Research Genesis

From Hypothesis to Null Result: the intellectual journey, preserved

Companion document to *Relative Entropy in a de Sitter Causal Diamond* (working paper v1.2)

July 2026

Abstract

This companion preserves what the main paper deliberately compresses: the diagrams, the dead ends, and the order in which attractive ideas were eliminated. The journey has become more interesting precisely *because* so much was discarded — each removal was forced by a definition, an identity, or an order-of-magnitude, never by taste. The evolution itself is the demonstration: not an argument that the conclusions are strong, but a record showing the discipline that produced them. The program began with an internet illustration sourced as a joke and ended, seven artifacts later, with a boxed null result ($\delta\Lambda_{\rm ent} = 0$) and exactly one open door ($Q = ?$). A framework designed so that it can succeed by proving itself wrong keeps its receipts; this document is the receipts.

Preface: why document a journey of eliminations

Papers report what survived. They rarely report the survivorship pressure — and in speculative territory the pressure *is* the result. Between the first poster and the final paper, five load-bearing ideas were killed: white-hole conduits, volume-integrated information, a fitted coupling β, black-hole-accumulation dark energy, and horizon bookkeeping as a cosmological mechanism. None was abandoned because it stopped being attractive; each was executed by something specific — no observable, no invariant definition, a rejected free parameter, a first-order decoupling, a factor of 10^{-18}. What remains is smaller than what began, and more valuable for it: the program can now end only in a no-go theorem or a derived exchange law, and either terminal state is science. The stages below give each artifact, what it added, and what it removed.

1 Stage 1 — The egg (internet illustration)

A meme in the lineage of the declassified 1983 "Gateway Process" analysis: toroidal universe, black-/white-hole conduit at the nucleus, holographic substrate, closed loop of "time-space evolving." No mathematics, no mechanism, no predictions. Sourced as a joke; adopted as a hypothesis generator. *Everything is on the table.*

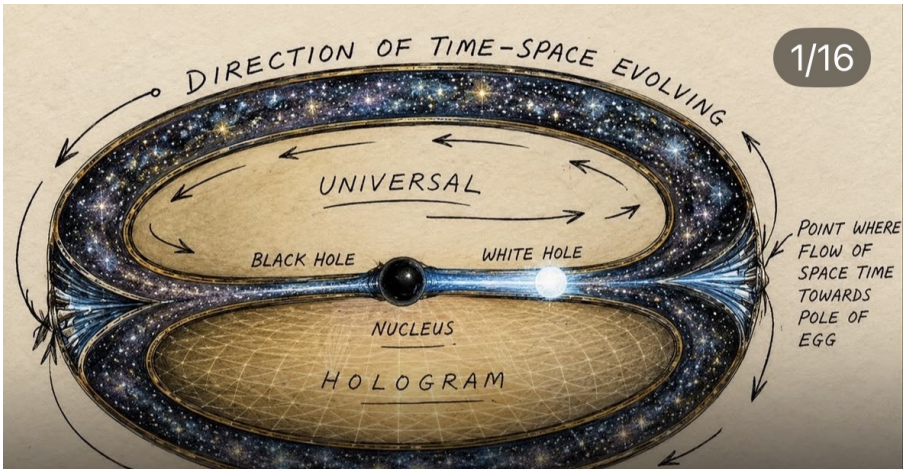

2 Stage 2 — Poster v1: the information cycle, first draft

Introduces the cycle Matter → Information → Geometry → Gravity → Structure → More Matter; black holes as information nodes; entropy accounting; and the emerging dark-energy ansatz $\Lambda_{\text{eff}} = \Lambda_0 + \beta I$.

Removed: white holes / black–white conduits. No mechanism by which they source information or affect Λ; no observables. The first casualty, and the easiest.

3 Stage 3 — Poster v2: refinements

Adds relative entropy and the horizon first law; expands the theoretical foundation (Ryu–Takayanagi, entanglement thermodynamics) and observational connections; begins to question its own volume-based information measure.

