and a $|-> f(w', a)$ coincide.

Rate–distortion theory says the same thing in information language: minimize $I(W; X)$ subject to $E[d(W, X)] \leq D$, with d chosen to penalize fitness-relevant errors, not pixel fidelity. None of this proves that human spacetime is such a p . It shows that “compression” is a precise class of maps, and that isomorphism is a

special case of that class. The Compression Hypothesis is the claim that bounded agents under selection implement maps in the lossy, action-tuned subclass — and that space, time, and objecthood are features of X ’s format, not of W .

M4. FBT and Markov blankets (pointers, not re-lectures)

Fitness-beats-truth. Prakash et al. (2021) prove, in their model, that when fitness functions are drawn from a broad (“generic”) class and $|W|$ grows, structure-preserving strategies lose to suitably tuned interface strategies with probability tending to one. That is a theorem about the model. It is not yet a theorem about Earth until one checks how often real fitness is generic rather than monotone in latent variables (Part VI). Critical limits (Bagwell 2023; Mattes 2026) are argued in §13; I do not rehash them. Used carefully, FBT supports only this: evolution does not guarantee isomorphism. That is the compression-friendly reading, not “perception is fiction.”

Markov blankets. Partition states into internal μ , sensory s , active a , and external ψ . A blanket condition is conditional independence

$$\mu \perp\!\!\!\perp \psi \mid \{s, a\}.$$

Internal systems only “see” the outside through $\{s, a\}$. That is the formal avatar of meeting the world at an interface (Part IIb develops the science). Free-energy / active-inference dynamics add that internal states minimize a variational bound on surprise. Logical limit: blankets and free energy scaffold the *weak* claim (encoding at a boundary). “Spacetime is format” is an extra interpretive step — not a corollary of conditional independence alone.

M5. Ledger: proved vs not proved

Claim	Status
No local HV of Bell type (assumptions + experiment)	Theorem + experiment
No non-contextual full value assignment ($\dim \geq 3$)	Theorem (KS)
ψ not mere classical ignorance (prep. independence)	Theorem (PBR)
Absolute facts + locality + free choice (LF package)	Active no-go program; packaging debated
Lossy fitness-tuned maps beat isomorphism under generic f	Theorem in FBT model only
Access to environment only via blanket states	Definition / conditional independence
“Compression” is a well-defined class of maps	Definition (rate–distortion / fitness)
Human spacetime / objects are interface format	Conjecture — not a theorem
Relational ground without pre-cut objects	Structural / philosophical hypothesis
Consciousness collapses the wave function	Not used; not supported here

Table M1. Formal status. The Compression Hypothesis lives in the lower rows: constrained and motivated by the upper rows, not deduced from them.

Bottom line. Math can kill free object-primacy in important senses, and it can define compression and interfaces so the hypothesis is not fog. It cannot, by itself, prove CH is the unique true ontology. The dual abuse is also wrong: that no math bears on this at all. Bell, KS, rate–distortion structure, FBT-in-model, and blanket formalisms all bite — each in its own job, not as one mega-proof.

Part II — The Interface: Hoffman and the Evolutionary Argument

11. The argument from the other direction


Part I attacked the object from physics. Donald Hoffman attacks it from evolutionary biology, and the interest of his argument is that it needs no quantum mechanics at all. If it works, it works on entirely classical premises — which makes its agreement with strand 1 non-trivial.

The intuition is simple enough to state in a sentence, which is both its strength and the source of most misunderstanding of it. Natural selection optimizes for fitness, and fitness is not truth.

12. The Interface Theory of Perception

Hoffman, Singh and Prakash set out the position formally in “The Interface Theory of Perception” (Psychonomic Bulletin & Review, 2015). The claim is that perception evolved as a user interface rather than as a window.

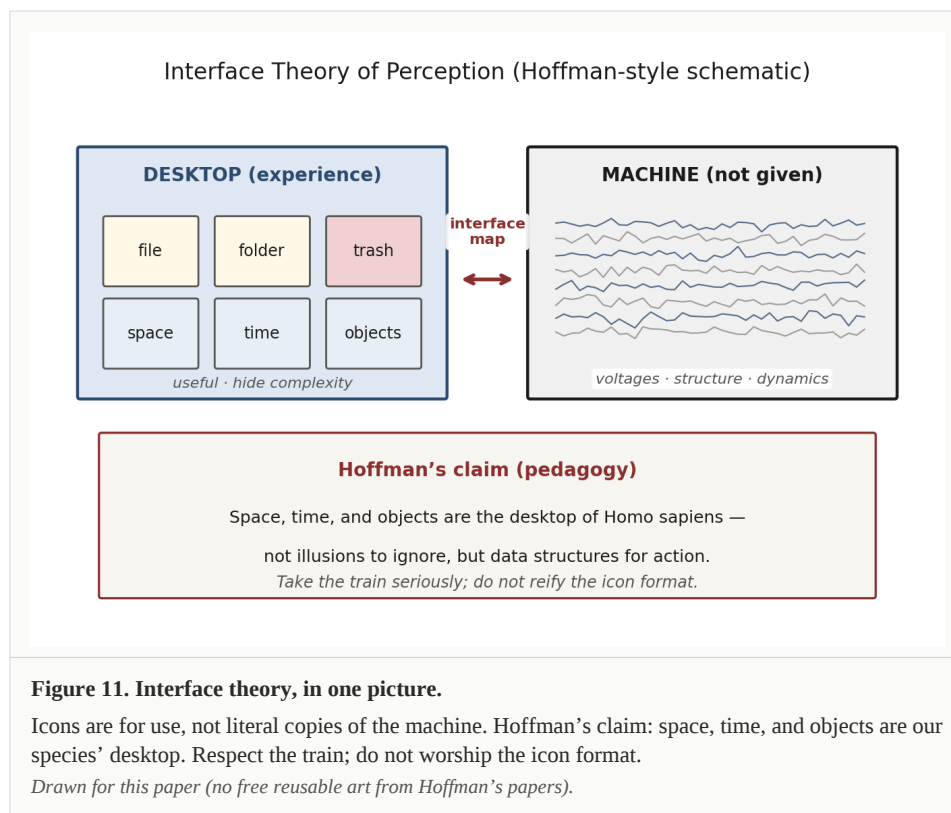
The metaphor is the desktop. A file icon is blue, rectangular, and in the lower-left corner of the screen. None of these is a property of the file. The file is a voltage pattern distributed across a storage medium; it is not blue, not rectangular, not in the lower left of anything. Yet the icon is not a lie. It is a functional interface, and dragging it to the trash really does delete the file. The interface is useful precisely because it hides the complexity of the machine — an interface that showed you the actual voltages would make you a worse user, not a better one.



Hoffman’s claim is that space, time, and physical objects are the desktop of Homo sapiens. Not illusions to be dispelled — you should take a speeding train as seriously as you take the trash icon — but data structures whose format reflects the demands of action rather than the structure of the world.

The relevant lineage matters. This is a broadly Kantian position, and Hoffman says so: the phenomenal world is structured by the constitution of the perceiver, and the thing-in-itself is not given. What Hoffman

adds to Kant is a mechanism (natural selection) and a proof strategy (evolutionary game theory), which is a substantive addition — Kant had a transcendental argument; Hoffman claims to have a mathematical one.



13. The Fitness-Beats-Truth theorem, precisely

Because this theorem is routinely overstated — including, in my judgement, by Hoffman in popular presentations — it is worth stating carefully what it does and does not say. The formal result appears as Prakash, Stephens, Hoffman, Singh and Fields, “Fitness Beats Truth in the Evolution of Perception” (*Acta Biotheoretica* 69, 2021), with earlier evolutionary-game simulations in the 2015 paper and elsewhere.

Setup. A world W with a finite number of states, equipped with some structure. An organism has a perceptual map from world states to perceptual states, and acts on the basis of perceptual states. A veridical (or “truth”) strategy is one whose perceptual map preserves the structure of W — roughly, an isomorphism or homomorphism onto the perceptual space. An interface strategy is one tuned to a fitness function that need not respect W 's structure at all. Fitness functions are drawn generically — that is, without assuming any special relation to the world's structure.

Result. As the number of world states grows, the probability that a veridical strategy has higher expected fitness than a fitness-tuned interface strategy tends to zero. In the authors' formulation, veridical perception is driven to extinction by non-veridical strategies tuned to fitness, generically and with probability approaching one.

Now the assumptions, which do a great deal of the work.

1. Fitness functions are generic. This is the crux. The theorem samples fitness functions from a distribution in which they bear no systematic relation to world structure. But real fitness functions are not

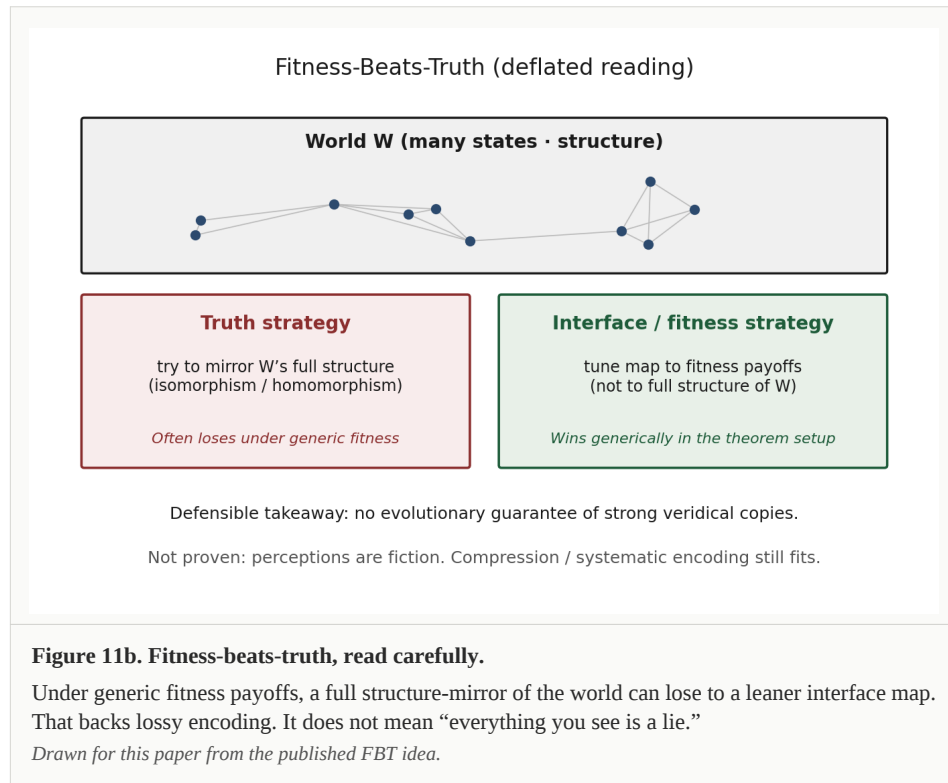
random draws from the space of possible functions — they are themselves products of the world’s structure. Nutritional value is not an arbitrary function of chemistry; predator lethality is not an arbitrary function of kinematics. Critics have argued, in my view correctly, that ecologically realistic fitness functions are frequently monotone in world variables, and that in the monotone case truth-tracking strategies do well. The theorem’s own authors acknowledge the monotone case as an exception; the dispute is about how large that exception is in nature. That dispute is empirical and still underdone (Part VI).

2. Perceptual cost is not the driver. The result is not merely the uncontroversial claim that veridical perception is metabolically expensive. It is a stronger claim about expected fitness even setting cost aside.

3. “Veridical” means structure-preserving in a strong sense. Here is the objection I consider strongest. There is an enormous space between “isomorphic to the world” and “unrelated to the world.” An interface can be systematically homomorphic on the fitness-relevant substructure — preserving orderings, adjacencies, causal relations — while being non-isomorphic overall. Such an interface is not veridical in the theorem’s technical sense but is massively truth-tracking in every sense we ordinarily care about. The theorem, on this reading, refutes perceptual isomorphism, which almost nobody held, while leaving perceptual reliability untouched.

I think this objection lands, and I think it is the most important limitation on strand 3. The FBT theorem does not establish that our perceptions are unrelated to reality; it establishes that they are unlikely to be structure-preserving copies of it. That is a weaker conclusion. But note carefully — it is still exactly the conclusion the Compression Hypothesis needs. A compression is precisely a non-isomorphic but systematically related encoding. Hoffman’s theorem, read in its defensible form, is an argument that perception is lossy compression, which is what I want, rather than fiction, which I do not need and do not believe.

