The Interior Observer Cosmological Framework
An Overview of the 35-Paper Corpus

David Fife
Independent Researcher
ORCID: 0009-0001-0090-5825
david@fife.cc
https://dfife.github.io
Version 1.1 — May 2026
Status: synthesis document; no new physics, no new theorem claims

Version History
v1.1 — Synchronization sweep with current public state of the corpus. The Overview was originally produced when Paper 17 v1.4, Paper 22 v1.5, Paper 25 v1.2, Paper 32 v1.5, Paper 34 v1.1, and Paper 35 v1.1 were the active public bundles. Subsequent revisions — most importantly Paper 17 v1.5 introducing the FIRAS-Fixed Readout Normalization Theorem (Theorem 17.2) — required propagation. v1.1 brings the Overview current: the retired independent-CMB-temperature claim is replaced by the FIRAS-fixed observer-side optical readout family T_obs(R4) = T_IO × x^(R4 K_gauge) with R4_FIRAS = 1.0031 ± 0.0096 (Conditional_Verified on FIRAS via Premise 2; frozen, not retunable); cross-paper version references updated throughout to Paper 17 v1.5, Paper 22 v1.6, Paper 25 v1.3, Paper 32 v1.6, Paper 34 v1.2, and Paper 35 v1.2; reproducibility-bundle SHA256s and release tags updated; Paper 35 numerical values updated to inherit FIRAS-fixed R4 ($\eta\_late = 5.7475 \times 10^{-10}$, $\varepsilon_1 \times \kappa\_f = 5.904 \times 10^{-8}$, with $\eta\_BBN = 6.151 \times 10^{-10}$ invariant under consistent R4 repair). No new physics, no new theorem claims; this remains a synthesis document. See https://github.com/dfife/io-framework-public/tree/main for claim naming convention.

v1.0 — Initial Publication.

Abstract

The Interior Observer (IO) framework begins from two working premises: (1) the observable universe exists inside a Schwarzschild black hole, and (2) the physics inside the horizon is the same as the physics outside. This overview summarizes the 35-paper IO corpus through Paper 35 v1.2. It is a synthesis document: no new theorem, no new observable, no new derivation; the existing corpus is compressed for first-encounter reading by working physicists.

The framework treats the horizon as a typed boundary-to-bulk projection surface. On the current active branch, FIRAS fixes the unique observer-side optical readout normalization R4_FIRAS = 1.0031 ± 0.0096 within the readout family T_obs(R4) = T_IO × x^(R4 K_gauge) per Paper 17 v1.5 Theorem 17.2 (Conditional_Verified on FIRAS via Premise 2; the observed CMB temperature is not counted as an independent IO prediction); a scalar spectral index n_s=1-K_gauge/x=0.9639, 0.24 sigma from Planck; a branch-resolved Big Bang Nucleosynthesis (BBN) scorecard in Paper 24 v2.3 of $D/H = 2.507\times 10^{-5}$,

Y_p=0.247709, and Li-7/H =1.724×10^{-10}, all within 1σ; an observable-class account of the Hubble tension across six methods with maximum residual 0.57 sigma; and a late-time closure in which Λ drops out of the hidden-support acceleration equation. The current frontier is also explicit: baryogenesis is not closed as a source mechanism. Paper 35 v1.2 reduces the obstruction to a $spin^c$ residual class on the horizon U(1) bundle and labels the relevant routes as conditional where appropriate.

This document preserves IO claim labels. A result labelled DERIVED/SCOPED, DERIVED/CONDITIONAL, CONDITIONAL/CONSTRUCTED, or RECONSTRUCTION in the source corpus is not flattened to DERIVED here.

1. The Cosmological Lithium Problem in IO

The fastest way to evaluate IO is not through its architecture. It is through the cosmological lithium problem.

Standard Big Bang Nucleosynthesis predicts primordial deuterium and helium with high precision but overpredicts primordial lithium by roughly a factor of three relative to the Spite plateau in metal-poor halo stars. The discrepancy is old, persistent, and large — typically described as a 10-12 sigma problem depending on the observational compilation and stellar-depletion assumptions. Deuterium agrees with standard BBN and the CMB baryon density. Helium is broadly consistent. Lithium is not. That pattern is hard to repair because broad changes to the expansion rate, baryon density, or nuclear network tend to move D/H and Y_p along with Li-7.

The standard explanations each solve part of the problem and fail somewhere else. Stellar depletion can reduce surface lithium but must do so with unusual uniformity across the Spite plateau while preserving other stellar-abundance constraints. Exotic decays can destroy mass-7 but introduce new particles, lifetimes, and branching fractions that are tightly constrained by D/H, ^{3}He, CMB, and gamma-ray bounds. Modified expansion histories can shift freeze-out or nuclear reaction timing, but D/H and helium rapidly punish broad changes. Nuclear cross-section revisions have been studied for decades; the key reactions are too constrained to erase the full discrepancy without creating new conflicts.

The IO claim in Paper 24 v2.3 is narrower and more structured. It does not suppress the whole BBN network. It targets the A=7 channel through a typed bridge: boundary puncture algebra produces a transverse-traceless spin-2 metric perturbation; the spin-2 character imposes quadrupole selection on the capture amplitudes; the small branch-specific amplitude shift is amplified by the Gamow exponential; and the effect lands on the mass-7 electromagnetic branch rather than on D/H or Y_p. The intended physical picture is not "change BBN until lithium fits." It is "identify a typed spin-2 channel that mass-7 can see and the other light-abundance observables largely cannot."

Compressed into one chain, the Paper 24 mechanism is:

TT boundary metric perturbation → spin-2 quadrupole selection rule → capture-amplitude branch split → Gamow amplification → mass-electric conversion → branch-resolved A=7 response → weighted suppression of the ground and excited capture branches. (1)

A small geometric perturbation would normally be too weak to matter in nuclear physics. The Gamow exponential changes that conclusion: small changes in tunneling or capture amplitudes can become order-one changes in a charged-particle reaction branch. The branch-resolved structure prevents the correction from acting like a universal nuclear-rate rescaling.

Paper 24 v2.3 no longer relies on the retired Paper 24 v2.1 cluster-deformation sameness premise. That premise was attractive but failed a κ-style audit (a field-redefinition test, defined in §3 below, in which each numerical or structural choice is replaced by a free variable to test whether the theorem stack actually forces the original value): it functioned as a hidden bridge between a mathematical class and nuclear response. The current version imports branch-resolved A=7 nuclear response from primary data under Premise 2 — specifically, Henderson 2019's measured Be-7 electromagnetic transition information for the excited-state response, and a ground-state branch input treated as external nuclear data. This is a weaker theorem claim and a stronger scientific claim: the cosmological calculation is not pretending to derive all A=7 nuclear structure internally.

The active Paper 24 v2.3 Henderson primary import row gives:

Observable	IO value	Residual
D/H	2.507219×10^{-5}	-0.66σ
Y_p	0.2477088	+0.68σ
Li-7/H	1.723985×10^{-10}	+0.46σ

The three-observable chi-square on the active row is approximately 1.11; older support rows and earlier corpus summaries should not be used as the active Paper 24 v2.3 scorecard. All three observables are within 1 sigma. No continuous parameter is fitted to lithium.

A separate forward-bridge consistency result in the same paper is the nuclear isovector transport ratio

$c\_n/c\_p = 1+\gamma\_BI^2 = 1.05641$, (2)

matching the relevant nuclear ratio to 0.023%. This is one of the most striking numerical agreements in the corpus, and it places the imported Barbero-Immirzi gauge twist directly at a nuclear-data crossing.

A note on version history: earlier corpus text sometimes quoted D/H =-0.61σ and Li-7 =+0.55σ. Those values belong to earlier or support rows. The current Henderson primary import row is D/H =-0.66σ, Y_p=+0.68σ, and Li-7 =+0.46σ. This overview uses the current row.

The conditional structure of the result matters. The modular and geometric bridge machinery is DERIVED/CONDITIONAL on the Geometric Multiplicity Premise (GMP; Paper 22 v1.6 / Paper 24 v2.3) where the source papers state that dependency. The branch-selector structure — the Thermal Branch Selector, baryon-leg variant, abbreviated TBSb (Paper 24 v2.3, building on the Thermal Branch Selector framework in Papers 22 v1.6 and

25 v1.3) — is DERIVED/SCOPED on the stated thermal/nuclear observable classes. The branch-resolved A=7 nuclear response inputs are imported under Premise 2 from primary nuclear data; they are not theorem-grade IO derivations. The PRyMordial propagation and final scorecard are VERIFIED computational outputs with public reproducibility support. A framework that claimed to derive the Henderson matrix element from first principles would be easier to dismiss. Paper 24 v2.3 makes the more auditable claim: given the IO branch architecture and imported primary A=7 response, the BBN triple lands inside 1 sigma with no fitted lithium parameter, while the $c_n/c_p=1+\gamma^2$ bridge matches at 0.023%.