Removed: volume information, $I_{\text{universe}} = \int \mathcal{I}(x)\, dV$. No rigorous definition in QFT or gravity; not slicing-invariant; in tension with the framework's own holographic principle. Information, done rigorously, is relational — S_{rel} on horizons — and never integrates to an observer-independent bulk density.

4 Stage 4 — Poster v3: focus and reduction

Removes speculative elements, emphasises horizon-based information, and promotes emergent dark energy to the central open problem, with an explicit six-step research program and a "missing step" box at the frontier.

Removed: β as a free parameter. The coupling must be derived, never fitted; an undetermined β makes the hypothesis unfalsifiable. The headline formula $\Lambda_{\text{eff}} = \Lambda_0 + \beta I$ is struck through pending derivation — the rule that later kills it.

5 Stage 5 — Poster v4: the framework map

The thought-experiment map: theoretical foundation, candidate predictions, observational tests, and the frontier (evolving dark energy) located precisely. The cycle is redrawn as the equations permit.

Removed: the black-hole accumulation model. Summing astrophysical horizon areas gives an energy density orders of magnitude too small, and information in black holes cannot couple to a global Λ at first order. (The Farrah et al. 2023 episode — proposed, tested, disfavored — was kept as the cautionary exhibit.)

6 Stage 6 — Working notes v0.1–v0.2: the technical program

First contact with exact mathematics: Fröb's modular Hamiltonian for dS diamonds, exact coherent-state relative entropy, entropy–area balance, the Jacobson–Visser first law. Attempts to compute β or $\delta\Lambda_{\rm ent}$ directly. Finds the horizon-anchored result $\delta\Lambda/\Lambda = S_{\rm rel}/S_{\rm dS} \sim 10^{-18}$ and, via the Bianchi identity, the unimodular route $\dot{\Lambda} = -8\pi G\, Q$.

Removed: horizon bookkeeping as the end of the story. $S_{\rm rel}/S_{\rm dS} \sim 10^{-18}$ is observationally irrelevant — seventeen orders of magnitude short of anything cosmological. Derived once, retired immediately. Whatever the mechanism is, this is not it.

7 Stage 7 — Paper v1.2: the first calculation

Performs the calculation the notes proposed, with exact modular data. Tests both diagnostic ansätze — $\partial S_{\rm rel}/\partial V_A$ over nested diamonds and $\partial^2 S_{\rm rel}/\partial\epsilon^2$ — and both collapse into matter-sector

terms already present in $\delta\langle T_{\mu\nu}\rangle$. The verdict is overdetermined by three independent structural facts (conformal tracelessness, the Bianchi identity, Λ's integration-constant status):

$$\boxed{\delta\Lambda_{\rm ent} = 0} \quad \text{under the stated assumptions.}$$

The frontier that remains

After seven artifacts and five eliminations, exactly one door is open:

$$\boxed{Q = ?}$$

Can modular dynamics produce a non-zero matter→vacuum exchange rate? If yes: an evolving Λ with a derived, parameter-free mechanism that must then survive DESI, Euclid and Roman. If no: a genuine no-go theorem for information-driven dark energy, in the lineage of Weinberg's 1989 no-go — a negative result that fences the problem for everyone who digs after. The next calculation in the program is the modular fluctuation $\langle \Delta K_D^2 \rangle$ in the Bunch–Davies state for a de Sitter diamond: the one step where nobody knows the answer in advance.

The journey in one line

A speculative picture → a big idea → many possibilities → successive refinements and eliminations → a technical program → a quantitative verdict (way too small) → a final verdict (for now): $\delta\Lambda_{\rm ent} = 0$.

The two honest endings (and how to tell which one matters)

Let us be scrupulously clear about something, because clarity here is the whole game: this programme has not established any breakthrough. What the paper demonstrates is a first-order null result under stated assumptions — no more, and, rather wonderfully, no less. What a *successful future* programme could imply is another matter entirely, and the distinction deserves to be framed and hung where visitors can see it.