Recent critical literature must be on the page. Bagwell (2023), in *Synthese*, argues that Hoffman-style evolutionary debunking of ordinary objects is self-defeating: the conclusion undercuts the scientific premises that support the argument, and appeals to “universal Darwinism” lead to a fatal dilemma. Mattes (2026), in *Biological Theory* (“The Case Against The Case Against Reality”), argues that once fitness functions are allowed to evolve in more realistic ways, the simulations and theorem fail to doom approximate perceptual realism — fitness bears truth rather than simply beats it. I take both seriously. They hurt Hoffman’s strong popular claim more than they hurt a deflated compression reading, but they forbid complacency. Anyone who cites FBT without citing its critics is doing advocacy, not assessment.

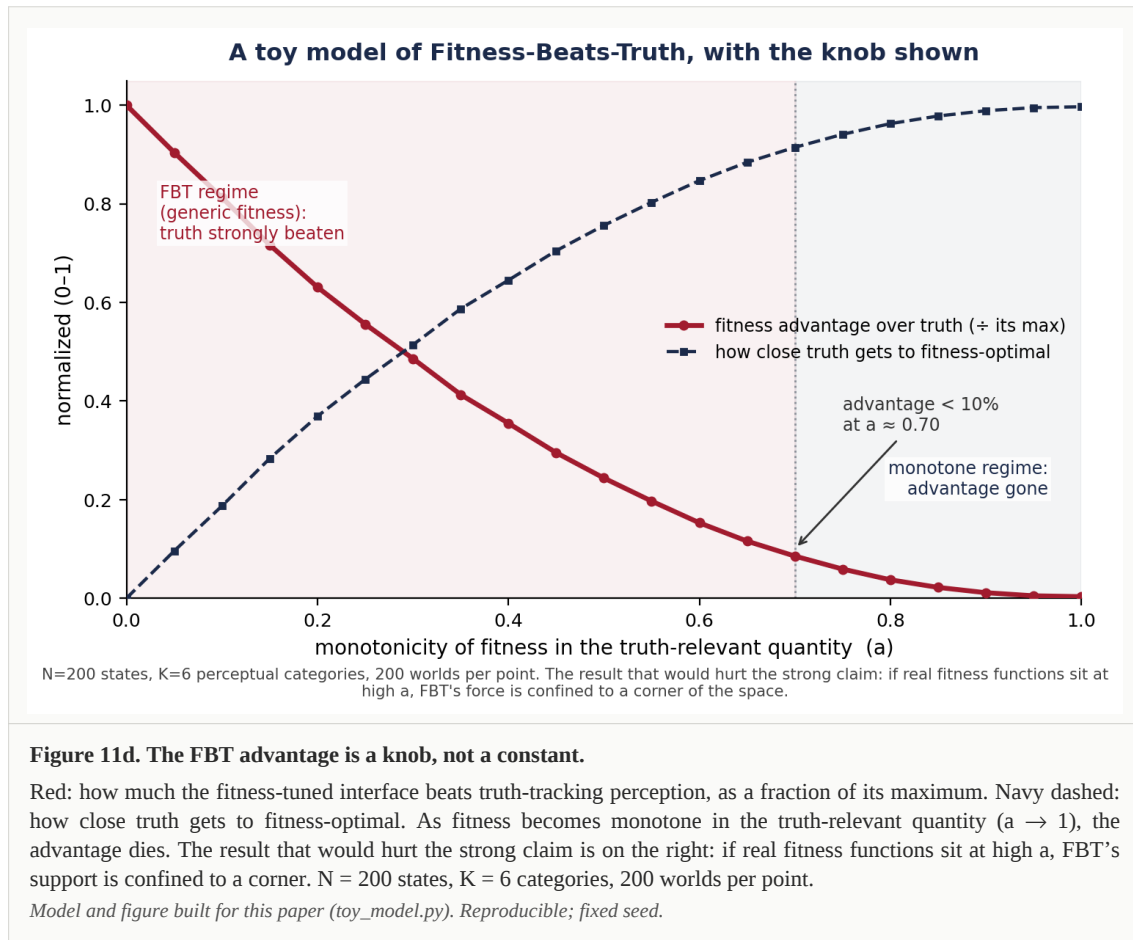


13b. A toy model with the knob shown

The genericity objection deserves better than prose, so here is a small model that makes it a thing you can turn a knob on. It is a cartoon, not a proof, and I have built it so that it can come out against me.

A world has many states, each carrying a hidden quantity that actually matters — call it a resource level. Fitness is a blend of that quantity and a purely generic component unrelated to it, mixed by a single dial a . At $a = 0$ fitness is generic — the regime the FBT theorem samples, where payoff bears no relation to world structure. At $a = 1$ fitness is exactly the resource level — the realist's regime, where what pays and what is true coincide. Real ecologies live somewhere between, and where they live is an empirical question, not a theorem. Perceivers are bounded: they have only a handful of perceptual categories, far fewer than there are states, and must commit to one action per category — the point of an interface is that it cannot afford to resolve everything. A truth strategy sorts states into categories by their resource level (perceive the world's own structure). A fitness strategy sorts them to group states of similar payoff (perceive whatever pays). Each pays for the mismatch it is forced into when unlike states fall in one category.

Sweeping the dial and asking how badly truth is beaten produces the figure. At $a = 0$ the fitness-tuned interface wins by its full margin; the truth strategy is at zero efficiency, exactly the FBT result. As fitness becomes monotone in the resource level, that advantage collapses, and by $a \approx 0.7$ it has fallen below a tenth of its maximum: truth-tracking perception has essentially caught up. The FBT advantage is not a constant. It is a function of one number — how monotone real fitness is — and that number is measurable, not stipulable.



This does two honest things at once. It shows the FBT argument is real where it applies — generic fitness genuinely crushes truth-tracking, and the desktop intuition is not empty. And it shows exactly which measurement would deflate it: estimate the monotonicity of real fitness landscapes (Part VI, sketch A), and you locate the world on this axis. If ecologies cluster near $a = 1$, the evolutionary strand offers little to the strong Compression claim, and I would have to say so. A model you cannot lose with is not a model. This one has a losing side, drawn in.

14. Conscious realism and the conscious-agent formalism

Hoffman's positive theory goes considerably further than the negative argument, and here the paper must slow down and become more sceptical.

In "Objects of Consciousness" (Hoffman and Prakash, *Frontiers in Psychology*, 2014), the proposal is conscious realism: the objective world consists of conscious agents and their experiences, and spacetime and objects are the interface through which agents represent one another. A conscious agent is defined as a structure involving spaces of experiences and actions and Markov kernels cycling among them. The formalism is deliberately minimal: it says that a conscious agent is a certain kind of stochastic process with an experience space, and nothing about what experience is.

Two properties are claimed. First, combination: networks of conscious agents can themselves constitute a conscious agent, offered as an answer to the combination problem that plagues panpsychism. Second, projection: spacetime and physical objects arise as a low-dimensional projection of the dynamics of the

agent network. In later work (Hoffman, Prakash and Prentner 2023) the program is extended toward combinatorial objects related to the amplituhedron. The proposal is a bridge from a consciousness-first ontology into the emergent-spacetime program.

I find the negative argument in §13 substantially persuasive in deflated form and the positive program in §14 substantially unpersuasive as currently developed, and I think a paper of this kind is worthless if it does not say so. The formalism postulates an experience space and says little about why any element of it has the character it has. That relocates the hard problem into an axiom. A formalism earns its keep by generating predictions or by explaining something previously unexplained, and to date the conscious-agents framework has done neither in a way that has been independently checked. The decorated-permutations bridge has not been taken up by the amplitudes community; absence of engagement over years is weak but real evidence. For this paper: use ITP and FBT as pressure against naive veridicality; quarantine conscious realism as a separate, optional ontology I do not adopt.

15. Self-undermining and related objections

If perception is systematically non-veridical, and evolutionary theory is itself constructed from perceptual evidence — fossils seen, genomes sequenced by instruments read with eyes — then the argument saws off the branch it sits on. This is structurally related to Plantinga’s Evolutionary Argument Against Naturalism, and it is a serious problem. Hoffman’s reply is that mathematical structure survives even if perceptual content does not: the theory is a formal object, and formal reasoning is not perception. I think this reply works only partially. It saves the mathematics; it does not obviously save the empirical premises that make the mathematics apply to us. At minimum, the interface theorist owes an account of why some of our cognitive outputs (mathematics, evolutionary theory) escape the interface while others (space, objects) do not, and I have not seen a satisfying one. Bagwell’s 2023 development of the self-defeat problem is the best recent sharpening of this knife. I return to it under Objections and flag it as unresolved rather than solved.

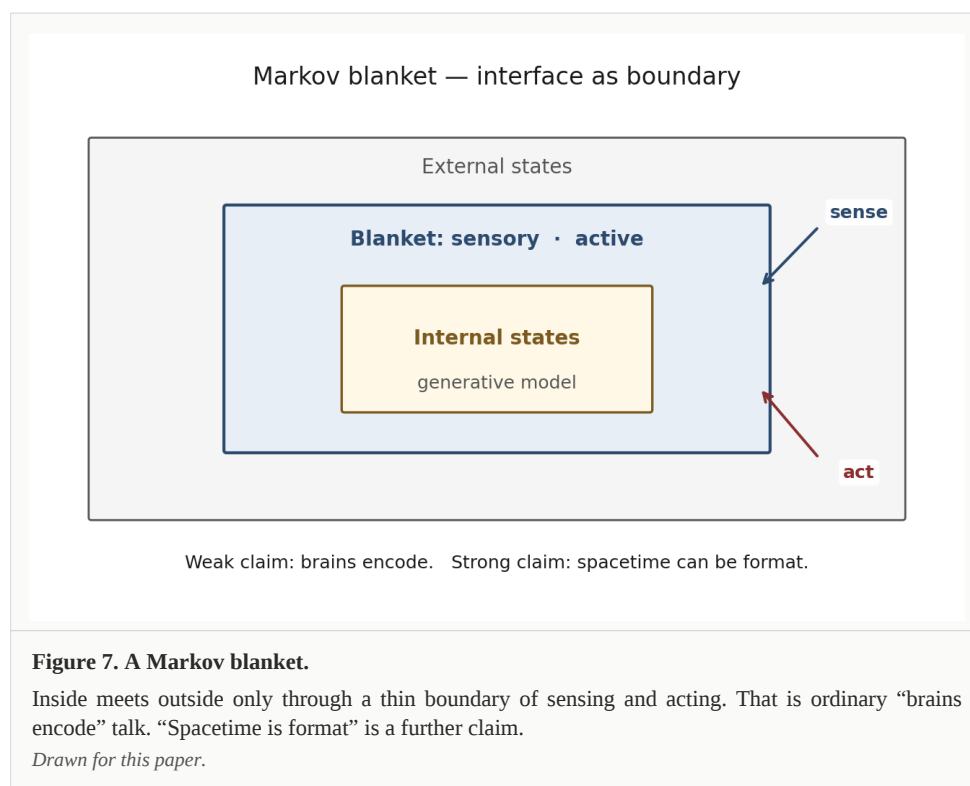
What survives, and it is not nothing: a rigorous, published, mathematically explicit argument that there is no evolutionary guarantee of perceptual veridicality in the strong sense, and that the burden of proof runs the other way from where naive realism assumes. That is a genuine contribution, independent of physics, pointing in a direction physics also points — when both are read carefully.

Part IIb — Predictive Processing, Affordances, and Plasticity

16. Compression needs a scientific backbone

A desktop metaphor is pedagogy. It is not yet a theory. The nearest mainstream scientific language for lossy, action-oriented encoding is predictive processing and the free-energy principle (Friston 2010; the broader predictive-mind literature associated with Andy Clark and others; Anil Seth’s “controlled hallucination” slogan for ordinary perception).

On these views, a living system maintains itself by minimizing surprise — or an upper bound on surprise, free energy — relative to a generative model of its sensory causes. Perception is inference. Action is a way of making sensations conform to predictions. The system is separated from its environment by a statistical boundary — a Markov blanket — and only ever meets the world through that boundary. That is interface talk with equations attached, and it does not require any commitment to quantum foundations or dual-aspect monism.