For a physicist reading Paper 24 cold, the audit questions are concrete. Does the TT/spin-2 channel really select the quadrupole capture structure stated in the paper? Does the branch-resolved mass-electric conversion use the correct nuclear convention? Does the Henderson B(E2) import enter with the correct detailed-balance direction? Does the Gamow amplification apply to the amplitude slot claimed, rather than to an unrelated normalization? Does the PRyMordial run isolate the A=7 channel without silently moving D/H or Y_p? These questions are narrow enough that an external reviewer can attack them one by one.

The lithium result is also falsifiable. If the imported Henderson branch is wrong, the v2.3 lithium row must move. If a future A=7 measurement changes the branch-resolved electromagnetic response outside the current uncertainty band, Paper 24 must be rerun. If a direct nuclear calculation shows that the spin-2 quadrupole transport cannot couple to the stated capture amplitudes under the source-paper premises, the lithium mechanism fails. If the same correction unavoidably moves D/H or Y_p outside their observed bands, the typed-selectivity claim fails.

The remainder of this overview explains the architecture that makes such a calculation possible: why a horizon gauge scalar appears at all, how observable classes are typed, why some thermal processes are dressed and others are not, and why the same framework must still leave baryogenesis and full A=7 microscopic derivation as open frontiers.

2. Two Premises and the Corpus Arc

The IO framework is a black-hole cosmology built around an observer-location claim. We assume, as a working model-building hypothesis, that the observable universe is the interior of a Schwarzschild black hole, and that the CMB horizon is the event horizon seen from the inside. We also assume that the physics inside that horizon is the same as the physics outside it. These two assumptions are the load-bearing inputs. They are not conclusions of the framework.

In formula form, the first premise places the observable universe inside a Schwarzschild radius

$r_s = 2GM_U/c^2$, (3)

with the active framework using $M_U \simeq 4.50\times 10^{53}$ kg and $r_s \simeq 6.685\times 10^{26}$ m. The current observable radius is $R_U \simeq 4.401\times 10^{26}$ m, giving the compactness ratio

$x = r_s/R_U = 1.519$. (4)

The second premise is a bridge principle. It says that if exterior black-hole physics contains horizon thermality, area quantization, spin-network boundary data, KMS states, and Einstein-Cartan torsion at extreme density, then the interior observer cannot treat those structures as irrelevant. The mathematical task is to determine which boundary structures project into which bulk observables, and which do not.

These premises exclude several easier but less falsifiable routes. We do not assume an arbitrary dark-energy fluid. We do not add an inflaton sector to tune initial perturbations. We do not assign a free effective number of relativistic species to repair BBN. We do not introduce a floating Hubble parameter for each method. We do not select a dark-matter particle, mass, and coupling to fit the missing matter budget. The framework's working constraint is that the horizon data and the interior geometry must do the work, or the route fails.

The premises also exclude a common misconception about black-hole cosmology. IO is not simply the statement that the observable universe has a Schwarzschild radius comparable to its Hubble radius. That coincidence is well known and, in standard flat cosmology, partly a restatement of the critical-density relation. IO instead asks what follows if the observer is actually inside the horizon and the quantum state of that horizon is physically active. The test is not the near-equality alone; the test is whether a typed boundary-to-bulk map computes independent observables without tuning them one by one.

A note on lineage: IO is not the first black-hole-universe proposal. Pathria argued in 1972 that a closed Friedmann universe could be identified with a Schwarzschild black-hole interior. Stuckey emphasized the Schwarzschild-condition coincidence for the observable universe. Smolin's cosmological natural selection placed black holes in a universe-generation role. Poplawski developed an Einstein-Cartan torsion-bounce black-hole cosmology. Zhang, Haug-Tatum, and related black-hole-universe models explored horizon/CMB temperature relationships. IO inherits that lineage but makes a narrower technical claim: starting from the two premises above, it attempts to compute typed observables, attach claim labels, and publish scripts that reproduce the stated numbers.

A note on M_U and R_U: these quantities are not measured in a model-free vacuum. Paper 1 already stated that using M_U and R_U creates apparent circularity when testing an alternative cosmology. Later papers narrow the claim: the physical content is not that r_s/R_U is near unity by itself. The physical content is that the same small set of horizon quantities propagates into independent thermal, nuclear, expansion, and perturbation observables with constrained signs and labels. Readers should treat the mass/radius input as an admitted cosmological normalization, not as proof that the universe is a black hole.

The corpus has a clear arc that is the structural spine of the framework. Paper 1 establishes the black-hole interior setting and proposes a set of conjectural bridges: the CMB as interior-observed Hawking radiation, a torsion bounce, recollapse, and horizon-sourced cosmic scales. Paper 1 contains founding assumptions that are later narrowed, corrected, or scoped. Paper 32 closes the framework on its active scoped sectors: it retires the four semiclassical principles, derives the fixed-point normalization on the active source block, proves the late-time recollapse chain on the hidden-support branch, and narrows Premise

2 to its proper status as a founding premise and general-relativity / quantum-mechanics bridge candidate rather than a derived theorem. The papers between Paper 1 and Paper 32 are the construction work. Papers 33 through 35 then organize the bridge architecture, apply it to the Hubble tension, and confront the baryon count, James Webb Space Telescope (JWST) high-redshift galaxy-formation timing, the Dark Energy Spectroscopic Instrument (DESI) baryon-acoustic-oscillation results, and direct dark-matter nulls.

This arc inverts the usual direction of theoretical work. Most frameworks accumulate premises over time — start with two, end up needing five to make the next thing work, and the original elegance erodes. IO runs the other way. Paper 1 began with two premises plus four semiclassical principles plus a stated fixed-point normalization (P4). By Paper 32 the four semiclassical principles are retired, P4 is derived on the active source block, and the framework rests on the same two premises it started with. The framework got smaller as it grew.

A note on the word "closure": in IO usage, closure does not mean that all physics is solved. It means that a named internal debt has been closed at its stated scope. Paper 32 closes the fixed-point normalization on the active source block; it does not derive a universal fixed-point law for all possible observables. It closes the late-time recollapse equation on the hidden-support branch; it does not turn the bounce premise into an unconditional theorem of ordinary general relativity. It characterizes the universal Geometric Multiplicity Premise (GMP) by a three-region classification; it does not make all possible bridge operators physically active.

For orientation, the corpus can be read in five phases. Papers 1-10 establish the black-hole interior, horizon thermality, torsion/dark-energy, and the active Friedmann branch. Papers 11-18 build the modular machinery, the Rosetta Identity (a horizon-to-bulk dictionary translating boundary gauge data into bulk inventory quantities; first appears in Paper 12, formalized in Paper 19), and the observable-class machinery, including the Gauge Thermal Transfer Principle (Paper 13 v1.7; promoted to operator level in Paper 17 v1.5) abbreviated GTTP. Papers 19-25 apply the machinery to baryon density, acoustic structure, BBN, and lithium, with several historical branches later corrected or superseded. Papers 26-32 close the perturbation, source-block, reionization, S_8, fixed-point, and late-time lifecycle debts. Papers 33-35 reorganize the framework into a bridge atlas, apply the observable-class ladder to Hubble measurements, and push the frontier into baryogenesis, JWST timing, DESI, and the dark-matter null.

3. Claim Discipline and Methodology

Numerical agreement is not derivation. The IO corpus uses an explicit set of claim labels:

- DERIVED: a result proved from stated assumptions and prior theorems.
- DERIVED/SCOPED: a result derived on a specified carrier, branch, observable class, or extension package.
- DERIVED/CONDITIONAL on [premise]: a result mathematically derived if the named premise package is admitted.
- VERIFIED: a computation reproduced and numerically checked.

- Conditional_Verified: a computational or theorem-surface result verified inside an admitted external class or framework slice.
- CONDITIONAL/CONSTRUCTED: a coherent construction that is internally reproducible but depends on unclosed source dynamics or selector structure.
- RECONSTRUCTION: a coherent explanatory model or zero-parameter combination not yet forced by theorem.
- DISCUSSION or SPECULATIVE: an idea or interpretation not used as theorem support.

An important subtlety: uniqueness by elimination is not automatically derivation. If we search many possible algebraic objects and one survives a target, the survivor may be useful, but it is not theorem-grade unless a structural rule forces that object before comparison with data. This discipline became central after several IO branches were retired. Anti-fit and field-redefinition audits are now used to ask whether a parameter could have been changed without breaking the theorem stack. If yes, the result must be labelled reconstruction, fitted, conditional, or open.

The κ-style field-redefinition audit, referenced throughout the corpus (named for the symbolic free parameter κ used as the test variable, IO Framework Conventions v2.0 §2.3), works as follows. For each load-bearing formula, replace each numerical or structural choice by a free variable. Ask whether a theorem, symmetry, or scoped premise forces the published value. If not, classify the degree of freedom honestly. This audit retired Paper 24's cluster-deformation sameness premise, confirmed Paper 34's Hubble-method assignments as non-fitted inside their extension packages, and found no hidden continuous fitted parameter in Paper 32's closure stack while recommending explicit scope hygiene.