If — and it is an "if" doing heroic load-bearing work — a later stage survived peer scrutiny and made new, confirmed predictions, the implications would be genuinely profound: a de Sitter version of "geometry from entanglement," extending the AdS successes to the universe we actually appear to inhabit; a parameter-free explanation for the cosmological constant, which would address one of the deepest embarrassments in theoretical physics; a working bridge between quantum information and cosmology, unifying programmes that are today connected by conceptual sympathy rather than shared mathematics; and a deeper understanding of general relativity — not its replacement, but its explanation, in precisely the way statistical mechanics explains thermodynamics without ever once firing it. And the marvellous thing — the genuinely thrilling thing — is that none of this requires the original poster to have been right. The groundbreaking outcome was never going to be "the egg was correct all along." It is, and always was: *the process of rigorously testing the idea either uncovers a genuinely new mechanism, or rules out an entire class of them.* Both are science. Only one gets a documentary.

Two edges of this want sharpening, and they should be built into the programme now rather than discovered later with regret.

First, the no-go branch: its value will scale with its generality, and generality must be designed in. The present null result is narrow — conformal matter, coherent states, first order — which

makes it a lemma, not yet a theorem. Weinberg's 1989 no-go became one of the most load-bearing results in the entire field precisely because it covered *all* dynamical adjustment mechanisms under mild assumptions: a negative result that fenced the problem for everyone who dug afterwards, cited more often than most positive proposals. Notes 2 and 3 should therefore be written not merely to check whether particular loopholes close, but to close the *widest class* each can honestly reach — so that if the programme terminates at the negative branch of Note 4, what it leaves behind is a fence, not a fence post.

Second — and here a moment of pedantry with a purpose — a "specific $w(z)$" is not, by itself, the decisive step, however satisfying it would feel to write one down. The background expansion history produced by matter–vacuum exchange can be made *exactly degenerate* with quintessence and the standard CPL parameterisation: a confirmed $w(z)$, even one matching a pre-registered prediction, would be claimed by every model in the class within the week. The genuinely decisive object is the *joint* prediction $\{w(z),\ f\sigma_8(z)\}$ — the growth-of-structure suppression correlated with the redshift of the phantom crossing — because that correlation is what an exchange mechanism forces and quintessence does not. If an exchange law ever emerges from Note 3, that pair, registered publicly before Euclid DR1 and Roman report, is the form any landmark claim must take. Anything less is an elegant construction politely queueing with its competitors.

So: two endings, both honest. A no-go theorem — increasingly general, quietly useful, the kind of result that saves other people decades. Or a unique, parameter-free exchange law that must then survive the mathematics, the referees, and finally the sky — physics being, mercifully, a democracy of evidence rather than of enthusiasm. Only the second could become a landmark, and even then only after independent verification, peer acceptance, and observational confirmation. Neither ending is what the poster hoped for. Which is, by now, the programme's proudest feature rather than its concession: **the value of the work no longer depends on whether the original idea survives.** The gallery that follows preserves the full artifacts, unretouched, so the eliminations — and the promises — can be audited.

The programme flow, as it stands

Two flows, kept deliberately separate. The first is the *physics* flow — the information cycle in the only form the equations permit, everything upstream of Q now exact and established, everything downstream standard cosmology. The second is the *ideation* flow — the decision tree the programme now walks, with every branch terminating in a publishable result. Note where the coin flips live: nobody knows the outcome of the shaded nodes in advance, and that is precisely what makes them worth computing.