Two claims must be kept apart, and I will keep repeating the distinction because it is where this paper is easiest to misread. **Weak claim:** brains and other self-organizing systems compress and predict; experience is a controlled construction constrained by sensory input. That claim is close to mainstream cognitive science. **Strong claim:** spacetime and objecthood themselves are features of the encoding format, not of the encoded structure. That is the Compression Hypothesis. Predictive processing supports the weak claim directly and makes the strong claim more natural. It does not prove the strong claim. A physicalist can accept generative models all day without accepting that spacetime is an icon.

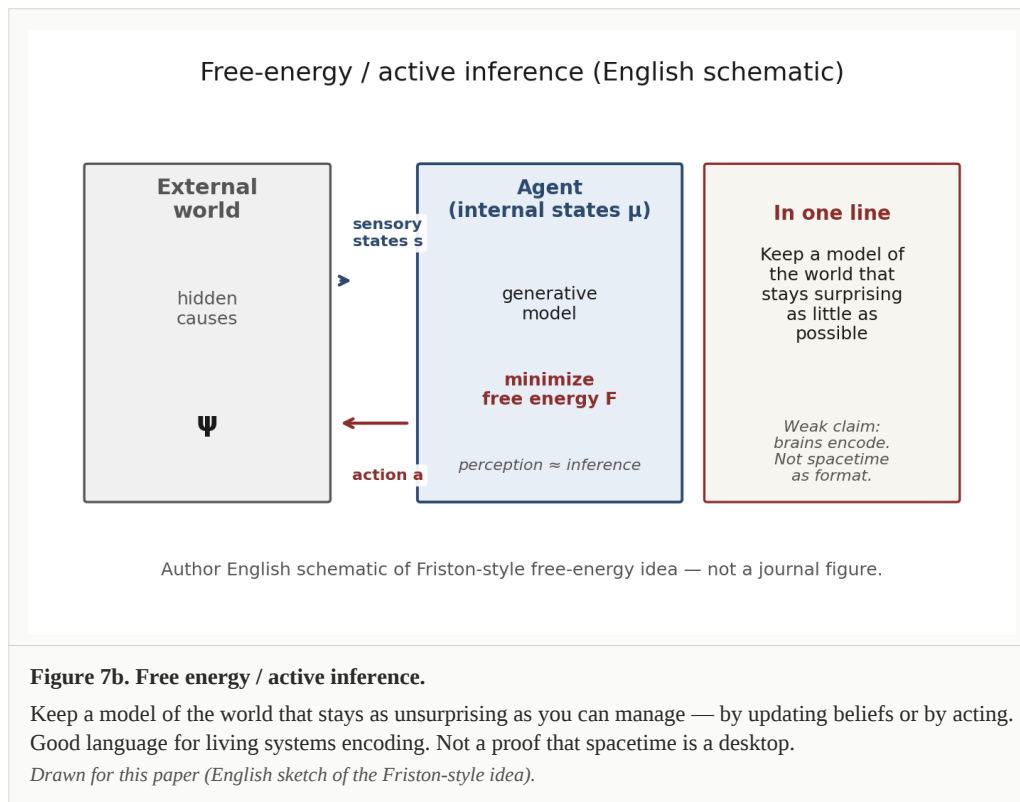
A hostile referee will put this sharply: the empirically live half of the paper is not the Compression Hypothesis, and the Compression Hypothesis is the half that is not empirically live. That is fair, and it

deserves a table rather than a dodge. If the strong claim is to be more than mood, it must be possible to state what would kill it while leaving the weak claim standing — otherwise the two are not really distinct, and the borrowed prestige of cognitive science is a confidence trick. Here is that table.

Observation / result	Weak claim (brains encode)	Strong claim (spacetime is format)
Neuroscience confirms predictive, generative-model perception in ever more detail	Confirmed	Consistent, but not confirmed — this was never the disputed part
Real fitness landscapes turn out predominantly monotone in trackable world variables (Part II model; VI-A)	Untouched	Weakened — the evolutionary support for format-level non-veridicality erodes
Emergent-spacetime programs fail to extend beyond anti-de Sitter toy universes	Untouched	Weakened — strand 2 loses its claim on our world
A physical objective-collapse theory (GRW/CSL, gravitational) is confirmed	Untouched	Badly hurt — collapse is then a mind-independent physical process; the interface is not where definiteness is made
Perceptual/representational structure is shown to track environmental geometry independent of action repertoire (VI-B)	Untouched	Hurt — the action-constraint that fixes the format loses its evidence
Born rule is derived from a fully realist, observer-free axiom set	Untouched	Hurt — the strongest motive for an observer-involving ontology is removed

Table A. What would kill the strong claim but not the weak one. Every row in the third column that reads “weakened” or “hurt” is a way this paper can lose that leaves cognitive science untouched — which is what makes the strong claim a claim and not a borrowed halo.

The table also disciplines the paper’s own enthusiasm. Notice that no row in the weak-claim column ever reads worse than “untouched.” That is the point: predictive processing is not evidence for the Compression Hypothesis, it is the mainstream floor the hypothesis stands on and tries to reach past. The strong claim earns its keep only in the third column, and only by being killable there.

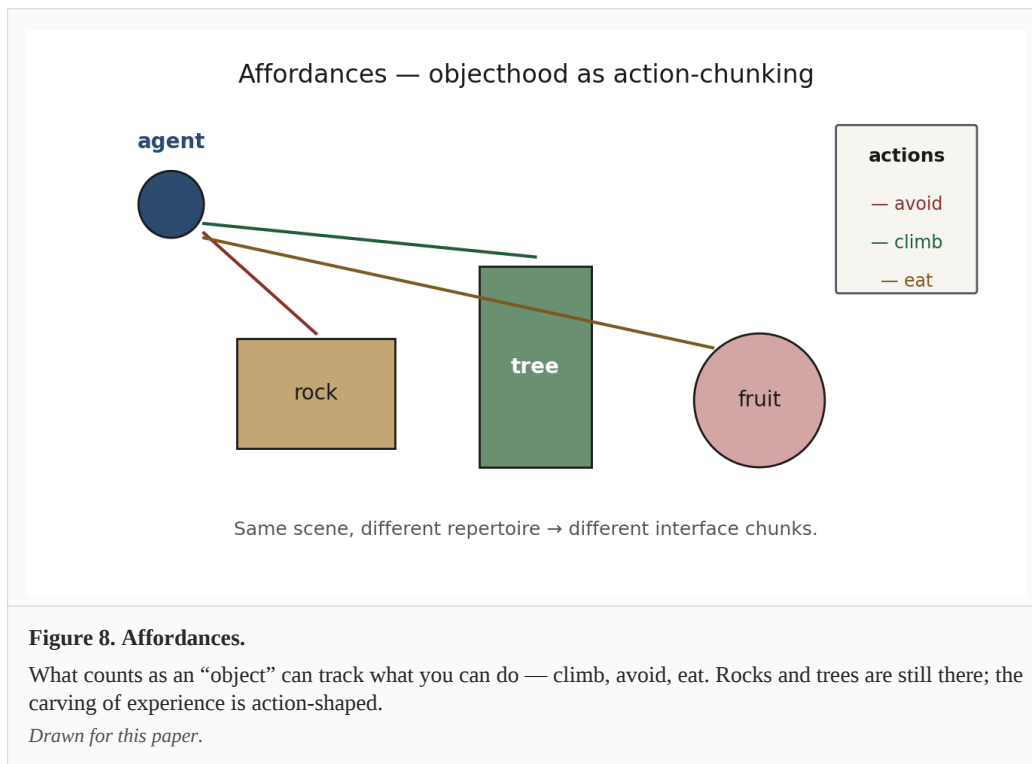


What “compression” means here, a little more formally: a map from a high-complexity relational domain to a lower-complexity representation optimized under an action or fitness constraint, preserving task-relevant mutual information at the expense of isomorphism. Rate–distortion theory is the information-theoretic skeleton. Evolutionary game theory and active inference are two ways the distortion measure gets fixed by life rather than by an engineer. This is still a sketch, not a derived theorem connecting FBT to free energy. It is the right sketch.

17. Affordances and objecthood

James Gibson’s ecological psychology treated perception as detection of affordances — what the environment offers an animal for action — rather than as construction of an internal picture of object properties. Ecological realists will insist that affordances are real relations in the world, not icons. I can accept that and still hold the Compression Hypothesis. The sweet spot is: real relational structure, non-object format. Objecthood, on the action-constraint reading, is what a manipulable, avoidable, edible, or navigable chunk of the world looks like when encoded by something that manipulates, avoids, eats, and navigates. That is Gibsonian in spirit even when it disagrees with Gibsonian realism about the full content of experience.


This matters for intersubjective agreement later. Shared action constraints in a shared environment produce correlated encodings without requiring that any encoding be a structure-preserving copy of a ready-made object inventory.



18. Plasticity of the format

If space, time, and objects are format, the format should be state-dependent. Dreams already show a full object-world without a matching external scene. Sensory substitution reorganizes spatial skills. Virtual reality (and Chalmers’ Reality+ style arguments) force questions about what “real presence” tracks. General anesthesia removes the render and restores it. None of these prove the Compression Hypothesis. All of them show that the ordinary desktop is not the only possible encoding a human system can run — and that feeling absolute is cheap. Dreams feel absolute while they last.

Chemically altered perception belongs in the same folder: plasticity data, not ontology. I treat it at length in Appendix B so it cannot contaminate Parts I–III. The short version: classic psychedelics perturb high-level priors and large-scale network dynamics; they produce alternative encodings that can feel more real than waking life; that feeling is not a truth-tracker; they do not hand you the unsplit ground.



Part III — Dual-Aspect Monism

19. Why this strand is here, and how much it weighs

A reader who has followed Parts I and II may reasonably balk. Bell’s theorem is a theorem; Jungian psychology is not. Including them in one argument invites exactly the category error I complained about in §1.1.

The reason to include it is that the fourth strand is not primarily an empirical claim about coincidences. It is a structural proposal about the relation between mind and matter, it was developed in collaboration with one of the most rigorous physicists of the twentieth century, and it has been reconstructed over the past three decades in a form precise enough to be criticized properly. The empirical claim (synchronicity) is the weakest part of it and I will argue that the structural claim can survive its collapse — while still refusing to weight this strand like Kochen–Specker. It is an ancestor and a parallel. Draft 1 sometimes let historical richness suggest coequal status. Draft 2.1 does not.

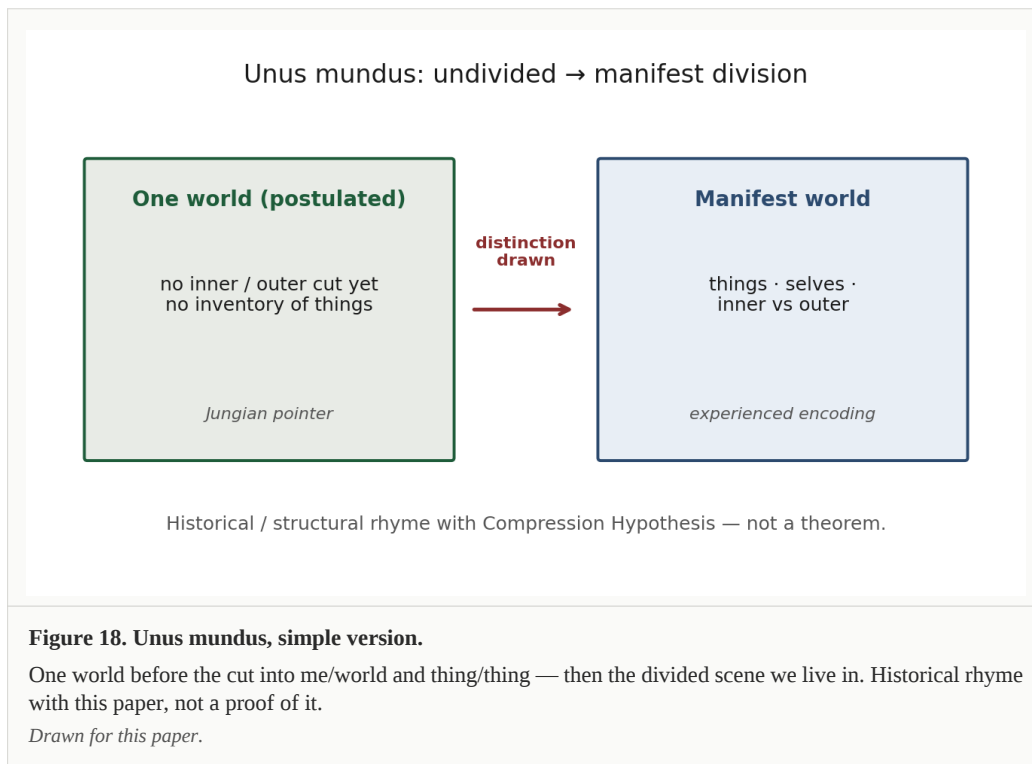
20. What Jung actually claimed

The *unus mundus* is not Jung’s coinage. It comes from Gerhard Dorn, a sixteenth-century alchemist and follower of Paracelsus, and it names the “one world” — the unitary, pre-differentiated ground from which the manifest world of distinct things arises. Jung takes it up principally in *Mysterium Coniunctionis* (CW 14, 1955–56), the last of his major works.