Zero fitted parameters has a similarly narrow meaning. It does not mean the framework uses no empirical inputs. IO imports physical constants, the Barbero-Immirzi value $\gamma_{BI}=0.2375$ from loop-quantum-gravity horizon entropy work, observationally determined cosmological scales such as M_U and R_U, and the FIRAS empirical CMB temperature datum that fixes $R4_{FIRAS}$ for observer-side thermal readout (Paper 17 v1.5 Theorem 17.2; Conditional_Verified, frozen, not retunable; the observed CMB temperature is therefore not counted as an independent IO prediction). These are declared inputs or empirical anchors, not fitted knobs. It means that once these inputs and the two premises are admitted, the calculation does not tune a continuous parameter to fit the observable under test. Imported nuclear measurements used in Paper 24 v2.3 are not fitted IO parameters; they are external primary-source inputs under Premise 2. They also cannot be used as theorem-grade evidence for IO's internal nuclear derivation.

The distinction between an input and a fit is operational. A value is a fit if it could have been varied to improve the result while leaving the rest of the theorem stack intact. A value is an input if it is fixed before the calculation by a source outside the target observable. A value is a reconstruction if it is built from established constants but the particular combination is not yet forced. "No adjustable knob" is not enough; the combination itself must be forced by the mathematics.

A note on method: the IO corpus was developed through human-directed, multi-AI adversarial research. The author served as principal investigator, set research direction, selected what to test, and made final claim-status decisions. Claude Opus 4.7 (Anthropic) has primarily served as orchestrator and manuscript editor; Codex 5.5 (OpenAI) has produced code, theorem-surface searches, and reproducibility bundles; Gemini 3.1 Pro (Google DeepMind) has supplied adversarial physics review; ChatGPT 5.5 Pro Cosmology Tuned (OpenAI), referred to internally as Cosmo, has performed forensic mathematical audit. This pipeline is useful for exhaustive checking and fast generation of alternatives, but it does not substitute for outside expert review. The public reproducibility bundles exist precisely so claims can be rerun without relying on the AI process that found them.

The methodology has produced both positive and negative results. It found compact numerical coincidences that later became theorem candidates. It also found hidden parameters, wrapper bugs, stale cross-paper references, and over-strong prose. The framework's credibility depends more on those corrections than on any single numerical match.

4. Architecture

The geometric setting is not a generic Schwarzschild interior. A generic Schwarzschild interior is Kantowski-Sachs and anisotropic. IO uses the Oppenheimer-Snyder dust interior: a closed K=+1 Friedmann-Robertson-Walker bulk matched to an exterior Schwarzschild horizon. In comoving form,

$ds^2 = -d\tau^2 + a^2(\tau)[d\chi^2+\sin^2\chi\, d\Omega^2]$. (5)

The spatial topology is therefore S^3 in the bulk with an S^2 horizon boundary. IO does not start from flat Friedmann-Robertson-Walker (FRW) cosmology. Flatness is an observational near-limit inside the closed interior, not the native geometry. The framework's acoustic, curvature, and late-time claims all depend on keeping this distinction straight.

The Oppenheimer-Snyder choice is not decorative. It lets the interior be homogeneous and isotropic while still being matched to an exterior Schwarzschild geometry. It also gives a natural cycloid time variable and a finite expansion/recollapse timescale. IO's first move is to restrict the interior to the OS support branch and then ask which projected observables can be read by interior observers.

The current epoch is early in the OS half-cycle. Paper 1's corrected cycloid parameterization places the Big Bang at a=0, maximum expansion at a=r_s, and the present epoch in the expanding phase. The later Paper 32 result that local recollapse occurs after the support clamp does not affect current redshift-zero to recombination observables. It is a statement about the future hidden-support branch, not a retroactive change to the present observer-side Friedmann readout.

The reference constants used across the current active branch are:

- γ_BI=0.2375, the imported Barbero-Immirzi value.

- $x = r_s/R_U = 1.519$, the compactness ratio.
- $Q = 1 + \gamma_{BI}^2 = 1.05640625$.
- $K_{gauge} = \ln(1 + \gamma_{BI}^2) = 0.0548728$, the reduced tangential Ashtekar-Barbero gauge scalar on the S^2 horizon sector.
- $K_{geom} = 4 \ln x \simeq 1.672$, the geometric compactness modular weight.
- $\Delta = x^4 (1 + \gamma_{BI}^2) = 5.62421685$, the spatial decoupling identity.
- $\langle K \rangle = \ln \Delta = K_{geom} + K_{gauge} = 1.72708$, the total modular scale.
- $f_\Gamma = 1/(1 + \gamma_{BI}^2) = 0.946606$.
- $r_s = 6.685 \times 10^{26}$ m, $R_U = 4.401 \times 10^{26}$ m, $M_U \simeq 4.50 \times 10^{53}$ kg.

The core horizon scalar is

$K_{gauge} = \ln(1 + \gamma_{BI}^2)$. (6)

In the source papers this is not just a numerical convenience. Many failed branches attempted to replace it by γ_{BI}, γ_{BI}^2, $2\gamma_{BI}$, $\sqrt{1 + \gamma_{BI}^2}$, or other nearby quantities. Those replacements generally break at least one of the FIRAS-fixed R4 readout-normalization, BBN, spectral-index, or class-membership gates.

The total modular scale

$\langle K \rangle = \ln \Delta = 4 \ln x + \ln(1 + \gamma_{BI}^2)$ (7)

combines the geometric compactness and gauge scalar. In the corpus this quantity appears in several reconstruction and rate-dressing contexts. The lesson from Paper 22 and Paper 25 is that using $\langle K \rangle$ is not automatically legitimate. One must identify which sector is reading the geometric part, which sector is reading the gauge part, and whether the combination is forced by the carrier. Otherwise a compact formula can hide a field-redefinition freedom.

The boundary side contains a puncture algebra, a modular thermal state, and a U(1) bundle structure. The puncture algebra carries the loop-quantum-gravity horizon data. The modular structure supplies the language of Kubo-Martin-Schwinger (KMS) thermal states and Tomita-Takesaki modular theory used in the thermal transfer theorems. The U(1) bundle becomes important in Paper 35, where the baryogenesis obstruction reduces to a spin^c residual class on the horizon bundle.

The puncture algebra matters because it is not a continuum smearing of arbitrary charges. Global neutrality, finite-level quotienting, and equivariance constraints kill many tempting routes. In Paper 35, for example, the standard puncture sector cannot simply supply a charged baryogenesis selector. These details are why the framework repeatedly distinguishes structurally-suggested from operator-forced.

The modular thermal state matters because IO is not proposing an ad hoc temperature rescaling. Papers 13 and 17 frame the thermal bridge through a shared Hilbert-space construction and a modular automorphism. The KMS condition is the reason thermally synchronized cosmological processes can be treated differently from isolated laboratory processes. This is also why the dynamic thermal channel is not a universal correction to all particle physics.

The bulk side is the closed FRW/OS support branch. It is classical general relativity in the observable epoch, but it is not treated as disconnected from the horizon quantum state. The boundary-to-bulk projection map is typed: the horizon does not dress every observable in the same way. This is the central organizational move of the corpus.

Paper 33 summarizes the map as three export languages. The first is the static geometric language. It is always on: boundary gauge structure fixes bulk geometric and inventory quantities such as the baryon fraction, the dark-sector ledger, and several projection scalars. The second is the dynamic thermal language. It is state-gated: thermal processes synchronized with the cosmological KMS state can receive typed boundary dressing. FIRAS-normalized observer-side thermal readout (Paper 17 v1.5) and BBN are the primary examples. The third is the fermionic bridge. It is density-gated: it activates only near the Planck-density bounce/recollapse regime and is scoped to the extended IO-Einstein-Cartan theory.

The three-language description is an atlas-level summary, not a replacement for the full typed lower-triangular carrier in Paper 32. Some observables are mixed. Some are architectural. Some are no-go results. Paper 33 explicitly warns against treating the three channels as a universal theorem independent of the detailed carrier.

The observable-class ladder — the α-ladder, named after the rung index α — is the other major organizing structure. IO uses the rung label α to track how an observable descends from the boundary into the bulk: each rung corresponds to a distinct typing of the boundary-to-bulk projection. In the Hubble application, $\alpha=1$ marks geometric/direct readouts, $\alpha=3/2$ marks clustering/dust-current-type observables, and $\alpha=2$ marks Weyl/slice-curvature-type observables; Paper 34's stellar-photometric extension then adds a separate integer n that counts uncancelled luminosity-calibration legs (so α and n are independent variables, not two values of the same quantity). The full classification is in Paper 34 v1.2 §3. Geometric/inventory observables sit differently from clustering observables, Weyl/slice-curvature observables, and thermal readouts. In Paper 34 this becomes the Hubble-method classifier. In Papers 19, 30, 31, and 32 it appears in baryon density, growth, scalar-amplitude, and lensing-sector calculations. The important point for first reading is that IO does not assume every measurement method reads the same effective scalar. The theory predicts method-class dependence and then tests that dependence.