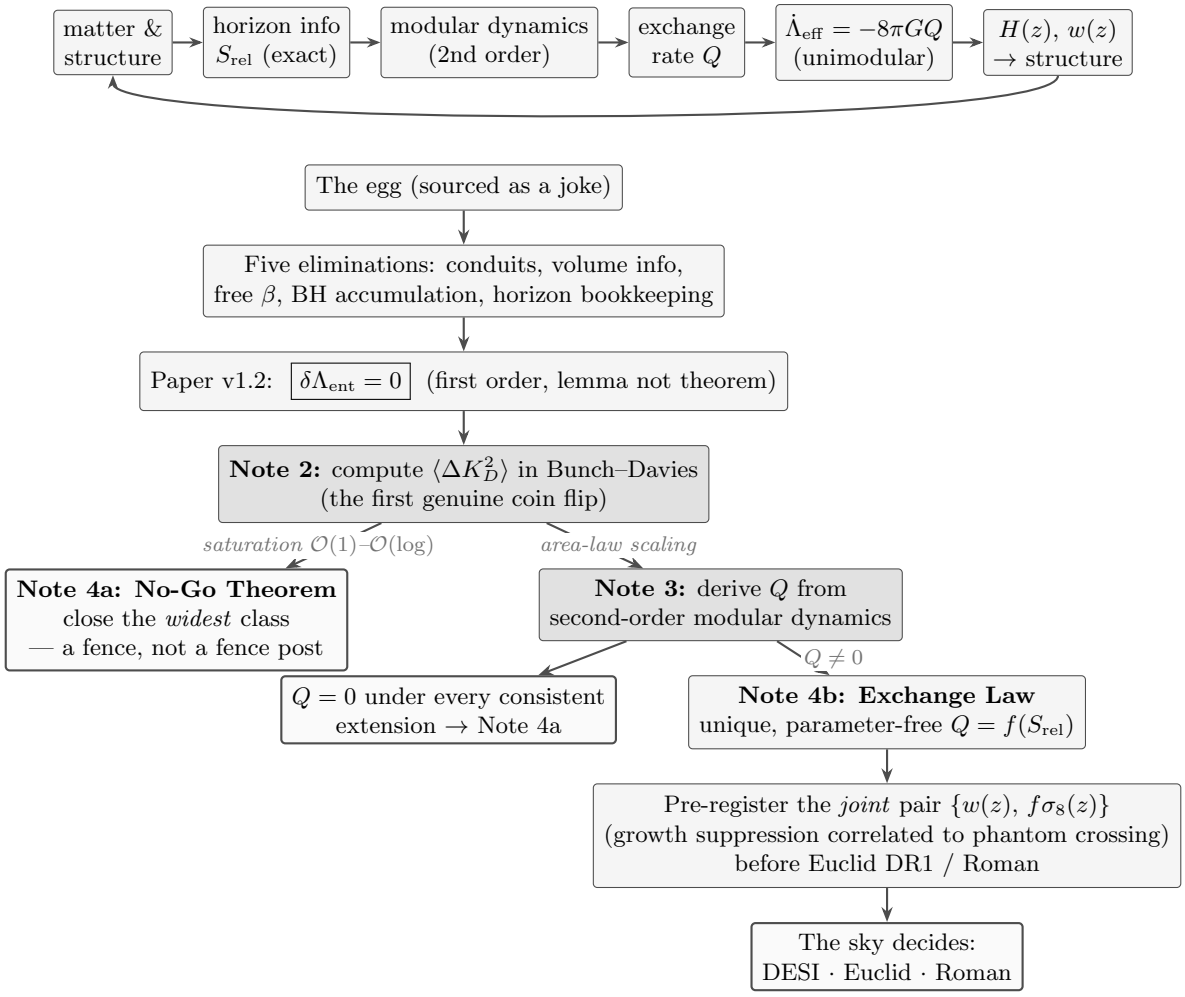

Every leaf is publishable; no branch ends in ambiguity. The left column is quietly useful and saves other people decades; the right column is the only path to a landmark — and even it must survive the mathematics, the referees, and finally the sky, in that order.