Jung’s version is this: psyche and matter are not two substances that somehow interact, nor is one reducible to the other. They are two manifestations of a single underlying reality which is in itself neither psychic nor physical. He writes that the *unus mundus* is the “potential world outside time,” the ground in which the distinction between inner and outer has not yet been drawn.

The crucial technical notion is the psychoid archetype. In the later work — especially “On the Nature of the Psyche” (CW 8) — Jung insists that archetypes are not contents of the psyche but formal ordering principles that are “psychoid”: transgressive of the psyche–matter boundary, capable of manifesting as psychic image or as physical arrangement. An archetype on this account is closer to a symmetry principle or a mathematical structure than to a mental image. It is not in the mind. The mind is one of the places it shows up.

This is a considerably more interesting claim than it is usually given credit for. The parallel with physics should not be overplayed: Jung’s archetype is not the amplituhedron. The structural rhyme is that both try to name a level where a distinction we take as basic has not yet been drawn, from which branches we treat as separate are projections. Rhyme is not identity.



A published diagram of the same idea is useful for readers who know the Jung literature under that name. Again: history of ideas, not a result.



21. Synchronicity: the weakest link, discarded

Jung published “Synchronicity: An Acausal Connecting Principle” in 1952, in a joint volume with Pauli’s essay on archetypal influences in Kepler. The joint publication is itself the point: this was intended as a collaboration between depth psychology and physics. Jung’s definition is a “meaningful coincidence of two or more events, where something other than the probability of chance is involved.” He proposed synchronicity as a fourth explanatory principle alongside space, time, and causality.

The famous scarab anecdote need not be retold at length. The case against is strong and I want it in full force. The reference-class problem: to claim a coincidence is improbable you must specify the class of events and the space of alternatives; neither is well-defined. Littlewood’s Law: rare events are common because there are so many opportunities. Apophenia and confirmation bias: human pattern-detection is tuned for high sensitivity and pays for it in false positives; we do not record the non-coincidences. Meaning is observer-imposed: Jung’s criterion is meaningful coincidence, and meaning is a relation between the pair and an interpreting mind, which makes the criterion subjective and the phenomenon effectively unfalsifiable. Jung’s own inconsistency: he describes it as “acausal,” then speaks as though the archetype does ordering work, which is causal talk in another vocabulary; he vacillated on whether synchronicity was rare or continuous.


My verdict: synchronicity as an empirical claim about anomalous statistical correlation is not established, is probably not testable in its current formulation, and should be dropped from any serious version of this argument. The Pauli Room 137 story is a perfect example of why remarkable anecdotes survive retelling while ten thousand unremarkable deaths generate none. I want nothing from that for the thesis.

22. Pauli: what a Nobel laureate was actually doing there

Wolfgang Pauli formulated the exclusion principle, predicted the neutrino, was among the sharpest critics in twentieth-century physics — the origin of the phrase “not even wrong” — and spent twenty-six years in correspondence with Jung, from 1932 until his death in 1958. The letters are collected in *Atom and Archetype* (Meier 2001). The relationship began therapeutically; it did not stay therapeutic. Pauli became an intellectual collaborator.

Pauli’s concern was that physics had achieved its objectivity by excising the observer, and that quantum mechanics had shown this excision to be incomplete — that the “detached observer” of classical science was no longer available. He sought what he called a neutral language: a mode of description prior to the split between physical and psychical, in which both could be expressed. His interest in Kepler was diagnostic rather than antiquarian: he wanted to show that even the founding moments of modern physics had been guided by archetypal, pre-rational intuitions about form and symmetry, and that the pretence of a purely empirical science was historically false.

This is a serious epistemological position. Pauli was not doing mysticism on the side. He thought the incompleteness of physics as physics required it. That does not make Jung’s scarab a data point. It does make the dual-aspect project historically non-frivolous.

	
Figure 17b. Wolfgang Pauli Physicist who spent decades writing with Jung. Face for history — not evidence.	Figure 17d. C. G. Jung Unus mundus / dual-aspect lineage. History only — not scarab stories as data.
Source: Wikimedia Commons, File:Wolfgang Pauli young.jpg (PD).	Source: Wikimedia Commons, File:Jung 1910.jpg (PD).

23. Atmanspacher’s reconstruction and OSR as ally

The most important development for the credibility of this strand is the reconstruction of the Pauli–Jung conjecture by Harald Atmanspacher, with Hans Primas and later Dean Rickles (Atmanspacher and Primas 1996; Atmanspacher 2012; Atmanspacher and Rickles 2022).

The reconstruction works like this. Distinguish the ontic level — how things are, independent of any observation — from the epistemic level — how things are known, relative to a means of description. Dual-aspect monism in the Pauli–Jung form says that the ontic level is psychophysically neutral: it is a single domain that is neither mental nor physical. The mental and the physical are two epistemic aspects, arising through a distinction drawn by an observer, in something like the way a single quantum system admits complementary descriptions that cannot be jointly instantiated.

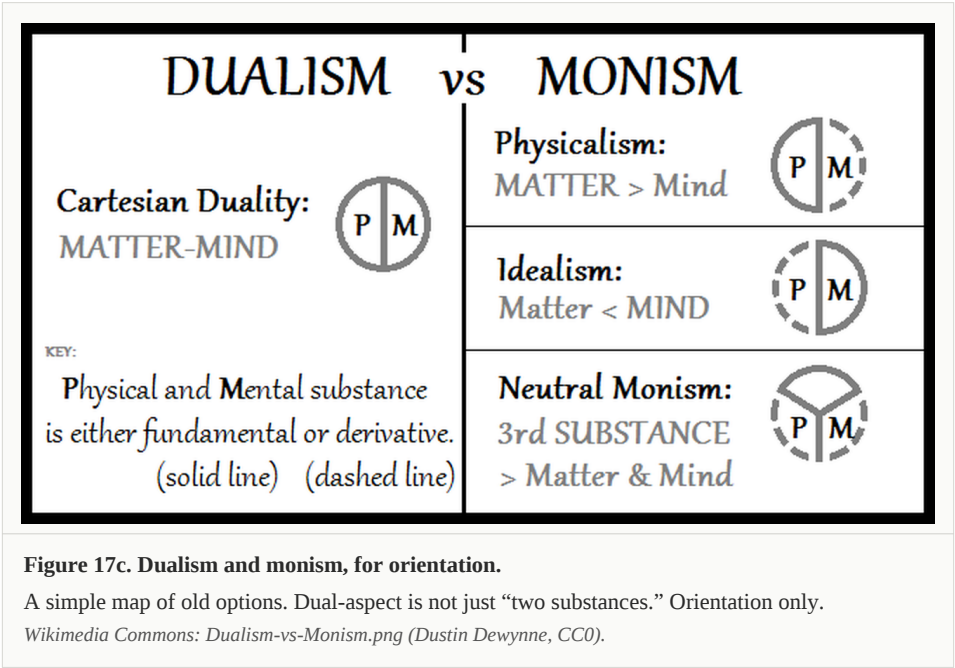
This is not Spinoza’s version, where the two attributes are parallel and complete. It is not Russell’s neutral monism, where the neutral stuff is sensory. It is not Chalmers’ naturalistic dualism, where phenomenal properties are additional fundamental properties. Atmanspacher’s formulation makes the distinctive claim precise: the mind–matter distinction is epistemic, drawn at the level of description, not found at the level of being — and correlations that appear “acausal” from within either aspect are what you would expect if both derive from a common unsplit ground.

It is still under-specified positively. “Neither mental nor physical” is a constraint, not a theory. My own desiderata demand a positive characterization of the ground. Dual-aspect currently fails that test as badly as many rivals. Claims of weak falsifiability via non-reproducible mind–matter correlations risk looking like unfalsifiability dressed as prediction; I am not optimistic about that research program.

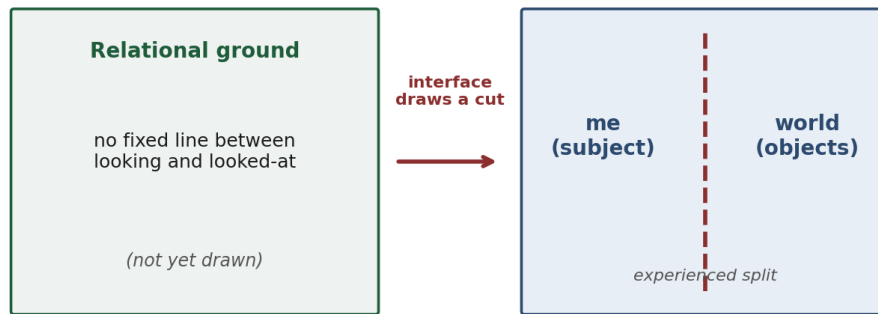
For a cleaner philosophical ally on relational primacy without Jung’s apparatus, see ontic structural realism (Ladyman and Ross 2007, Every Thing Must Go): relations without fundamental relata, metaphysics naturalized from physics. That is the ground-language I prefer when forced to name one. Position the

Compression Hypothesis as: OSR-style structure about the ground, plus interface theory about the format of experience. Dual-aspect becomes optional historical rhyme rather than load-bearing pillar.

First a crude map of substance options — dualism is not the only alternative to “matter only” — then the dual-aspect sketch this strand actually needs.



The perceiver / perceived split is part of the encoding



Useful for action. Not required as the final structure of the ground.

Figure 13. The me / world split.

The cut between “me” and “world” is useful for acting. It need not be the last word about what is underneath.
Drawn for this paper.

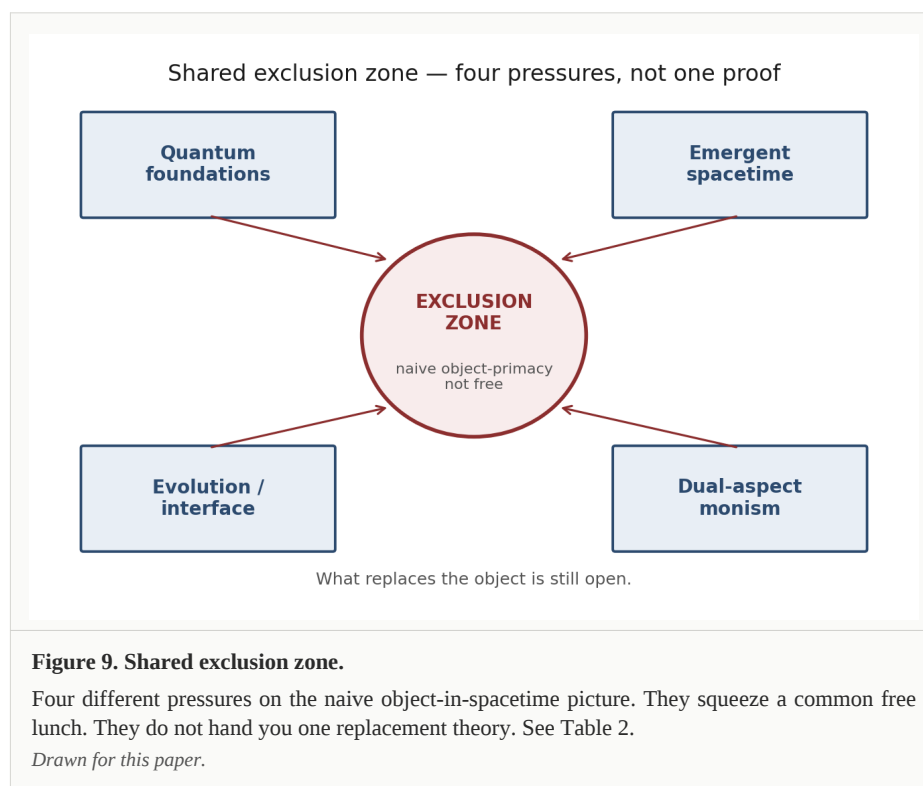
Part IV — Exclusion Zone, Divergence, and the Compression Hypothesis

24. Four negative pressures, not one thesis

Set the strands side by side. Draft 1’s consilience rhetoric sometimes implied identity of negative results. That was the main intellectual threat to the paper. Here is the careful version.