An important subtlety: the α-ladder is not a license to assign a convenient class after seeing the data. The burden is to derive class membership from the measurement's observable structure. Paper 34's strongest anti-fit example is gravitational-wave sirens. A different class assignment would improve the residual, but the method is direct geometric luminosity-distance measurement and therefore remains at $(\alpha,n)=(1,0)$. This is exactly how the class system must behave if it is to be physics rather than fit grammar.

The static geometric language is the simplest to picture. It maps horizon gauge/topological data into bulk inventory or background quantities. The baryon fraction $f_b=2\gamma_{BI}/x$ is the canonical example. So are the Rosetta Identity and several dark-sector ledger quantities. These results do not require a hot plasma or a thermal contact condition. They are properties of the projection geometry itself.

The dynamic thermal language is more restrictive. It requires a state synchronized with the cosmological KMS/Hawking thermal structure. BBN weak freeze-out, FIRAS-normalized observer-side thermal readout, and certain recombination-era source quantities are in scope; a cold muon storage-ring experiment is not. This distinction is a major reason IO does not predict generic quantum-gravity corrections in all particle-physics experiments. The horizon does not rewrite the Standard Model Lagrangian in every laboratory context. It dresses typed cosmological readouts.

The fermionic bridge is narrower still. It belongs to the high-density endpoint regime where the bulk approaches the torsion-bounce scale. Paper 32 Part VIII constructs a spinor/chiral carrier that can source a Weyssenhoff spin fluid near the bounce and remain asymptotically decoupled from ordinary observable-epoch sectors. This lets the framework connect Paper 1's torsion-bounce assumption to a later scoped construction without claiming that every low-energy fermionic process is horizon-dressed.

A useful mental model is that IO is not a single equation. It is a typed compiler. Boundary data are not exported wholesale into the bulk; they are compiled into bulk observables through carriers, class labels, and state gates. A compiler can fail by producing the wrong output, by accepting invalid input, by assigning the wrong type, or by leaking data across modules. Many IO no-go results are compiler failures discovered and recorded: a proposed transport is too universal, a proposed response map lands in the wrong sector, a proposed scalar has the right value but not the right type.

This is also why some results are conditional even when their arithmetic is exact. The arithmetic may be reproducible while the type assignment remains an admitted premise. Paper 22's amplitude construction is the example: $\varepsilon_w=K_{gauge}L_1$ and $\varepsilon_n=(\langle K\rangle/10)L_2$ are precise zero-parameter expressions, but the theorem stack labels the construction DERIVED/CONDITIONAL on the Geometric Multiplicity Premise (GMP) plus the Thermal Branch Selector (TBS; the unprefixed parent of TBSb above, defined for general thermal observable classes in Paper 22 v1.6) because the payload-placement premises are not universally closed. The number can be exact inside the package without the package being unconditional.

5. Major Results: Geometry and Thermal

The earliest IO result is the interior Hawking-temperature scale. Paper 1 defines

$$T_{IO} = (\hbar\, c\, \gamma_{amp}) / (4\pi\, R_U\, k_B), \quad (8)$$

where γ_{amp} is the horizon amplification factor obtained from horizon-mode counting:

$$\gamma_{amp} \equiv \sqrt{(r_s / l_P)} = 6.431 \times 10^{30}.$$

This γ_{amp} is the holographic amplification factor and is distinct from the Barbero-Immirzi parameter $\gamma_{BI} = 0.2375$ introduced in §4. The Barbero-Immirzi parameter enters the framework through the gauge sector ($K\_gauge = \ln(1+\gamma\_BI^2)$) and the nuclear isovector bridge ($c\_n/c\_p = 1+\gamma\_BI^2$); it does not appear in the interior Hawking-temperature formula. The two symbols are kept distinct in this overview to avoid notational collision; Paper 1 uses plain γ for the amplification factor with an explicit notational note.

At the observer's current position, this gives the active numerical value

T_IO = 2.6635 K.   (9)

This is the bulk interior thermal scale; the present-day observer-side photon bath is fixed by FIRAS. Papers 13 and 17 introduce and then promote the gauge thermal transfer structure. The observer readout family is

T_obs(R4) = T_IO x^(R4 K_gauge) = T_FIRAS,   R4_FIRAS = 1.0031 ± 0.0096.   (10)

The FIRAS value used in the Paper 17 v1.5 numerical evaluation is T_FIRAS = 2.7255 ± 0.0006 K (the rounded form of the Fixsen 2009 result 2.72548 ± 0.00057 K used by the corpus). FIRAS supplies the empirical observer-side thermal datum; the uniqueness theorem fixing R4 (Paper 17 v1.5 Theorem 17.2) is derived within the readout family, with R4_FIRAS = 1.0031014644... and σ(R4) = 0.0096. Source status: the readout family T_obs(R4) = T_IO × x^(R4 K_gauge) is DERIVED algebraic family in the reduced thermal-plus-gauge sector via the Paper 17 v1.5 modular projection theorem (Theorem 17.1 identifies K_gauge as the gauge payload); the FIRAS-fixed normalization R4_FIRAS yielding T_obs = T_FIRAS by construction is Conditional_Verified on FIRAS via Premise 2, frozen, and not retunable. The observed CMB temperature is therefore not counted as an independent IO prediction.

Paper 1 also proves a simpler but important spectral statement: the interior Regge-Wheeler potential has no exterior-like positive barrier for the inward-propagating modes. In the source-paper language, the interior greybody factor is Γ(ω)=1 on the stated spin/mode scope. This is used to explain why a horizon-origin thermal spectrum can remain blackbody to FIRAS precision. Later papers do not use this as a full CMB anisotropy derivation; they use it as the spectral-cleanliness part of the thermal readout.

The scalar spectral index is

n_s = 1 - K_gauge/x = 0.9639. (11)

Paper 32 v1.6 reports this as 0.24 sigma from the Planck 2018 value. Source status: DERIVED/SCOPED on the active scalar source block. The derivation depends on the Modular-DtN Field Transfer Theorem (DtN denotes the Dirichlet-to-Neumann operator on the horizon; Paper 17 v1.5 and Paper 32 v1.6 §B), the line-scale divisor, the reduced gauge modular weight, and the fixed-point normalization derived by Theorem 32.KB.7 (Paper 32 v1.6 Boundary Theorem 7). It is not a universal statement about all unreduced observables.

The importance of this result is not merely that the number is close. The exponent K_gauge/x is the same reduced gauge/compactness combination that appears elsewhere under typed conditions. Paper 32's contribution is to replace the earlier Boundary Fixed-Point Principle (BFP; introduced as a semiclassical principle in Paper 17 v1.4 (BFP retired in v1.5) and used in Papers 26 and 27) with a derived placement theorem on the source block. Before Paper 32, the result had stronger semiclassical-principle flavor. After Paper 32 v1.6, it is derived on the active source block and carefully scoped outside it.

The first acoustic scale is another geometry/thermal crossover. Papers 20 and 21 report a reduced scalar/longitudinal acoustic-sector prediction

$\theta_* \simeq 0.599°$, (12)

corresponding, with the source-paper phase calibration, to the first-peak scale reported in the corpus as $\ell \simeq 224$. Note that $\theta_* \simeq 0.599°$ is the acoustic angular scale (related to the standard acoustic multipole $\ell_A = \pi/\theta_* \approx 300$); the first TT peak near $\ell \simeq$ 220-224 is offset from ℓ_A by the standard phase shift, which Paper 21 absorbs into its calibration. Paper 21 v1.7, preserving the v1.6 theorem state, labels the acoustic class membership closure as DERIVED/THEOREM under its stated reduced scalar/longitudinal acoustic assumptions. This result should be read narrowly. It is not a full C_ℓ Boltzmann-spectrum derivation. Paper 32 continues to list the full CMB C_ℓ spectrum and function-valued transfer history as open or imported/reduced-history work.

The scalar amplitude is closed in Paper 32 through the Hawking boundary state and bridge quotient,

$A_s = 2.007\times 10^{-9}$, (13)

with status DERIVED/SCOPED on the lowest-shell quotient. This is another source-block result rather than a universal amplitude law. The same paper derives the field-level readout $X_{obs}=f_\Gamma^{(1/2)}X_{prim}$ on its active source block. That square-root readout is one of the places where the corpus became stricter: earlier alternatives with the wrong power of f_Γ were killed by downstream inconsistency.