Gallery I: the evolution map

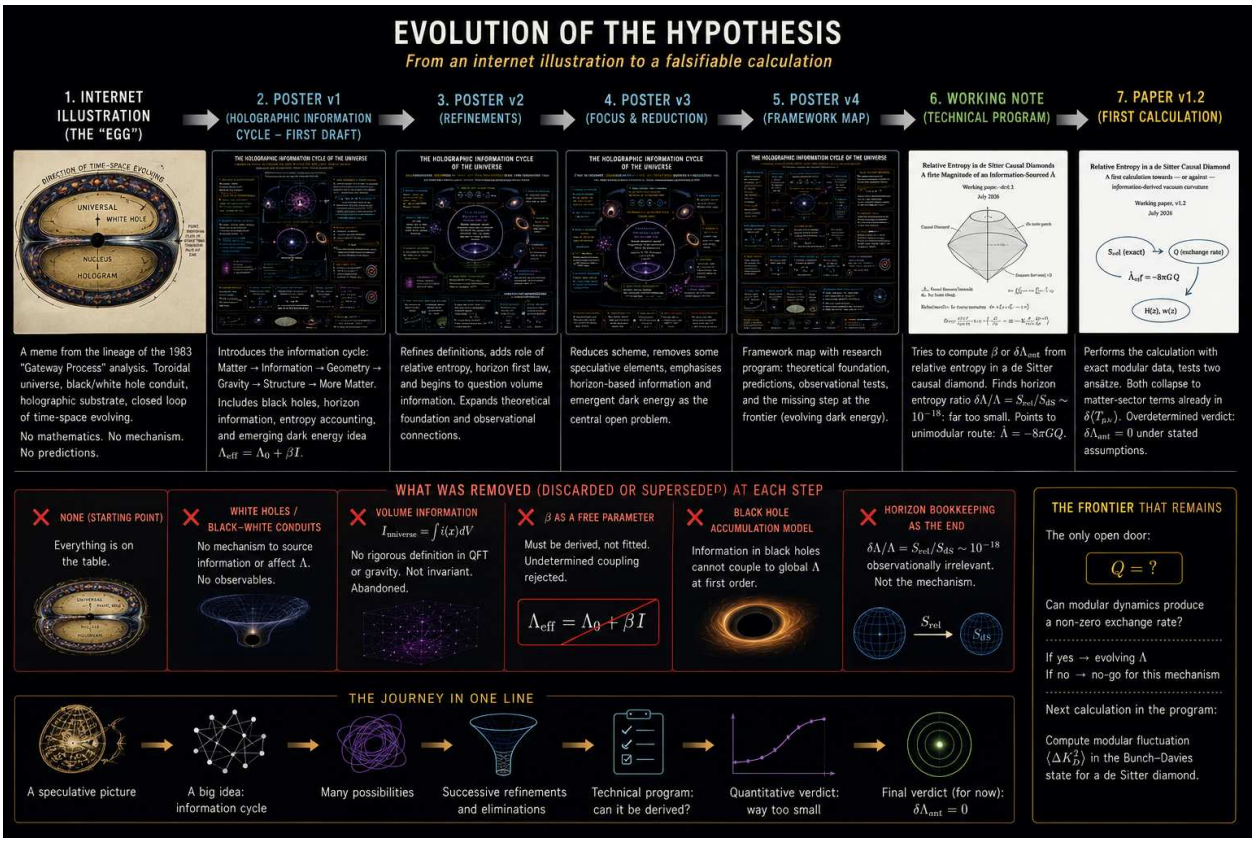

The seven-stage evolution with the removal track: what was discarded or superseded at each step, and the frontier that remains.

Gallery II: Poster v1 (first draft)