Strand	Method	What it presses against
Quantum foundations	No-go theorems + experiment	Non-contextual, local, absolute object-properties as free lunch
Emergent spacetime	Holography, AdS/CFT, entanglement geometry	Spacetime as unquestioned fundamental stage
Evolutionary epistemology	Game theory / FBT (deflated)	Guaranteed structure-preserving perception
Dual-aspect monism	Conceptual reconstruction	Mental/physical as ontic split

Table 1. Shared exclusion zone — pressures, not identical theorems.



Independence of methods is real. Bell was not thinking about perception. Maldacena was not thinking about Jung. Hoffman’s theorem needs no quantum mechanics. Pauli and Jung were writing in the 1950s about a problem the other three had not posed. There is limited shared premise and limited shared literature. That makes shared pressure interesting. It does not make four proofs of one Compression Hypothesis.

Kochen–Specker contextuality is not “spacetime is a desktop.” Emergent geometry in AdS is not Fitness-Beats-Truth. Dual-aspect is not a no-go theorem. The honest consilience is narrower: naive object-primacy is under pressure from several directions. That is still worth saying. It is not the same as four independent demonstrations of one positive thesis.

25. What each strand still allows

Exclusion without “still allows” is incomplete. Here is the divergence that matters more than the convergence.

Strand	Still allows (examples)	Chief remaining liability
QM no-gos	Everett branching objects; Bohm particles; QBist agents; RQM relative facts	Does not pick interface ontology
Emergent spacetime	Fundamental quantum info without any perception theory	AdS $\neq$ our universe; toy models
FBT / ITP	Approximate realism if fitness is monotone; ecological realism	Self-defeat; generic-fitness assumption
Dual-aspect	Idealist or physicalist completions; under-specified ground	Positive content thin; weak tests

Table 2. What the exclusion zone does not settle.

Every entry in a “what is fundamental” column across the field is different. The convergence is on the destruction of free object-primacy; the reconstruction is wide open. This is the actual state of the field, and any paper that presents a unified answer as forced is lying. The Compression Hypothesis is one conjecture about the shape of the replacement — not the unique implication of the zone.

Consciousness science specifically is immature. IIT 4.0 (Albantakis et al. 2023) was hit the same year by an open letter signed by many researchers characterizing it as pseudoscience — a charge others rejected as conflating “probably false” with “unscientific.” The Cogitate adversarial collaboration pitting IIT against Global Neuronal Workspace published in Nature in 2025 and delivered a genuinely mixed verdict: challenges to both, decisive vindication of neither. That is what a healthy but immature science looks like. Nobody should be building an ontology on it, including me.

26. The Compression Hypothesis, stated

With the ledger clear — and with Part I-M’s reminder that formal results constrain rather than uniquely force the full package — here is the positive proposal.

The experienced world — extended in space, ordered in time, composed of separable objects bearing definite properties — is a lossy, action-oriented encoding generated at the interface between a bounded system and its environment. Space, time, and objecthood are properties of the encoding format. What is encoded is a relational structure that does not have object-form, in which the distinction between the perceiving and the perceived has not been drawn.

Stated as a lump, that sentence is easy to wave at and hard to attack, which helps no one. A hostile reader is owed the seams. The hypothesis is a conjunction, and its conjuncts are not equally load-bearing or equally secure. Here it is taken apart, each piece marked for how much weight I am actually putting on it — core (the thesis dies without it), optional (a natural extension I would give up before the core), or speculative (offered, not leaned on).

Conjunct of the Compression Hypothesis	Status
Experience is a lossy, action-oriented encoding produced by a bounded system	Core
The encoding is systematically related to its ground, not a fiction (compression $\neq$ illusion)	Core
Objecthood is format — “thing with definite properties” is a feature of the encoding	Core

Conjunct of the Compression Hypothesis	Status
The ground is relational, without fundamental object-form (working ontology: ontic structural realism)	Core
Space is part of the format rather than the ground	Core / strong
Time, too, is part of the format (the arrow and the block may both be encoding)	Optional
The subject/object cut is itself drawn by the encoding, not found in the ground	Optional
Dual-aspect monism as a gloss on “mental” and “physical” as two encodings of one ground	Optional / historical
Any reading on which a drug, an anomaly, or intention exits or rewrites the format	Rejected — not part of the thesis

Table 3. The hypothesis as a labelled stack. Attack it a conjunct at a time. The core rows are where the fight actually is; the optional rows I will trade before the core; the last row is disowned.

One consequence of the table is worth stating outright, because it answers the fair charge that the ground is a “playlist” of gestured-at options. When forced to name the ground positively — and desideratum 3 says I must — I name one: ontic structural realism, the view that relations are ontologically prior to relata, that at bottom there are structures and no things (Ladyman and Ross 2007). The Compression Hypothesis is then a two-part package: **structural realism about the ground, plus interface/compression about the format**. Dual-aspect monism is kept as historical rhyme in the optional rows, not as the preferred name of the ground. That is a real commitment, and it is falsifiable-adjacent: if the best physics of the ground turns out to be irreducibly about individual objects with intrinsic properties after all, this package is wrong at the core, not merely underdetermined.

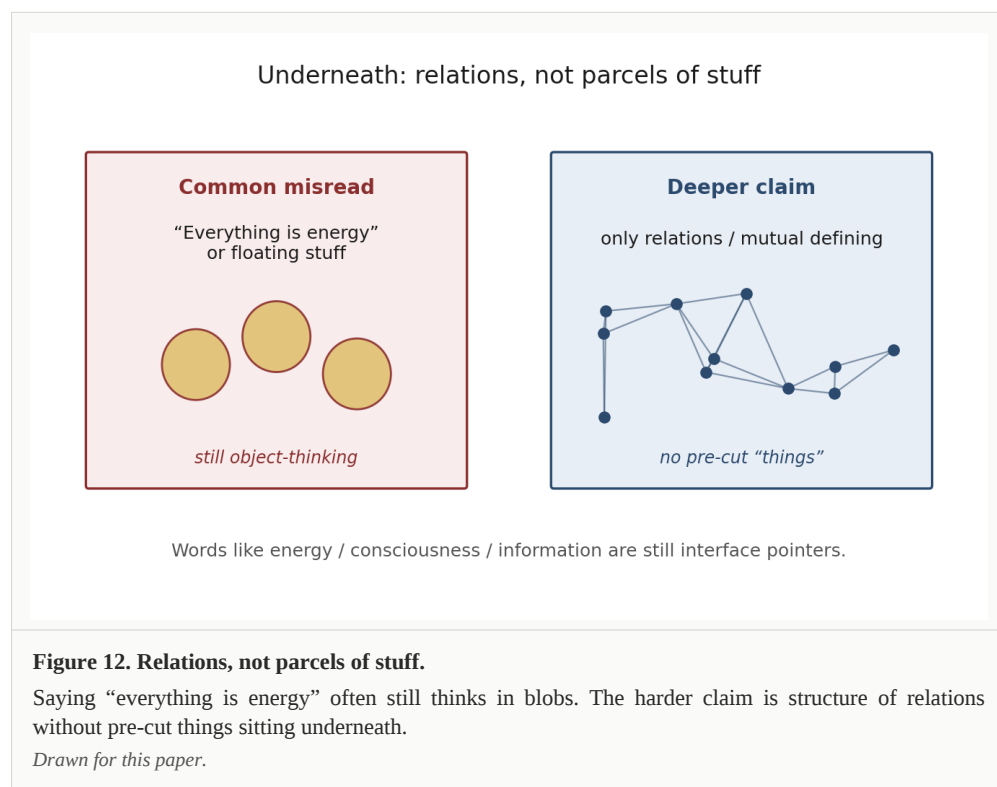
Six clarifications, each of which is load-bearing:

- **(i) A compression is not an illusion.** Lossy compression is a systematic, mathematically characterizable relation. The encoded signal constrains the encoding; the encoding supports reliable inference back to fitness-relevant features of the signal. Realism about the world’s constraining power is fully preserved. What is denied is that the format of the encoding is the format of the world.
- **(ii) It is not anthropocentric.** The interface is generated by bounded systems, not by human beings, and not necessarily by conscious ones. A thermostat has an interface in a thin sense. A bacterium has an interface. The relevant boundedness is informational — a finite system cannot represent an environment of greater complexity than itself, and must therefore compress. This connects directly to Friston’s Markov-blanket formalism and to Rovelli’s relational reading, in which any physical system can play the observer role. Consciousness is not required for the interface. What consciousness may be required for is experience of the interface, which is a different claim.
- **(iii) Compression is under an action constraint, which fixes the format.** This is where Hoffman’s argument does its real work and where the hypothesis gets its specificity. Encodings are not arbitrary; they are optimized. What they are optimized for is the maintenance of the system, which means the format will be organized around affordances — around what can be done — rather than around structure-preservation. Objecthood is what a manipulable, avoidable, edible chunk of the world looks like when encoded by something that manipulates, avoids, and eats. This makes a testable-in-principle prediction: interface structure should vary with action repertoire, and comparative work has documented related effects for decades. Artificial agents now allow experimental manipulation of action spaces (Part VI).
- **(iv) Why the deep structure is hard to see.** You cannot inspect the format from inside the format using the format’s own primitives. Every instrument we build is an object in space; every reading we

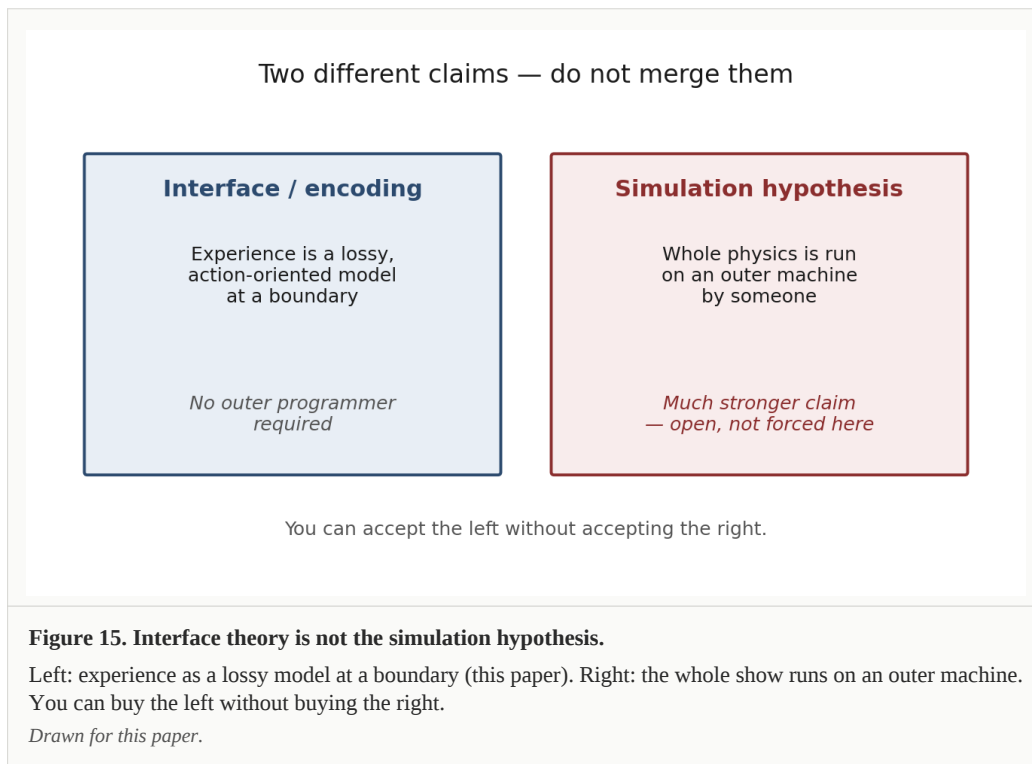
take is perceived through the interface. This is not a convenient excuse — it is a structural prediction, and it explains the otherwise odd fact that the routes we have found to the deep structure are formal ones: mathematics, which is not perceptual, and the no-go theorems, which work by deriving contradictions rather than by looking. It is also not a free pass to unfalsifiability; Part VI exists because a theory that cannot be surprised is not working.