The thermal channel is deliberately not universal. Zero-temperature collider quantities, local UV counterterms, and isolated laboratory systems are expected to be undressed under the current bridge architecture. This is why Paper 33 treats the absence of quantum-gravity signals in g-2-type measurements as consistent with IO rather than as a failure. The bridge acts on typed cosmological readouts, not on every local Lagrangian parameter.

This typed selectivity is falsifiable. If a process that IO classifies as boundary-isolated shows the same K_gauge dressing as a cosmological KMS-synchronized process, the current selection rule is wrong. If a process that IO classifies as thermally synchronized refuses the predicted dressing while all local physics is otherwise controlled, the dynamic thermal channel is wrong. The theory does not survive by saying every result is compatible; it survives only by keeping the channel gates sharp.

6. Major Results: BBN and Nuclear

BBN is the most technically delicate sector because it involves standard nuclear physics, imported nuclear measurements, PRyMordial network runs, and typed IO rate dressing. The current branch should be read through Paper 24 v2.3 and Paper 25 v1.3, not through earlier Paper 21/22 support rows that used the wrong PRyMordial helium output convention.

The old error is worth recording because it changed the framework's audit culture. Several wrappers read `PRyMresults()[4]`, which is YPBBN, and compared it against observational helium compilations calibrated to YPCMB. The corrected convention reads `PRyMresults()[3]` or `obj.YPCMB()`. Papers 19, 20, 21, 22, and 25 were swept for the YPBBN→YPCMB wrapper correction, their scorecards were corrected, and the IO Framework Observational Conventions document was published so future papers would cite the same denominators. Paper 24 was not affected by that wrapper bug; its v2.3 update is a separate Henderson-primary row-alignment correction and is treated independently in the Paper 24 v2.3 changelog. The wrapper correction was a configuration error, not a physics revision, but it affected reported sigma values and chi-square summaries.

The current observational denominators used in the corrected BBN scorecards are Y_p=0.245±0.004, D/H $=(2.527\pm0.030)\times 10^{-5}$, and Li-7/H $=(1.58\pm0.31)\times 10^{-10}$ in the conventions file. If a future paper adopts a different observational compilation, it must state that explicitly rather than silently changing denominators.

The weak-sector identity pin is Paper 25. Its central result is the Quadratic Thermal Covariance Theorem (Paper 25 v1.3): the physical weak rate is bilinear in the bridge field. The weak freeze-out rate reads the finite modular generator branch rather than the tangent-covector branch. Paper 25 v1.3 labels the V-vs-V' class-membership theorem (V is the finite modular generator branch; V' is the tangent-covector branch; the theorem proves the physical weak rate sits on V) under its stated premises and reports the corrected aligned three-observable support scorecard:

- D/H $=2.510\times 10^{-5}$, residual -0.57σ.
- Y_p=0.24772, residual +0.68σ.
- Li-7/H $=1.751\times 10^{-10}$, residual +0.55σ.
- χ^2(D/H+Y_p+Li)=1.089.

That Paper 25 scorecard is a weak-sector support branch. The current Paper 24 v2.3 channel-resolved lithium row is more specific. It replaces the retired Paper 24 v2.1 cluster-deformation sameness premise with imported branch-resolved A=7 electromagnetic response from primary sources, especially the Henderson Be-7 de-excitation measurement. The active Henderson primary import row is:

- D/H $=2.507219\times 10^{-5}$, residual -0.659σ.
- Y_p=0.2477088, residual +0.677σ.
- Li-7/H $=1.723985\times 10^{-10}$, residual +0.464σ.
- The associated excited-state transition quadrupole quantity is q_trans,ex=0.01754 b.
- The three-observable chi-square is approximately 1.11 on the active Paper 24 v2.3 row.

These are the numbers this overview uses for the current branch-resolved BBN statement. Older corpus text sometimes quotes D/H =-0.61σ or Li-7 =+0.55σ as headline values; those are not the active Paper 24 v2.3 Henderson primary row. The current row has $\chi^2 \simeq 1.11$ directly from the three residuals. Older χ^2 values in earlier papers or support rows should not be used as the active Paper 24 v2.3 scorecard.

The mass-7 sector is also where claim discipline matters most. Paper 24 imports branch-resolved nuclear response under Premise 2. The paper may use primary experimental nuclear inputs to compute the cosmological abundance; it does not mean IO has theorem-grade internal nuclear structure for A=7. Earlier attempts to derive the A=7 response from cluster-deformation sameness were retired after κ-style audit. The source-paper language is now deliberately more modest: IO supplies the typed channel architecture and rate dressing; accepted A=7 nuclear response is imported as physical data.

The imported excited-state response enters through an electromagnetic transition-strength conversion rather than through a fitted lithium abundance. Paper 24 v2.3's Henderson primary row uses the measured B(E2) branch with the corrected de-excitation convention, projects it through the amplitude-weighted Coulomb kernel, and propagates the resulting excited-state response through the branch-resolved Li-7 calculation. The fact that the final Li-7/H remains in the observational band is important. The equally important statement is that the central value is not literally unchanged from the old derivation-based premise. The v2.3 paper says this explicitly.

This is an example of a source-data firewall. Henderson 2019 [37], evaluated nuclear lifetime data, mirror form-factor measurements, and R-matrix/AZURE-style discovery scaffolds can be used as primary experimental or discovery scaffolding where properly labelled. They cannot be used to claim IO has derived all A=7 matrix elements from the horizon. That mathematical bridge remains outside the current theorem stack.

The Modular Intertwiner Theorem (Paper 24 v2.3) and the thermal branch-selector status are labelled DERIVED/SCOPED in the relevant papers. The branch-resolved nuclear response inputs are not. That distinction should survive every retelling.

The neutron/proton mirror-ratio result is a separate, important bridge. Paper 24 reports the value already given in Eq. (2): $c\_n/c\_p = 1+\gamma\_BI^2 = 1.05641$, with a mismatch of about 2.35×10^{-4}, or 0.0235%, against the relevant nuclear ratio. This is one of the most striking numerical agreements in the corpus, but its boundary should be stated clearly: it is a forward bridge consistency result tied to the Barbero-Immirzi gauge twist, not a complete microscopic derivation of all A=7 nuclear structure.

A note on exhaustion: the A=7 class-search program is valuable precisely because it killed over-strong claims. Several candidate algebraic and numerical class structures (reduced generator-coordinate method, random algorithmic class, Schur-decomposition class, Berggren-basis proxy, R-matrix proxy, current-surrogate class) often moved the target surface, but randomized projector and orientation controls moved it more. The conclusion was not that nuclear physics is exhausted. The conclusion was that the raw aggregate scoring surface was too permissive, and that authoritative nuclear data can be used as a discovery scaffold only if it is quarantined from theorem-grade evidence. This is a methodological success even though it retired earlier confidence.

7. Major Results: Hubble Tension and Late-Time

Paper 34 v1.2 gives the current IO Hubble-tension account. The formula is

$$H\_ext(\alpha,n) = H\_0 \cdot f\_\Gamma^{(1-\alpha)} \cdot x^{(n/2)K\_gauge}, \quad (14)$$

In copy-safe plain-text form: H_ext(α,n) = H_0 × f_Γ^(1−α) × x^((n/2)K_gauge).

where H_0=67.5758565 km/s/Mpc is the active Paper 10 legacy projected-observer baseline, α is the observable-class rung, and n counts uncancelled photometric calibrator legs in the stellar-photometric extension.

The six-method scorecard is:

Method	IO prediction	Published value	Residual	Status
Planck CMB baseline	67.58	67.4±0.5	+0.35σ	DERIVED/SCOPED
GW sirens	67.58	69.9±4.1	-0.57σ	DERIVED/SCOPED
TRGB direct	70.26	70.39±1.94	-0.07σ	DERIVED/SCOPED within stellar-photometric extension
TDCOSMO	71.39	71.6±3.6	-0.06σ	DERIVED/SCOPED within centered lens-normalization extension
SH0ES	73.04	73.0±1.0	+0.04σ	DERIVED/SCOPED within stellar-photometric extension
TRGB+SN	73.04	73.18±0.88	-0.15σ	DERIVED/SCOPED within stellar-photometric extension

The maximum residual is 0.57 sigma. The SH0ES row is not fitted after the fact; Paper 34 v1.2 includes an anti-fit audit showing, for example, that GW sirens remain assigned to (α,n)=(1,0) even though another row would be numerically closer. Structural class membership overrides residual minimization. The κ-style audit found no hidden continuous fitted parameter in the H_ext formula, but it required explicit scope language. The result is derived inside two extension packages, not as a universal all-method H_0 theorem on the pre-extension stack.