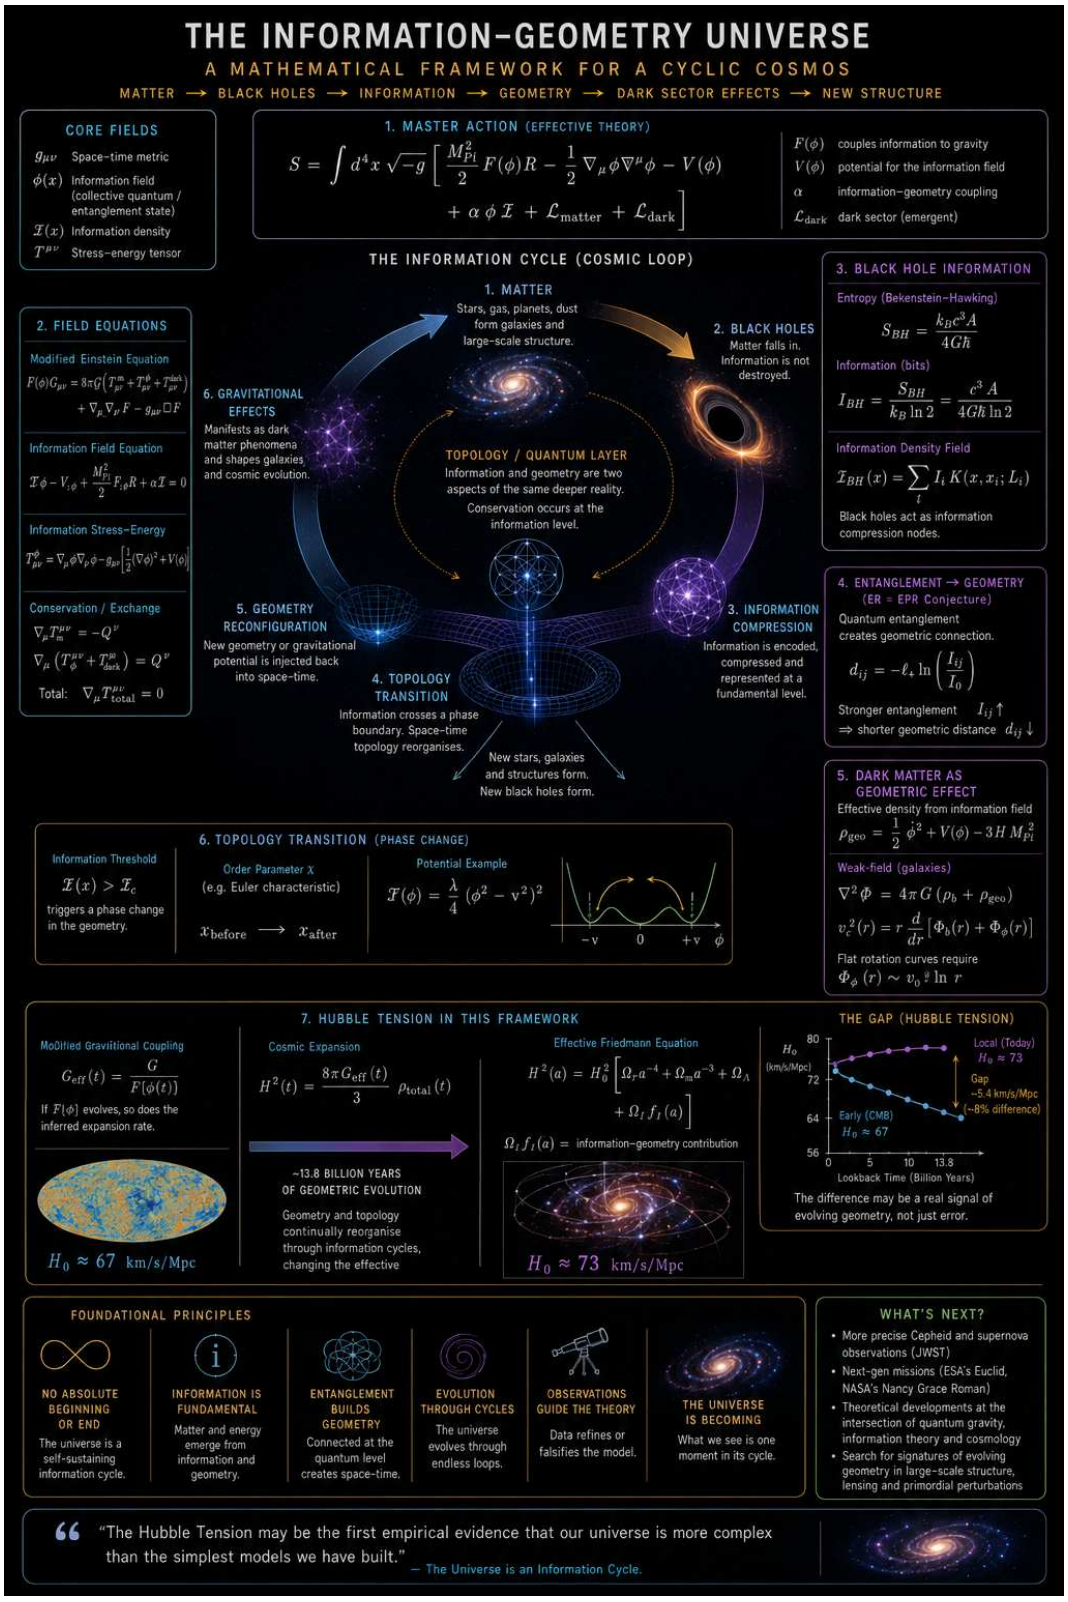

Gallery III: Poster v2 (refinements)