- **(v) How it sits with the four strands without straining.** Contextuality is what you expect when properties are generated at the encoding interface rather than read off. Emergent spacetime is what you expect when the format is spatial and the content is not — remembering that the bridge is synthesis, not theorem. FBT, deflated, is the argument that the format is fitness-tuned. Dual-aspect monism is the claim that “mental” and “physical” are two encodings of one unsplit ground — on my version, “physical” is the encoding of the ground under the action constraint, and “mental” is the encoding under the constraint of self-modelling. Ontic structural realism names the ground as structure without fundamental things.

- **(vi) It does not solve the hard problem, and I will not pretend it does.** Why there is something it is like to run an interface is exactly as mysterious after this hypothesis as before it. The Compression Hypothesis is a thesis about the structure of the objective world and the format of experience, not an explanation of phenomenal consciousness. Conscious realism attempts the latter by axiom; I regard that as relocating the problem rather than solving it, and I do not want to make the same move more quietly.



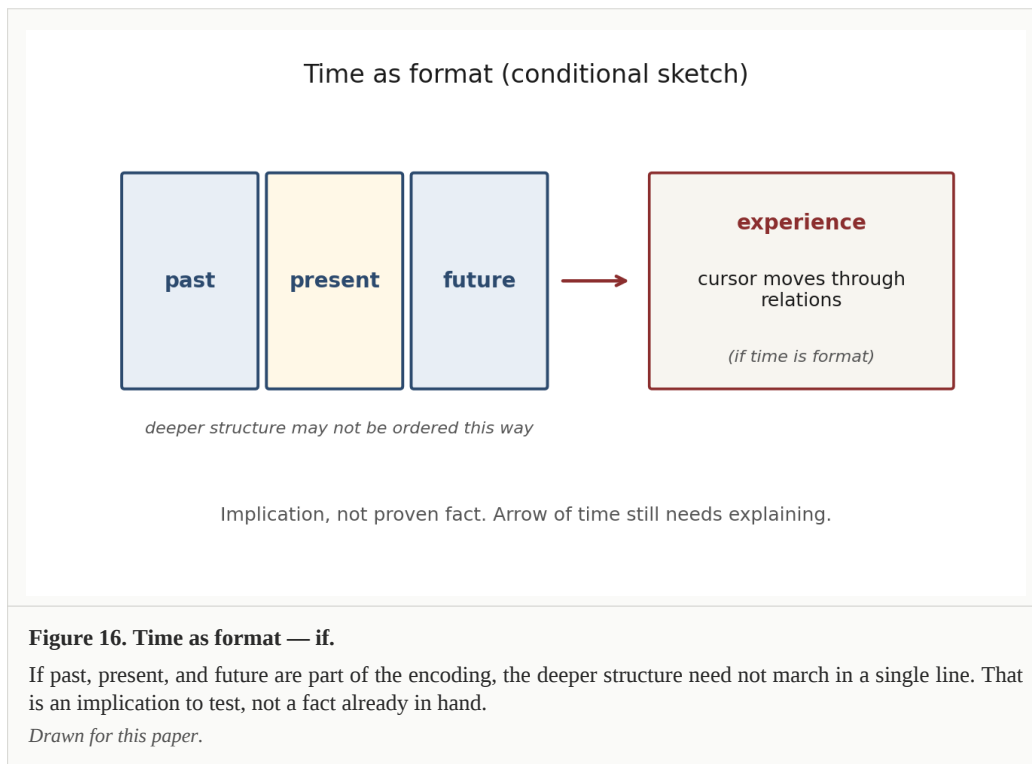
A common mix-up is worth blocking now, while the thesis is still on the page.




27. Time as format

Inside the ordinary encoding, time appears to flow. There is a past that is gone, a present that is happening, and a future that has not yet arrived. That sequential feeling is extremely useful for action: you can plan, remember, and react.

If time itself is part of the encoding format, then the deeper relational structure does not have to be ordered the same way. Past, present, and future may not be three different places that things move through. They may be different aspects of the same structure, the way different pages of a book all exist even though you read them in sequence. Physics already has related ideas — block-universe pictures in which past, present, and future exist equally in a four-dimensional spacetime; thermodynamic accounts of the arrow of time as special initial conditions and entropy increase. The interface view goes one step further by suggesting that even a spacetime block may still be an encoding rather than the final structure.



Important limits. This is a natural implication of treating time as format. It is not a proven fact about the deeper reality. We have no direct access to the deeper level. The asymmetry of time and the stability of the past remain desiderata any theory must meet (Part V): the render does not flicker; memory and record agree. Infinite branching timelines are not required by the Compression Hypothesis and should not be smuggled in from many-worlds or science fiction. Relativity’s light-cone picture is about causal contact in spacetime geometry — useful background, not a slogan that “time is only a desktop.”



28. Intersubjective agreement — a mini-model

If facts are interface-relative, why do observers agree so thoroughly? This is the classic killer objection to interface and QBist-adjacent views. A theory that makes agreement mysterious has traded one problem for a worse one.

Sketch, not solution. Bounded systems that maintain themselves in a shared environment under overlapping action constraints will evolve or learn encodings that are highly correlated, because encodings that fail to track the same fitness-relevant regularities are eliminated. When two interfaces couple — conversation, joint attention, shared instruments — consistency conditions appear as further interactions, not as a view from nowhere (this is close to how RQM and some QBist discussions handle agreement). Scientific consensus is then a further, socially and instrumentally amplified coupling of interfaces under shared experimental constraints.

This does not require that any interface be isomorphic to the ground. It requires that interfaces be systematically related to the same constraints and to each other. Agreement is explained by shared environment, shared action repertoire of the species, and mutual coupling — not by transparent access to objects-in-themselves. The sketch is incomplete; desideratum 2 remains open. But the problem is not automatically fatal once compression rather than fiction is the frame.

Part V — Desiderata and Objections

29. Six desiderata for any theory in this space

If the foregoing is right, then the reconstruction problem — what replaces the object — is the live question. Here are the constraints I think any candidate must satisfy. I offer them as the most useful concrete output of this paper.

- **1. Recover the Born rule non-circularly.** Any observer-involving ontology must derive quantum probabilities without assuming them. Everett has struggled with this for decades; QBism declares it normative rather than derived; conscious realism has not seriously attempted it. This is the sharpest technical filter available.
- **2. Explain intersubjective agreement.** If facts are interface-relative, why do observers agree so thoroughly? Section 28 is a sketch, not a completion. Consciousness-first views generally have not done this work.
- **3. Specify the ground positively, not by negation.** “Neither mental nor physical” is a constraint, not a theory. The ground must be characterized by what it is — a combinatorial structure, an information-theoretic object, a space of agents with specified dynamics, the modal structure of physical theory — precisely enough to be wrong. OSR is a step; dual-aspect alone is not enough.
- **4. Preserve the asymmetry of time and the stability of the past.** Interfaces are dynamic; the world is stubbornly consistent. Any theory must explain why the render does not flicker, and why memory and record agree.
- **5. Make at least one novel prediction that could fail.** This is the requirement most of the field fails, and the one I take most seriously against my own position. Part VI is my attempt.
- **6. Explain why the interface is unified.** Whatever generates the render produces a single, integrated field of experience rather than a heap of independent channels. This is the binding problem, and it is where IIT — whatever its status — is asking the right question.

30. Objections

“This is a category error dressed as a synthesis.” The objection: Bell’s theorem is mathematics; Jung’s archetypes are literary criticism. Placing them in the same argument commits the equivocation that “reality is not what it seems” means the same thing in each — when in physics it means “observables lack non-contextual value assignments” and in Jung it means something closer to “the world is meaningful.”

The response: I accept the danger and I have structured the paper against it. Strands are weighted explicitly; §10 states exactly what physics establishes; synchronicity is discarded; dual-aspect is demoted; exclusion zone language replaces false consilience. But the objection over-reaches if it claims no claim can span both domains. Atmanspacher’s dual-aspect monism is a philosophical thesis about the ontic/epistemic distinction, stated precisely enough to be criticized, and its relation to relational quantum mechanics is structural rather than metaphorical: both say that the description is relative to a cut, and neither claims the cut is drawn by a mind in the mystical sense. The load-bearing claim in strand 4 is not an empirical claim about coincidences. If the objection is that Jung’s evidence is worthless, I largely agree. If it is that his structural proposal is meaningless, that is too strong.

“The physicists who do this for a living do not agree with you.” The objection: most working physicists regard interpretive questions as unimportant, and of those who care, many are Everettians or instrumentalists. The consciousness-adjacent readings are a minority position held disproportionately by

non-physicists.

The response: This is true and it is the strongest sociological evidence against me. I offer three qualifications, none of which fully answers it. First, the negative results in §10 are not minority positions; they are consensus. Second, the modal position among working physicists is not “the object is fine” but “I do not think about it” — survey data on interpretation preferences show a field that is genuinely split and largely disengaged (Schlosshauer et al. 2013 and later polls), which is weaker evidence against than active rejection would be. Third, the emergent-spacetime program is mainstream, well-funded, and staffed by people with no interest in consciousness, and it independently concludes that spacetime is derivative in its models. But I will not pretend the sociology favours the full Compression Hypothesis. It does not, and a reader who weights expert consensus heavily should discount this paper substantially.

“The argument undermines itself.” The objection: if perception is a compression and the interface hides the world, then the neuroscience, evolutionary biology, and physics used to establish this were all conducted through the interface. The conclusion destroys its premises.

The response: This is the deepest objection and I do not think it can be fully answered. Three partial responses. First, the compression framing is significantly less self-destructive than Hoffman’s stronger anti-veridicality claim. A compression preserves fitness-relevant structure; science is a fitness-relevant enterprise conducted on the compressed data; there is no contradiction in a compression being reliable about the features it encodes, including features of its own operation. Second, and more importantly: the argument is largely negative, and negative arguments are more robust to this objection. The no-go theorems do not require that perception be veridical. They require only that the mathematics be valid and that experimental records be reliable relative to one another — a consistency requirement, not a correspondence requirement. Deriving a contradiction from a set of assumptions does not presuppose that any assumption is true of the noumenon. Third, I accept a genuine residual cost: the positive thesis is more vulnerable than the negative one, and the conclusion “the deep structure is relational and non-spatial” is a claim I reach using cognitive tools whose reliability I have just questioned. Bagwell’s self-defeat critique is the best recent development of this problem for strong interface skepticism. I flag this as unresolved rather than resolved.

“It makes no novel predictions.” The objection: the Compression Hypothesis predicts everything already observed and forbids nothing. It is philosophy dressed in the vocabulary of science.

The response: Largely conceded at present, with a qualification and a plan. The qualification: the hypothesis is a framework claim, at the same level as “the world is made of matter in motion,” and framework claims are not usually falsified directly but by the progressive success or sterility of the programs they sponsor. The plan is Part VI, which attempts to say what would count against it, including two near-term methods sketches. A theory that cannot be surprised is not doing work; I take that seriously enough to write the next part.

“You want this to be true.” The objection: an ontology in which consciousness is not an accident and the world is not indifferent is consoling. Motivated reasoning is the null hypothesis for anyone who arrives at it.