The two extension packages are part of the claim. The stellar-photometric extension introduces a scalar radiative photometric carrier and counts uncancelled luminosity-calibration legs; this distinguishes TRGB direct from SH0ES and TRGB+SN. The centered lens-normalization extension fixes the mass-sheet and anisotropy selectors for TDCOSMO under boundary-state centering assumptions. Without those extensions, the corresponding method assignments would not be theorem-grade. With them, Paper 34 labels the assignments DERIVED/SCOPED and ships the audit that catalogues the premises.

A note on the phrase "Hubble tension resolved": the paper title uses that language. This overview uses more careful prose. IO gives a scoped observable-class account in which six method classes fall within 0.57 sigma of their assigned predictions with no continuous fitting. That is a strong claim. It is not a proof that every possible Hubble measurement, calibration pipeline, or future systematic belongs to the existing six-row table. New methods must be classified structurally before comparison.

Late-time closure is in Paper 32. The observer-side scalar optical branch has a finite domain boundary

$x_{crit}=Q^{-1/4}=0.98637546$. (15)

At that boundary, the scalar optical channel terminates. Below it, observer-side quantities such as the projected equation-of-state parameter are undefined, not merely unknown. This is a statement about the observer readout map, not about the disappearance of the hidden geometric support.

On the hidden support variable R, the active local acceleration equation is

$\ddot{R}=-(c^2 r_s)/(2R^2)$. (16)

At the support clamp $R=r_s$ this gives

$\ddot{R}(r_s)=-c^2/(2r_s)=-6.722\times 10^{-11}$ m s^{-2}. (17)

The cosmological constant drops out because the torsion-Λ term enters the velocity integral as an R-independent constant. Differentiating the velocity integral removes that constant from the acceleration equation. Source status: DERIVED on the physical hidden-support local OS branch, not on the observer-side Friedmann equation.

This is one of the most compact results in the corpus, and one of the most likely to be misread. Paper 32 is not claiming that the observed cosmological constant vanishes today. The observer-side Paper 10 legacy branch remains the correct current-epoch projected cosmology. The dropout happens in the future hidden-support local acceleration equation after the observer-side scalar optical channel terminates. Present observations from z=0 to recombination are not recomputed with Λ=0.

The recollapse timescale is

$\Delta\tau_{crunch}=\pi r_s/(2c)=110.99$ Gyr, (18)

with the symmetric full cycle

$\Delta\tau_{cycle}\simeq 222$ Gyr. (19)

The bounce-attachment statement is not the same as the recollapse derivation. Recollapse is DERIVED on the hidden-support branch under the support clamp. Bounce attachment is DERIVED/SCOPED within the extended fermionic/chiral IO-Einstein-Cartan theory of Paper 32 Part VIII and remains conditional on Paper 1's torsion-bounce premise. The hard-restart selectors β=1 and λ=1 enter as cycle-consistency conditions, not as independent derivations.

8. Reproducibility Infrastructure

The IO corpus ships paper-versioned reproducibility bundles for its load-bearing papers. Each bundle is archived as an immutable GitHub Release with a SHA256 checksum on the source tarball. The current bundles are:

- `paper17-v1.5` (modular projection / FIRAS-fixed readout): tarball SHA256 b807de6d5ef624a371dc09007049b40bee483e86d40272e0a01f81aa24bb8a33.
- `paper22-v1.6` (TBS rate-dressing): tarball SHA256 2e3ef109f25b342c1d3160f23b699191cf541ca56b221bce5d2a7d2a6ba1a06d.
- `paper24-v2.3` (BBN / lithium): tarball SHA256 81166b4332cc6814908c0bc46ee846cb7f27613571c2011b9e7f6619c2720651.
- `paper25-v1.3` (BBN weak-sector identity pin): tarball SHA256 1e3ee6cbacadb6c0820b8ff60bcc5a6022b9094103f714a943d6a64723003119.
- `paper32-v1.6` (closure): tarball SHA256 fcde39da0d96594715f62298bfdca6e034fc359a1697a268c6d53ed40152e03e.
- `paper34-v1.2` (Hubble tension): tarball SHA256 bfbb1f74a9f98d684e0e5f0d72a926dcdbb956e4fccaabc1dd124aab435d2597.
- `paper35-v1.2` (four problems): tarball SHA256 5201bd85c8addb0d3fe028ccea2e9b7ddaea24e183e1406519dafd2a5f01fadd.

Each bundle contains numbered scripts, frozen JSON outputs, environment notes, a manifest, and a validator. The standard validation invocation is `python3 papers/paperNN/vX.Y/scripts/NN_validate_expected_outputs.py`. A reviewer without PRyMordial installed can run the frozen-output validator to check arithmetic reproducibility against the manuscript values.

The reproducibility layer has three levels. Arithmetic reproducibility: do the scripts reproduce the stated numbers? Provenance reproducibility: do the scripts use the same constants, data files, conventions, and branch choices stated in the manuscript? Structural reproducibility: is the mathematical object being computed actually the object the prose claims? The public bundles address the first two. The κ audits and scope ledgers address the third, but they remain arguments to be reviewed, not machine-checkable facts.

The multi-AI pipeline is a reproducibility and error-discovery mechanism, not a credential claim. Papers 19, 20, 21, 22, and 25 exposed a real PRyMordial wrapper bug: older scripts read YPBBN where observational compilations require YPCMB. The correction sweep changed published scorecards and forced version bumps on those papers. Paper 24 was not affected by that wrapper bug; its v2.3 cluster-deformation sameness premise was independently retired by κ-style audit and replaced with primary-source imported electromagnetic response. Paper 34 and Paper 32 added scope/open-premise sections after field-redefinition audits. These are not cosmetic changes; they are examples of the pipeline finding and recording failure modes.

Running a bundle and seeing PASS does not prove IO is correct. It proves that the manuscript's numerical statements are reproducible from the shipped code and assumptions. That is the minimum standard for engagement, not the final standard for acceptance. The next level is to examine whether the assumptions are physically admissible and whether the claim labels are strong enough for the prose.

9. Open Problems and the Baryogenesis Frontier

IO is not closed in every sense. It is closed only on the scoped sectors claimed by the source papers. The main open problems are not afterthoughts; they identify where a future critic or collaborator should apply effort.

The baryon count and the baryogenesis mechanism are different problems. Paper 35 v1.2 derives the late-time baryon-to-photon ratio

$\eta_{late}=5.7475\times 10^{-10}$ (20)

with status DERIVED/SCOPED plus imported standard constants and a mean-baryon-mass convention. It also gives the BBN microphysical bulk-frame ratio

$\eta_{BBN}=6.151\times 10^{-10}$ (21)

as DERIVED/SCOPED under the temperature assignment theorem. The agreement between η_BBN and the Planck observer-extracted value is a cross-typed consistency observation, not a same-observable confrontation.

The source mechanism is not closed. The Paper 35 baryogenesis investigation audits 48 theorem surfaces: 15 clean and 33 Conditional_Verified. The obstruction reduces to a single topological datum, the $spin^c$ residual class n on the horizon U(1) bundle. On the current carrier package, n=0 and the CP-odd pseudoscalar Q_5 vanishes. Because the baryon-asymmetry source factor is multiplicative in Q_5, the admitted current carrier cannot generate visible baryon asymmetry. This is a no-go result, not a completed mechanism.

The result is stronger than “we do not know the mechanism” and weaker than “the mechanism is impossible.” It says the current carrier package is too neutral. The route dies for a specific topological reason, narrowing the future search to charged surface data or relative inflow data rather than allowing arbitrary baryogenesis model-building.

Paper 35 identifies two exact future routes. The surface route requires a symmetry-breaking surface state (ω_sb,q) with $q\in\mathbb{Z}_k\backslash 0$. The inflow route requires relative $spin^c$ data (A_rel,m_odd) on a bulk manifold Y with $\partial Y=S^2$. Neither is currently promoted to theorem-grade closure. This is the central baryogenesis frontier.

The chiral source-era diagnostic is explicitly labelled CONDITIONAL/CONSTRUCTED. It defines

$g_\chi=K_{gauge}^4=9.07\times 10^{-6}$, (22)

a diagnostic freeze-out scale $T_f\simeq 2.20\times 10^{13}$ GeV, and

$\eta_{chiral}=7.04\ K_{gauge}^8=5.787\times 10^{-10}$. (23)

The sub-percent agreement with the late-time route is interesting, but the freeze-out dynamics theorem, baryon-number transfer operator, and chirality/sign selector remain unclosed. It cannot be presented as an independent theorem-grade baryogenesis derivation.