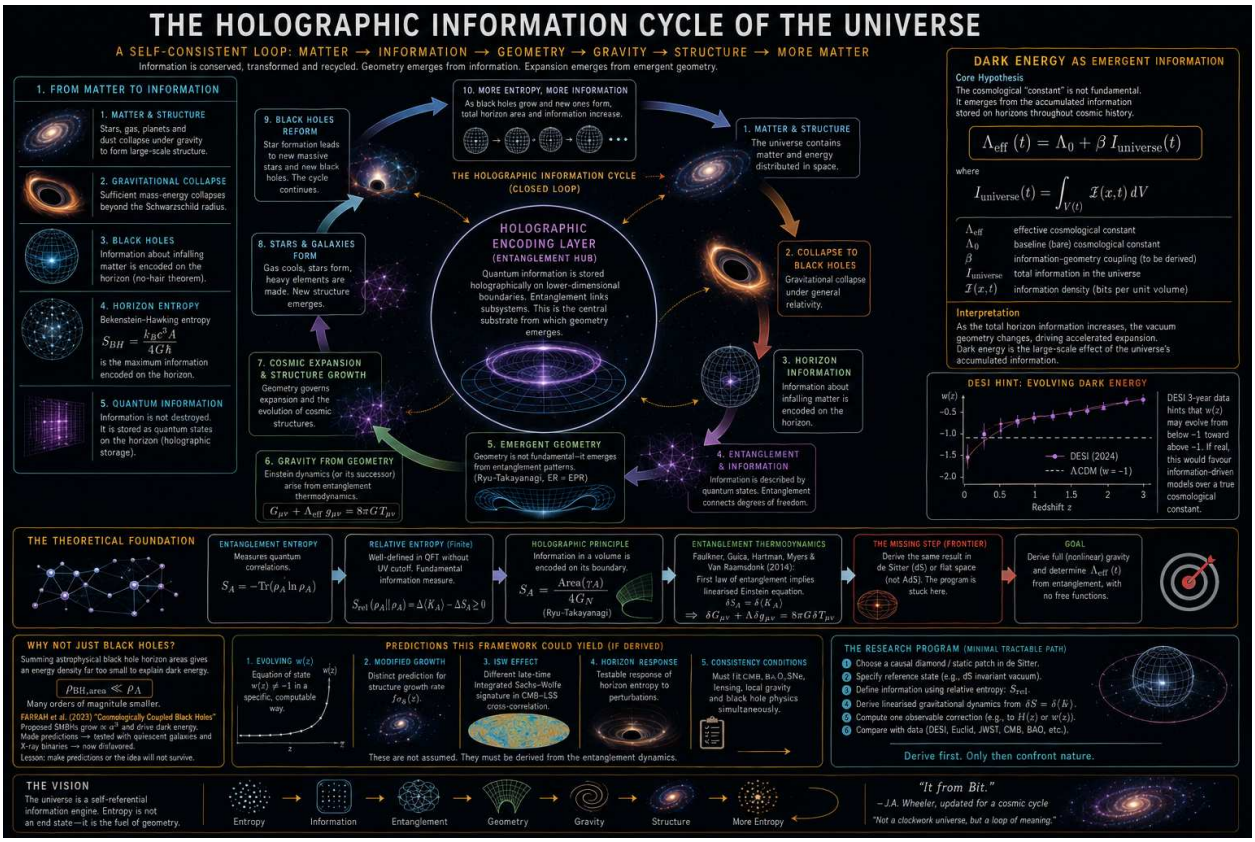

Gallery IV: Poster v3 (focus and reduction)