The response: Correct as a warning and irrelevant as an argument. Motivation explains why someone holds a view; it does not evaluate the view. But it is a real reason for elevated scrutiny, and it is why I have made a point of stating the strongest version of every objection, discarding synchronicity, refusing the collapse-by-consciousness reading, demoting dual-aspect, moving high-strangeness to an appendix, and conceding the sociology and prediction problems. I would add that the framework is not obviously

consoling. An ontology in which the self is an artifact of a compression scheme, in which facts are perspectival, and in which the objects one loves are format rather than content, is not straightforwardly comforting. It is at least as vertiginous as materialism.

Part VI — What Would Discriminate

A framework that cannot be surprised is not doing work. Here, seriously offered, are lines along which the positions could be pulled apart. None is a decisive experiment for the whole package; several are within reach.

1. Extended Wigner’s friend with a macroscopic, arguably-cognitive friend. The existing experiments use photons as friends, which is why they establish less than the headlines claim. Scaling the “friend” up — to a system with genuine information-processing and memory, and eventually to something we would be prepared to call an observer — is a hard but well-defined experimental program, actively pursued. If Local Friendliness violations persist as the friend becomes cognitively non-trivial, the absoluteness of facts is in serious trouble. If they vanish at some scale, that scale is itself a discovery of the first order.

2. Objective-collapse parameter space. Continued underground and low-temperature searches for spontaneous-collapse radiation constrain GRW/CSL and gravitational collapse models. This is the cleanest falsification frontier in foundations: it can rule out the “collapse is physical, no observer needed” family more tightly, which would substantially raise the standing of observer-involving views by elimination — without proving any particular interface theory.

3. De Sitter holography. If the holographic correspondence can be established for a universe like ours rather than only for anti-de Sitter space, emergent spacetime moves from “coherent in toy models” to “true of the world.” This is among the most consequential open problems in theoretical physics and progress on it directly changes the weight of strand 2.

31. Methods sketch A — testing FBT’s genericity assumption

The toy model in Part II (§13b, Fig. 11d) is the analytic skeleton of this sketch: it shows the FBT advantage as a function of a single monotonicity parameter and marks where that advantage dies. What the model stipulates — the distribution of real fitness functions along that axis — the experiment would measure. This is the most tractable high-value item for strand 3 and, to my knowledge, still underdone. Rather than sampling fitness functions generically, measure or model them: characterize actual fitness landscapes in tractable organisms or digital evolution systems with physics-like worlds; determine what fraction are approximately monotone in latent world variables; re-run evolutionary games with the empirical distribution. Mattes (2026) is a step in this direction — his coevolving-fitness simulations already push against the generic assumption — and the sketch is the systematic version of that move.

Protocol outline. (i) Choose a domain where fitness can be defined operationally (foraging models, toxicity landscapes, simple multi-resource environments). (ii) Estimate or simulate fitness as a function of latent state variables agents might in principle track. (iii) Classify landscapes as monotone / approximately monotone / rugged with respect to those variables. (iv) Run competitive evolutionary dynamics comparing structure-preserving strategies to fitness-tuned interface strategies under the empirical measure, not a uniform measure over all functions. (v) Pre-register the classification criteria and the success metric.

Interpretation. If real fitness functions are predominantly monotone, FBT is a theorem about a thin corner of nature and strand 3 substantially weakens. If they are not, strand 3 strengthens considerably. Either result is progress. This is computational and theoretical biology more than philosophy.

32. Methods sketch B — embodied AI action repertoires

The Compression Hypothesis predicts that perceptual format tracks action repertoire rather than environmental structure alone. Comparative work on species with different effectors is one path. Embodied artificial agents are another: for the first time, a system where the “interface” can be inspected from outside while the action constraint is varied experimentally.

Protocol outline. (i) Shared environment with fixed observation bandwidth. (ii) Agents trained with systematically varied action spaces (e.g. locomotion only vs locomotion plus manipulation vs no embodiment). (iii) Inspect learned world-models / latent representations for object-like, space-like structure. (iv) Test whether that structure tracks action repertoire more than raw environmental geometry when the two are experimentally dissociated. (v) Controls: matched reward complexity; ablations on embodiment; multiple seeds and architectures.

Interpretation. Success supports action-shaped encoding — almost ecological psychology with better instruments. Alone it does not prove spacetime is non-fundamental in physics. It does make the positive hypothesis surpriseable in a domain that did not exist when the desktop metaphor was coined. I regard this as the most promising near-term line for the Compression Hypothesis specifically.

6. Derive the Born rule from an interface postulate. The theoretical prize. Any observer-involving framework that produces the Born measure from its own primitives, without assuming it, would move from philosophy toward physics in a single step. QBism’s reconstruction program — deriving quantum theory from informational axioms — is the most serious attempt in this direction.

7. Predictions from dual-aspect monism. Atmanspacher and colleagues have argued that mind–matter correlations of the kind their framework predicts should be non-reproducible on demand in a characterizable way — decaying under repeated measurement rather than stabilizing. Whether that is a genuine prediction or an unfalsifiability device dressed as one is exactly the sort of thing that could be settled by pre-registered work, and has not been. I list it for completeness and remain sceptical.

Conclusion

The claim of this paper is not that consciousness creates reality. It is not that quantum mechanics vindicates ancient wisdom. It is not that Jung was right about coincidences or that Vallée was right about anything. It is not that mushrooms open a door out of the desktop.

The claim is that several investigations, which had limited reason to agree, put shared pressure on a picture: the object in spacetime as ground floor. Quantum foundations establish constraints with theorems. Quantum gravity offers explicit models in which geometry is reconstructed from entanglement. Evolutionary epistemology, carefully read, argues that perception is not selected for isomorphism. Dual-aspect monism, carefully reconstructed, treats the mental/physical cut as epistemic. Ontic structural realism says the same quiet thing from philosophy of physics: structure first. Predictive processing supplies the scientific language of lossy, action-oriented encoding without settling the strong metaphysical claim.

What replaces the object is unknown. The candidates — the universal wave function, relational states, agents' expectations, entanglement structure, conscious agents, a psychophysically neutral ground, a rule-space — are mutually inconsistent, and most of them are underdeveloped. The Compression Hypothesis I have proposed is a way of stating what they share: that the world we experience is a bounded system's action-oriented encoding of a structure that does not have the encoding's form. It is a conjecture. I have tried to say precisely what would count against it, and to concede, where honest, that it currently predicts too little.

Two closing observations. The first is that the shape of the answer has become surprisingly stable across all this disagreement. Whatever is fundamental, it appears to be relational and structural rather than substantial. Every serious candidate — entanglement, information, agents' interactions, the ontic structural realists' relations-without-relata, Jung's psychoid ordering principles stripped of numerology, Wheeler's bit — describes a world of relations rather than of things bearing properties. The dissolution of the object is not being replaced by a smaller object. It is being replaced by structure.

The second is that the question is no longer disreputable. For most of the twentieth century, asking what reality is rather than how it behaves marked you as insufficiently serious. That has changed, for a reason that has nothing to do with fashion: the technical results forced it. You cannot prove Kochen–Specker and then decline to have a view about whether properties are intrinsic. You cannot reconstruct a bulk geometry from boundary entanglement and then insist that space is basic without further argument. The metaphysical question was reopened by physics, not by philosophy alone, and it is now a live empirical enterprise with a growing list of things it could get wrong.

That is the most one can honestly say. It is also, I think, considerably more than most people realize.

A last practical note, for the reader who has come this far. If you take only three things from this paper, take these. First: the naive object is under pressure from more than one direction, and that pressure is real even if the directions do not prove the same positive thesis. Second: compression is not illusion — science can be objective about constraints without requiring that the format of experience be the format of the ground. Third: refuse the package deals. Consciousness-causes-collapse, synchronicity-as-fact, high-strangeness-as-proof, and psychedelic-exit-from-the-desktop are not required by the interesting thesis and repeatedly destroy its credibility when attached to it. The interesting thesis is hard enough without them.

Appendix A — High Strangeness as Stress Test (Vallée)

This appendix is the one most likely to cost the paper its credibility, and I want to be exact about its logical role.

I am not offering anomalous phenomena as evidence for the Compression Hypothesis. The evidential quality is far too poor, and a thesis supported by Parts I–IV would be weakened, not strengthened, by attaching an unresolved empirical controversy to it.

The section exists for two other reasons. First, the interface framework makes an unusual and reasonably sharp prediction about what anomalous experience should look like if the framework is true, and it is worth seeing whether that prediction is satisfied — a framework that predicts nothing about anything is not doing work. Second, this domain is where the epistemic hazards of the entire enterprise are most visible, and examining them here inoculates the rest of the argument.

A.1 Vallée’s position

Jacques Vallée is not a hobbyist. He holds a master’s degree in astrophysics, a PhD in computer science from Northwestern, worked on the Mars map program at NASA, was a principal investigator on ARPANET-era network projects at SRI International and Stanford, and has had a long subsequent career in venture capital. He is a scientist who took an unpopular problem seriously, which is a different thing from a crank.

The folklore continuity thesis. In *Passport to Magonia* (1969), Vallée compared modern UFO-encounter reports with the historical record of fairy encounters, Marian apparitions, incubi, and the folkloric abduction literature. The structural parallels are extensive: paralysis, missing time, luminous beings, abduction to another realm, disorientation on return, physiological aftereffects, the recurring theme of a substance offered and refused. His conclusion was that the modern reports are not new; they are the current cultural dress of a very old and recurring pattern of experience.

The control system hypothesis. In *The Invisible College* (1975) and *Messengers of Deception* (1979), Vallée pushed further. He argued that the phenomenon behaves less like a visitation and more like a control system — in the cybernetic sense, a thermostat — operating on human belief over long timescales. The observation driving this: the phenomenon’s manifestations are absurd, and the absurdity is systematic rather than incidental. Encounters routinely violate physics, and routinely involve behaviour that makes no sense as reconnaissance, contact, or invasion. Vallée’s move is to take the absurdity as functional. A signal designed to be undeniable to the experiencer and unprovable to everyone else would produce exactly this pattern: a permanent, self-sustaining reservoir of belief that can neither be dismissed nor confirmed, which is precisely the condition that most efficiently modifies a culture’s conception of what is possible.

Rejection of the extraterrestrial hypothesis. Vallée argues that the sheer volume of reported encounters — orders of magnitude beyond what any plausible survey program would require — is inconsistent with nuts-and-bolts interstellar visitation, and more consistent with something informational or “interdimensional.” He has been consistent on this for fifty years, frequently to the irritation of the UFO community itself.

A.2 What the interface framework predicts here

If the Compression Hypothesis is right, then the world of stable objects in spacetime is a rendering produced by our interface, and the rendering is normally seamless because selection has made it so. What would an anomaly look like on this picture?

It would look like a rendering failure. Specifically, it would be expected to have three features: (1) Physically inconsistent but perceptually vivid — not a clean violation of laboratory physics so much as a violation of the format: objects that fail to conserve momentum, change size discontinuously, occupy positions without traversing intervening space. This is what a rendering glitch looks like from inside the render, and it is a persistent feature of high-strangeness reports. (2) Resistant to instrumentation — if the anomaly occurs at the level of the interface rather than in the compressed content, then instruments — which are themselves objects within the render, and whose readings we perceive through the same interface — should give inconsistent or absent confirmation. This is a striking and under-discussed feature of the phenomenon: evidence quality has not improved in decades despite enormous improvement in sensor availability. (3) Structurally stable across cultures, cosmetically variable — the deep grammar of the encounter constant while surface content tracks local cosmology: fairies, then airships, then flying saucers, then greys. This is precisely Vallée’s Magonia finding.

I want to be very clear about the logic here: satisfied predictions of this weakness do not confirm a theory. Every one of these three features is also predicted, and predicted more economically, by the null hypothesis — that anomalous experience is generated by sleep paralysis, hypnagogia, misperception, dissociation, and cultural scripting, all of which produce vivid, uninstrumentable, culturally-dressed experiences with a stable deep structure. The psychological hypothesis and the interface hypothesis make nearly identical predictions, which means the data do not discriminate them. That is the honest statement. The interface framework is not embarrassed by high strangeness, which is a weak virtue, and is not supported by it.