The leptogenesis target reduction is DERIVED/CONDITIONAL on the standard external hierarchical thermal-leptogenesis class and an admitted Poplawski target slice. Paper 35 gives

$\varepsilon\_1\kappa\_f=5.904\times 10^{-8}$, (24)

with $M\_1\geq 5.99\times 10^{8}$ GeV for $\kappa\_f=1$, and $\kappa\_f\geq 2.72\times 10^{-5}$ if $M\_1=T\_f=2.20\times 10^{13}$ GeV. The chiral source-era diagnostic scale $T\_f = K\_gauge^4 M\_Pl = 2.20039\times 10^{13}$ GeV (Paper 35 v1.2) is unchanged under FIRAS-fixed R4 (uses powers of K_gauge and the diagnostic Planck scale, not optical R4); the Popławski η-target compatibility scale shifts slightly to 2.20016×10^{13} GeV under R4 repair. These are forced targets inside an external class. They are not internal IO source theorems. Standard hierarchical thermal leptogenesis still has positive-dimensional source freedom: the same η can be produced by multiple source parameter choices. Paper 35's source-manifold no-go records this rather than pretending η alone selects a unique mechanism.

The JWST result is timing-pressure relief, not full galaxy-formation closure. Paper 35 reports that the IO formation clock gives about 46-48% more time than ΛCDM at z=10 to 16: for example 0.695 Gyr versus 0.470 Gyr at z=10, and 0.433 Gyr versus 0.295 Gyr at z=14. This eases the clock-time pressure for early massive galaxies. It does not solve dust, feedback, seed black-hole, or stellar-population modeling. The result is useful if the dominant tension is clock time. It is insufficient if the dominant tension is astrophysical assembly physics.

The dark-matter null is a CONDITIONAL FORECAST scoped to the IO geometric dark-sector interpretation. The ledger quantity is

$f\_b=2\gamma\_BI/x=0.3127$. (25)

If the interpretation is correct, continued nulls in canonical xenon WIMP searches are expected. A robust direct-detection signal above roughly 5 GeV in the cited channels would falsify the IO geometric dark-sector interpretation. The current theorem stack does not prove an unconditional no-particle theorem; Paper 33 records that a primitive typed dark-sector carrier remains conditionally open. "Dark matter is geometry" is a possible interpretation of the IO inventory ledger, not yet a theorem that no local dark-sector carrier can exist.

The bounce/restart package remains scoped. Paper 32 derives local recollapse on the hidden-support branch. The torsion bounce is inherited from Paper 1/Poplawski-type Einstein-Cartan physics and implemented in the extended fermionic/chiral IO-Einstein-Cartan theory. Observer reactivation is proved possible, but full dictionary return depends on selector principles. The bounce is not derived in the same unconditional sense as the local recollapse equation.

The full CMB C_ℓ spectrum remains open. IO has source-block n_s, scalar amplitude, acoustic-scale, optical-depth, and reionization-import machinery. It does not yet ship a full

IO-native Boltzmann transfer code that reproduces the complete angular power spectrum from first principles. A future full-spectrum calculation would need an IO-native primordial source, a recombination history, perturbation transfer equations on the active closed branch, lensing and reionization treatment, and likelihood comparison against Planck or successor data. The current papers provide pieces of that pipeline, not the complete object.

Universal GMP beyond realized typed bridge classes is characterized, not globally closed. Paper 32's universal GMP theorem is a three-part classification: bridge sector closed, non-bridge complement inapplicable, and abstract bypass operators no-go. The word "inapplicable" is a structural classification inside the typed architecture, not an excuse to avoid observables.

A note on premise inheritance: IO Framework Conventions v2.0 requires conditional dependency inheritance. If a theorem depends on another theorem that is conditional on GMP + TBS, the downstream theorem inherits that conditionality even if its local statement looks algebraic. Papers 22 and 25 showed how easy it is for conditional status to disappear in cross-paper prose. The overview follows the stricter convention.

The most productive external work would likely fall into four categories. Independent reproduction: rerun the public bundles, replace the observational denominators with alternate accepted compilations, and see how much the scorecards move. Theorem audit: pick one scoped result such as H_ext, n_s, or the lithium branch map and run an independent field-redefinition test. Physics replacement: build a better A=7 microscopic response calculation, a better IO-native CMB transfer pipeline, or a better baryogenesis carrier and compare it against the current frontier. Falsification: look for measurements that IO classifies sharply, such as geometric-standard-siren H_0, future xenon/dark-matter nulls, or high-redshift formation clocks.

The framework is only scientifically useful if it tells other researchers where it can be broken. This section is part of the pitch, not a caveat after the fact.

10. Reading Paths into the Corpus

If you work on Einstein-Cartan torsion, start with Paper 1, Paper 32 Part VIII, and Paper 35. Paper 1 establishes the torsion-bounce premise and black-hole interior setting. Paper 32 shows how the fermionic/chiral bridge attaches the recollapse branch to the bounce regime, with its scope stated explicitly. Paper 35 tests how far that architecture goes in baryogenesis and identifies the remaining spin^c obstruction. Paper 1 contains founding assumptions that are later narrowed; Paper 32 is the necessary update because it separates derived recollapse from conditional bounce attachment.

If you work on loop quantum gravity or the Barbero-Immirzi parameter, read Papers 17, 24, 27, and 30. Paper 17 is the modular projection theorem behind $K\_gauge=\ln(1+\gamma\_BI^2)$. Paper 24 contains the $c\_n/c\_p=1+\gamma\_BI^2$ nuclear bridge and the branch-resolved lithium calculation. Papers 27 and 30 place the same constant into optical-depth, baryon-density, and structure-sector contexts. The interesting LQG question is whether the same imported horizon parameter appears in unrelated observables with the correct typed powers and signs. If those appearances reduce to hidden fitting, the framework fails. If they are structurally forced, the horizon gauge sector becomes physically consequential far below the Planck scale.

If you work on BBN or nuclear cosmology, read Papers 21, 22, 24, and 25 in version order. Paper 21 records the older response-map hunt and its supersession. Paper 22 builds the rate-dressing amplitude construction, now labelled DERIVED/CONDITIONAL on GMP + TBS. Paper 25 pins the weak-sector V-vs-V' identity and quadratic thermal covariance. Paper 24 v2.3 is the active branch-resolved lithium result using imported primary A=7 response. A reader who skips the version history will misunderstand why some earlier attractive numbers are no longer headline claims.

If you work on the Hubble tension, read Papers 29, 31, and 34. Paper 29 establishes the active projected-observer baseline branch. Paper 31 supplies the f_Γ and scalar/lensing readout context. Paper 34 v1.2 classifies six H_0 methods across three rungs and provides the κ-audited public reproducibility bundle. The main question is anti-fit. Do not begin by asking whether 73.04 is close to SH0ES. Ask whether Cepheid+SN must be (α,n)=(2,2) before the SH0ES value is examined. Ask whether TDCOSMO's mass-sheet and anisotropy selectors are truly fixed inside the centered lens-normalization extension. Ask why GW sirens remain at the low baseline despite a closer numerical row.

If you work on JWST high-redshift timing, read Papers 28 and 35. Paper 28 supplies the homogeneous OS formation-clock theorem. Paper 35 applies the clock to z=10 to 16 objects and states the result as timing-pressure relief rather than full galaxy-formation closure. The IO clock gives more time; it does not compute star-formation efficiency, feedback, dust, metal enrichment, or seed black-hole growth.

If you want the closure argument, read Paper 32 v1.6 first. It is the framework-closing paper: fixed-point normalization, late-time chain, recollapse, universal GMP classification, open-premise ledger, and public reproducibility bundle. Paper 32 is also the place to audit the phrase "zero fitted parameters." Its κ-style audit covers the recollapse acceleration, x_{crit}, the 111/222 Gyr cycle, Theorem 32.KB.7, the universal GMP classification, the founding-premise status of Premise 2, the retirement of the four semiclassical principles, and the closure of the original open problems. The audit found no hidden continuous fitted parameter, but it required scope and open-premise hygiene.

If you want the GR-QM bridge architecture, read Paper 33. It organizes the corpus into static geometric, dynamic thermal, and fermionic/density-gated export languages, while preserving the warning that the full theory is a typed lower-triangular carrier, not a slogan.

If you want the most recent state, read Paper 35 v1.2. It is not the closure paper; it is the frontier paper. It updates cross-paper BBN values, documents the baryogenesis obstruction, applies the IO clock to JWST, corrects stale DESI/dark-energy framing, and labels the dark-matter null as conditional forecast.

For all reading paths, use the current version numbers. Earlier versions remain citable historical artifacts, but active claims should be taken from the latest corrected paper and, where available, the immutable reproducibility bundle. This is especially important for Papers 19, 20, 21, 22, and 25, where the YPBBN-to-YPCMB wrapper correction changed scorecard values; for Paper 24 v2.3, where the Henderson-primary row alignment defines the current lithium scorecard; and for Papers 32, 34, and 35, where public bundles and κ audits now define the cleanest entry points.

Code and Data Availability

The public reproducibility repository is `https://github.com/dfife/io-framework-public`. It contains versioned bundles for the current public-bundle papers cited in §8, including Papers 17, 21, 22, 24, 25, 32, 34, and 35 where applicable. Each bundle has a README, manifest, environment notes, scripts, frozen JSON outputs, reports, and a validator where applicable.

The live framework website is `https://dfife.github.io/`. The Zenodo community for the paper corpus is `https://zenodo.org/communities/interior-observer/`.

This overview introduces no new computation. All numbers quoted here are taken from the active source papers and their public bundle outputs where available.

References

[1] Fife, D. *Interior Observer Paper 1 v3.5: A Black Hole Cosmology with the CMB as Event Horizon Radiation*. Zenodo. https://doi.org/10.5281/zenodo.18854813

[2] Fife, D. *Interior Observer Paper 10 v2.0: The Complete Holographic Friedmann Equation*. Zenodo. https://doi.org/10.5281/zenodo.18905811

[3] Fife, D. *Interior Observer Paper 12: Static-Geometric Bridges and the Rosetta Identity*. Zenodo. https://doi.org/10.5281/zenodo.18936507

[4] Fife, D. *Interior Observer Paper 13 v1.7: The Gauge Thermal Transfer Principle*. Zenodo. https://doi.org/10.5281/zenodo.18945480

[5] Fife, D. *Interior Observer Paper 17 v1.5: The Modular Projection Theorem (with FIRAS-Fixed Readout Normalization Theorem)*. Zenodo. https://doi.org/10.5281/zenodo.19045892

[6] Fife, D. *Interior Observer Paper 18: Modular and Observable-Class Upgrade*. Zenodo. https://doi.org/10.5281/zenodo.19054257

[7] Fife, D. *Interior Observer Paper 19: Rosetta Formalization, Baryon Scalarization, and the α-Ladder*. Zenodo. https://doi.org/10.5281/zenodo.19101802

[8] Fife, D. *Interior Observer Paper 20: Acoustic-Scale Construction and Reduced Sound-Horizon Operator*. Zenodo. https://doi.org/10.5281/zenodo.19145121

[9] Fife, D. *Interior Observer Paper 21 v1.7*. Zenodo. https://doi.org/10.5281/zenodo.19173253

[10] Fife, D. *Interior Observer Paper 22 v1.6*. Zenodo. https://doi.org/10.5281/zenodo.19198706

[11] Fife, D. *Interior Observer Paper 23 v1.3*. Zenodo. https://doi.org/10.5281/zenodo.19210262

[12] Fife, D. *Interior Observer Paper 24 v2.3: The Lithium Problem Solved: Channel-Resolved BBN, Quadrupole Isolation, the Modular Intertwiner Theorem, and a Conditional Bridge from Quantum Gravity to Nuclear Isospin*. Zenodo. https://doi.org/10.5281/zenodo.19219282

[13] Fife, D. *Interior Observer Paper 25 v1.3: BBN Weak-Sector Identity Pin and the Quadratic Thermal Covariance Theorem*. Zenodo. https://doi.org/10.5281/zenodo.19242682

[14] Fife, D. *Interior Observer Paper 27 v1.5*. Zenodo. https://doi.org/10.5281/zenodo.19333220

[15] Fife, D. *Interior Observer Paper 28 v1.3*. Zenodo. https://doi.org/10.5281/zenodo.19359276

[16] Fife, D. *Interior Observer Paper 29 v1.2*. Zenodo. https://doi.org/10.5281/zenodo.19367583

[17] Fife, D. *Interior Observer Paper 30 v1.3*. Zenodo. https://doi.org/10.5281/zenodo.19376058

[18] Fife, D. *Interior Observer Paper 31 v1.2*. Zenodo. https://doi.org/10.5281/zenodo.19415916

[19] Fife, D. *Interior Observer Paper 32 v1.6: The Cosmological Constant Drops Out: A Closed-Universe Cosmology with Zero Fitted Parameters*. Zenodo. https://doi.org/10.5281/zenodo.19433686

[20] Fife, D. *Interior Observer Paper 33 v1.1: Three Languages of the Horizon: How an Event Horizon Unifies General Relativity and Quantum Mechanics*. Zenodo. https://doi.org/10.5281/zenodo.19440227

[21] Fife, D. *Interior Observer Paper 34 v1.2: The Hubble Tension Resolved: Six Methods, Three Rungs, Zero Fitted Parameters*. Zenodo. https://doi.org/10.5281/zenodo.19558163

[22] Fife, D. *Interior Observer Paper 35 v1.2: Four Problems, One Horizon*. Zenodo. https://doi.org/10.5281/zenodo.19561708

[23] Fife, D. *IO Framework Observational Conventions v2.0*. Methodological authority for IO claim labels, BBN observational denominators, and conditional dependency inheritance. Cited throughout the v2.x corpus. https://dfife.github.io/data/conventions_v2.md

[24] Fife, D. *IO Framework public reproducibility repository*. GitHub. https://github.com/dfife/io-framework-public

[25] Fife, D. *Paper 24 v2.3 reproducibility bundle*. GitHub Release `paper24-v2.3`. https://github.com/dfife/io-framework-public/releases/tag/paper24-v2.3

[26] Fife, D. *Paper 32 v1.6 reproducibility bundle*. GitHub Release `paper32-v1.6`. https://github.com/dfife/io-framework-public/releases/tag/paper32-v1.6

[27] Fife, D. *Paper 34 v1.2 reproducibility bundle*. GitHub Release `paper34-v1.2`. https://github.com/dfife/io-framework-public/releases/tag/paper34-v1.2

[28] Fife, D. *Paper 35 v1.2 reproducibility bundle*. GitHub Release `paper35-v1.2`. https://github.com/dfife/io-framework-public/releases/tag/paper35-v1.2

[29] Pathria, R. K. "The Universe as a Black Hole." *Nature* 240, 298-299 (1972). https://doi.org/10.1038/240298a0

[30] Smolin, L. "Did the Universe Evolve?" *Classical and Quantum Gravity* 9, 173-191 (1992).

[31] Stuckey, W. M. "The Schwarzschild Black Hole as a Gravitational Mirror." *American Journal of Physics* 62, 788-795 (1994).

[32] Poplawski, N. J. "Cosmology with torsion: An alternative to cosmic inflation." *Physics Letters B* 694, 181-185 (2010).

[33] Fixsen, D. J. "The Temperature of the Cosmic Microwave Background." *The Astrophysical Journal* 707, 916-920 (2009). https://doi.org/10.1088/0004-637X/707/2/916

[34] Planck Collaboration. "Planck 2018 results. VI. Cosmological parameters." *Astronomy & Astrophysics* 641, A6 (2020). https://doi.org/10.1051/0004-6361/201833910

[35] Cooke, R. J., Pettini, M., and Steidel, C. C. "One Percent Determination of the Primordial Deuterium Abundance." *The Astrophysical Journal* 855, 102 (2018). https://doi.org/10.3847/1538-4357/aaab53

[36] Henderson, S. et al. "First measurement of the B(E2;3/2^-→ 1/2^-) transition strength in ^{7}Be: Testing ab initio predictions for A=7 nuclei." *Physical Review C* 99, 064320 (2019). https://doi.org/10.1103/PhysRevC.99.064320

[37] Riess, A. G. et al. "A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km/s/Mpc Uncertainty from the Hubble Space Telescope and the SH0ES Team." *The Astrophysical Journal Letters* 934, L7 (2022).

[38] Freedman, W. L. et al. "Status Report on the Chicago-Carnegie Hubble Program." *The Astrophysical Journal* (2024). https://doi.org/10.3847/1538-4357/adce78

[39] DESI Collaboration. DESI DR2 BAO and cosmological constraints data products and papers, cited in Paper 35 v1.2 bundle.

[40] LZ Collaboration. "Dark matter search results from 4.2 tonne-years of exposure of the LUX-ZEPLIN (LZ) experiment." arXiv:2410.17036.

[41] XENON Collaboration. "First dark matter search with nuclear recoils from the XENONnT experiment." *Physical Review Letters* 131, 041003 (2023). arXiv:2303.14729.

[42] PandaX Collaboration. "Dark matter search results from 1.54 tonne-year exposure of PandaX-4T." arXiv:2408.00664.

[43] Carniani, S. et al. "Spectroscopic confirmation of a galaxy at redshift z=14.32." *Nature* 633, 318-322 (2024).

[44] Curtis-Lake, E. et al. "Spectroscopic confirmation of four metal-poor galaxies at z=10.3-13.2." *Nature Astronomy* 7, 622-632 (2023).

[45] Bunker, A. J. et al. "JADES NIRSpec spectroscopy of GN-z11." *Astronomy & Astrophysics* 677, A88 (2023).

[46] Sakharov, A. D. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe." *JETP Letters* 5, 24-27 (1967).

[47] Davidson, S. and Ibarra, A. "A lower bound on the right-handed neutrino mass from leptogenesis." *Physics Letters B* 535, 25-32 (2002).

[48] Buchmuller, W., Peccei, R. D., and Yanagida, T. "Leptogenesis as the origin of matter." *Annual Review of Nuclear and Particle Science* 55, 311-355 (2005).