A.3 Official record, factually

For completeness on the empirical state of play as of the mid-2020s: the 2021 ODNI preliminary assessment reviewed 144 reports and offered explanations for few. The All-domain Anomaly Resolution Office was established in 2022. Congressional hearings in 2023 heard testimony from David Grusch alleging a crash-retrieval program, testimony that has not been substantiated by any produced evidence. AARO’s Historical Record Report Volume 1 (March 2024) found no evidence that any US government investigation had verified extraterrestrial technology, and attributed a substantial portion of the mythology to classified conventional programs and to circular reporting. AARO’s subsequent annual reports have continued to catalogue several hundred new cases per year, the large majority resolved or resolvable as balloons, drones, and sensor artifacts, with a residual unresolved fraction. Congressional interest has continued.

The correct summary is that nothing has been established. A residual of unexplained cases is what one expects from any large observational dataset with poor instrumentation and is not evidence of anything in particular. Anyone telling you that disclosure has occurred, or that the residual is significant, is going beyond the record.

A.4 The critique of Vallée, and the general hazard

The decisive objection to the control-system hypothesis is that structurally it is unfalsifiable. A control system that operates by producing evidence that cannot be verified will produce exactly the same data as no control system at all plus ordinary human psychology. Vallée’s framework absorbs confirming and disconfirming evidence with equal ease, which is the classic Popperian pathology. And the folklore continuity that he reads as evidence of an external agency is at least as well explained as a constant in human neurology and narrative — the more parsimonious reading, and one that requires nothing outside the skull.

And this points at the general hazard, which applies to the whole paper. Frameworks of the kind I am defending are unusually good at absorbing anything. If spacetime is a compression, any anomaly is a rendering artifact; any absence of anomaly is a well-functioning render. A theory that cannot be surprised is not doing epistemic work. I take this seriously enough that Part VI is devoted to it — and seriously enough that this material is an appendix, not a pillar.

Appendix B — Chemically Induced Plasticity

Classic serotonergic psychedelics — psilocybin first, LSD second, DMT as the extreme end of the same pharmacological family — alter the experienced structure of space, time, objects, and the self–world boundary. They are sometimes offered as proof that ordinary reality is unreal or that an infinite ground has been contacted. That is a non sequitur. This appendix states what they are good for and what they are not.

B.1 What the science supports

Psilocybin produces large, temporary changes in functional brain connectivity. Siegel et al. (2024, *Nature*) reported massive desynchronization across cortical and subcortical networks relative to a control stimulant (methylphenidate), with default-mode network disruption implicated in the loss of ordinary self–world framing. Earlier work and systematic reviews document reduced within-DMN coherence under classic psychedelics, correlating with reports of ego-dissolution. The brains of people under high-dose psilocybin can look more similar to each other than to their own unaltered baselines — individuality temporarily flattened at the network level.

The leading theoretical model in predictive-processing terms is REBUS — “Relaxed Beliefs Under pSychedelics” (Carhart-Harris and Friston 2019, *Pharmacological Reviews*): precision on high-level priors is relaxed; bottom-up information flows more freely; the hierarchy that normally stabilizes self, narrative, and object constancy is loosened. That is a model of how an encoding’s weights change, not a model of exit from encoding. Critiques of single-pathway REBUS exist (set and setting, multi-network effects, enactive multi-causality); the model remains the best available bridge from pharmacology to the language of this paper.

B.2 What the Compression Hypothesis may say

Classic psychedelics do not open a window onto the unsplit ground. They perturb the compression scheme — lowering precision on high-level priors that normally stabilize self, object, and narrative space — producing alternative encodings that can feel absolute while remaining encodings. That is evidence of interface plasticity, not of exit from the interface.

They are, in that limited sense, cleaner stress tests than folkloric high strangeness: dose-controllable, reversible, measurable with imaging and behavior. They hit constructs this paper cares about. They do not discriminate the Compression Hypothesis from ordinary physicalist neuroscience. Physicalism already expects chemicals to scramble a constructed percept. Plasticity is free under both pictures. Changing icons is not rewriting hardware: chemistry rewrites the desktop for a while; shared regularities and consequences remain.

B.3 What is not supported

- Noetic certainty (“more real than real”) as a truth-tracker. Dreams produce the same certainty. Intensity is a feature of the altered encoding.
- DMT entity encounters as evidence of external beings. Competing explanations include agent-detection biases, cultural scripting, REM-like intrusion, and predictive overfitting under high learning rates. Unestablished as ontology.
- Any bridge from REBUS to AdS/CFT, Kochen–Specker, or dual-aspect monism. Brain priors are not spacetime ontology.

- The idea that “all is energy” or “all is one” has been shown. Those are pointers inside the encoding, not captures of the ground.

Philosophically, Letheby’s naturalized account of psychedelic epistemology (Philosophy of Psychedelics, 2021) is useful discipline: therapeutic or epistemic benefits of certain experiences need not license cosmic metaphysics. I am not adopting Letheby’s full naturalism as a package; I am adopting the separation between phenomenological force and metaphysical license.

Bottom line for the paper: Appendix B is optional color on plasticity. It is not a pillar. Prefer psilocybin citations to DMT culture. Prefer REBUS language to “breakthrough” language. Prefer dreams and anesthesia as co-equal plasticity examples when speaking to sceptical audiences.




Figure credits

Most diagrams in this draft I drew for the manuscript (Python/matplotlib; see `generate_figures.py`). They are teaching pictures, not experimental results and not proof by themselves. Each caption says what you are looking at, what it does not settle, and where it came from.

Open-license images used in this draft (Wikimedia Commons; File: names in captions): JPEG example (JJ Harrison, CC BY-SA 3.0); double-slit light (inductiveload, PD); decoherence diagram (CC BY-SA 3.0); Solvay 1927 (PD); Schrödinger's cat (CC BY-SA 3.0); two-channel Bell test; entanglement swapping (CC0); Bloch sphere (CC BY-SA 3.0); AdS/CFT sketch (Dimension10, CC BY-SA 3.0); computer desktop (AJ Ashton, CC0); light cone (CC BY-SA 4.0); dualism map (Dewynne, CC0); Pauli and Jung portraits (PD); Jung unus mundus diagram (CC BY-SA 4.0). Rival-answers map and scorecard (Figs A–B), free-energy sketch (7b), and Hoffman ITP/FBT drawings are mine. CC BY-SA: credit the author; share-alike if you remix. Do not lift dual-aspect or Hoffman figures from journals without permission.

If you reuse a figure, keep the limit line in the caption, credit this manuscript for the author drawings, and credit Commons for the rest.

References

- Albantakis, L., et al. (2023). Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLoS Computational Biology*, 19(10), e1011465.
- Arkani-Hamed, N., & Trnka, J. (2014). The Amplituhedron. *Journal of High Energy Physics*, 2014(10), 30.
- Atmanspacher, H. (2012). Dual-aspect monism à la Pauli and Jung. *Journal of Consciousness Studies*, 19(9–10), 96–120.
- Atmanspacher, H., & Primas, H. (1996). The hidden side of Wolfgang Pauli. *Journal of Consciousness Studies*, 3(2), 112–126.
- Atmanspacher, H., & Rickles, D. (2022). *Dual-Aspect Monism and the Deep Structure of Meaning*. Routledge.
- Bagwell, J. N. (2023). Debunking interface theory: why Hoffman’s skepticism (really) is self-defeating. *Synthese*, 201, 25.
- Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics*, 1(3), 195–200.
- Bell, J. S. (1987). *Speakable and Unspeakeable in Quantum Mechanics*. Cambridge University Press.
- Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of “hidden” variables, I and II. *Physical Review*, 85, 166–179, 180–193.
- Bohm, D. (1980). *Wholeness and the Implicate Order*. Routledge.
- Bong, K.-W., et al. (2020). A strong no-go theorem on the Wigner’s friend paradox. *Nature Physics*, 16, 1199–1205.
- Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3), 316–344.
- Cavalcanti, E. G. (2021). Implications of Local Friendliness violation for quantum causality. *Entropy*, 23(8), 925.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Chalmers, D. J. (2022). *Reality+: Virtual Worlds and the Problems of Philosophy*. Norton.
- Cogitate Consortium (2025). Adversarial testing of global neuronal workspace and integrated information theories of consciousness. *Nature*, 642, 133–142.
- Donadi, S., et al. (2021). Underground test of gravity-related wave function collapse. *Nature Physics*, 17, 74–78.
- Eddington, A. S. (1928). *The Nature of the Physical World*. Cambridge University Press.
- Everett, H. (1957). “Relative state” formulation of quantum mechanics. *Reviews of Modern Physics*, 29, 454–462.
- Frauchiger, D., & Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications*, 9, 3711.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Fuchs, C. A., Mermin, N. D., & Schack, R. (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics*, 82(8), 749–754.
- Ghirardi, G. C., Rimini, A., & Weber, T. (1986). Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34, 470.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.
- Hensen, B., et al. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526, 682–686.
- Hoffman, D. D. (2019). *The Case Against Reality: Why Evolution Hid the Truth from Our Eyes*. Norton.
- Hoffman, D. D., & Prakash, C. (2014). Objects of consciousness. *Frontiers in Psychology*, 5, 577.
- Hoffman, D. D., Prakash, C., & Prentner, R. (2023). Fusions of consciousness. *Entropy*, 25(1), 129.
- Hoffman, D. D., Singh, M., & Prakash, C. (2015). The interface theory of perception. *Psychonomic Bulletin & Review*, 22, 1480–1506.
- Howard, M., et al. (2014). Contextuality supplies the “magic” for quantum computation. *Nature*, 510, 351–355.
- Jung, C. G. (1952/1955). Synchronicity: An acausal connecting principle. In *The Interpretation of Nature and the Psyche* (with W. Pauli). Routledge.
- Jung, C. G. (1955–56). *Mysterium Coniunctionis*. Collected Works, Vol. 14. Princeton University Press.
- Jung, C. G. (1947/1954). On the nature of the psyche. Collected Works, Vol. 8.
- Kochen, S., & Specker, E. P. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17, 59–87.
- Ladyman, J., & Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.
- Letheby, C. (2021). *Philosophy of Psychedelics*. Oxford University Press.

-
- Maldacena, J. (1998). The large N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics*, 2, 231–252.
- Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes. *Fortschritte der Physik*, 61, 781–811.
- Manning, A. G., et al. (2015). Wheeler’s delayed-choice gedanken experiment with a single atom. *Nature Physics*, 11, 539–542.
- Mattes, J. (2026). *The Case Against The Case Against Reality*. Biological Theory.
- Meier, C. A. (Ed.) (2001). *Atom and Archetype: The Pauli/Jung Letters, 1932–1958*. Princeton University Press.
- Okon, E. (2025). Reassessing the strength of a class of Wigner’s friend no-go theorems. (Discussion of Local Friendliness comparative strength.)
- Prakash, C., Stephens, K. D., Hoffman, D. D., Singh, M., & Fields, C. (2021). Fitness beats truth in the evolution of perception. *Acta Biotheoretica*, 69, 319–341.
- Proietti, M., et al. (2019). Experimental test of local observer independence. *Science Advances*, 5(9), eaaw9832.
- Pusey, M. F., Barrett, J., & Rudolph, T. (2012). On the reality of the quantum state. *Nature Physics*, 8, 475–478.
- Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35, 1637–1678.
- Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters*, 96, 181602.
- Schlosshauer, M., Kofler, J., & Zeilinger, A. (2013). A snapshot of foundational attitudes toward quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 44, 222–230.
- Siegel, J. S., et al. (2024). Psilocybin desynchronizes the human brain. *Nature*, 632, 131–138.
- Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics*, 36, 6377–6396.
- Tegmark, M. (2000). Importance of quantum decoherence in brain processes. *Physical Review E*, 61, 4194–4206.
- ’t Hooft, G. (1993). Dimensional reduction in quantum gravity. *arXiv:gr-qc/9310026*.
- Vallée, J. (1969). *Passport to Magonia: From Folklore to Flying Saucers*. Regnery.
- Vallée, J. (1975). *The Invisible College*. Dutton.
- Vallée, J. (1979). *Messengers of Deception*. And/Or Press.
- Van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. *General Relativity and Gravitation*, 42, 2323–2329.
- Wallace, D. (2012). *The Emergent Multiverse*. Oxford University Press.
- Wheeler, J. A. (1989). Information, physics, quantum: The search for links. In *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics*, Tokyo.
- Wiseman, H. M., Cavalcanti, E. G., & Rieffel, E. G. (2023). A “thoughtful” Local Friendliness no-go theorem. *Quantum*, 7, 1112.
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75, 715–775.