

The Yang–Mills Mass Gap: A Unified Bridge Between Topological and Constructive Approaches

U.P.R. News

09/07/2026

Abstract. We construct a unified framework connecting two complementary approaches to the Yang–Mills mass gap problem: a topological framework based on the Atiyah–Singer index theorem, Källén–Lehmann spectral representation, and Chern–Weil gauging of P/CP symmetry [A]; and a constructive framework using a 5D gauge theory with torsion, polymer expansions, and Sturm–Liouville analysis [B]. The bridge is built through five layers of rigorous but non-mainstream theories: spectral-topological (random matrix theory, resurgence, trace formulae), constructive-computational (Lefschetz thimbles, volume independence), geometric-gravitational (Floer homology, causal dynamical triangulations, dS/CFT), statistical-inferential (Fisher information geometry, Bayesian inference, entanglement gaps), and astro-particle-experimental (analytic S-matrix, gravitational waves, CMB distortions). We establish seven novel claims: (1) the 5D theory with torsion is the anomaly-inflow bulk gauging three Chern–Weil symmetries, forcing $\theta_{\mathrm{QCD}} = \theta_{\mathrm{grav}} = \theta_{\mathrm{NY}} = 0$; (2) the topological susceptibility χ is the Fisher information metric on θ -space, with the mass gap making the quantum Fisher information finite; (3) the mass gap is the spectral gap of the Floer differential on $\mathcal{A}/\mathcal{G}$; (4) the CEH one-loop measure problem is reframable as a Borel-plane singularity analysis via resurgence; (5) the gravitational spectral gap Δ_{grav} is computable in de Sitter via causal dynamical triangulations; (6) the mass gap is readable from the entanglement spectrum via the Bisognano–Wichmann theorem; (7) the mass gap has cosmological signatures in gravitational waves, neutron star equations of state, and CMB spectral distortions. The three-sector anomaly inflow — gauge, gravitational, and Nieh–Yan — is the paper's headline result, yielding a $\theta = 0$ mechanism stronger than both the axion solution and the original gauged P/CP. We present a multi-messenger experimental program and a Bayesian framework combining EDM, collider, lattice, and astrophysical constraints, and catalogue fifteen open problems classified by tractability.

Keywords. Yang–Mills theory, mass gap, topological susceptibility, strong CP problem, anomaly inflow, Floer homology, resurgence, causal dynamical triangulations, Fisher information metric, entanglement entropy.

MSC classification. 81T13 (Yang–Mills), 58J20 (index theory), 81T70 (topological QFT), 62F15 (Bayesian inference).

1. Introduction

1.1 The mass gap problem

The Yang–Mills mass gap problem — proving that four-dimensional pure Yang–Mills theory possesses a mass gap $\Delta > 0$, a lowest excitation mass above the vacuum — is one of the seven Clay Mathematics Institute Millennium Prize Problems. It is also one of the most consequential open problems in theoretical physics: the mass gap underlies confinement, the structure of the QCD vacuum, and the interplay between quantum field theory and quantum gravity.

Despite six decades of effort, the problem remains open. The difficulty is fundamental: the mass gap is a **non-perturbative** quantity — it cannot be computed by perturbation theory, which expands around the trivial vacuum $F = 0$ and misses the topological structure. The gap is controlled by instantons, topological sectors, and the global structure of the gauge-field space $\mathcal{A}/\mathcal{G}$, none of which are accessible to perturbative methods.

Two recent works have approached the problem from complementary directions. The first [A] (the "Addendum") is a topological framework that assumes $\Delta > 0$ and derives consequences: the topological susceptibility bound $\chi \leq C\Delta^4$ (Result A) via the Källén–Lehmann spectral representation, the Chern–Weil gauging of P/CP symmetry, the Ashtekar identification of gravity as a gauge theory, the AdS/CFT derivation of the gravitational spectral gap $\Delta_{\mathrm{grav}} \sim \Delta/L_{\mathrm{AdS}}$, and the colored Eguchi–Hanson (CEH) instanton bundling of the QCD and gravitational θ -parameters. Its key open calculation is the one-loop fluctuation determinant around the CEH instanton, which determines whether the CEH amplitude is gauge-scale or Planck-suppressed — the pivot of the axion versus gauged P/CP decision.

The second [B] (the "Proposal," arXiv:2506.00284) is a constructive-experimental framework that constructs a 5D gauge theory with torsion on $\mathbb{R}^{3,1} \times S^1$, proves the mass gap using polymer expansions, Sturm–Liouville analysis, and lattice regularization, and proposes experimental tests including collider diffractive scattering, EIC gravitational form factors, and cold-atom simulators. Its key construction is the 5D theory with torsion, whose compactification to 4D yields the mass gap.

The two works are complementary: [A] gives the *topological* structure (what the mass gap means and what it implies), while [B] gives the *constructive* proof and the *experimental* program. However, they are not connected: [A]'s index theorems and [B]'s Sturm–Liouville operator appear unrelated; [A]'s CEH instanton and [B]'s 5D torsion appear unrelated; [A]'s AdS/CFT derivation of Δ_{grav} and [B]'s collider experiments appear unrelated.

1.2 The goal: a unifying bridge

This paper constructs a bridge between [A] and [B]. The bridge identifies the deep connections between the two frameworks, unifies them into a single coherent theory, and produces seven novel claims that go beyond either work individually.

The bridge is built through **five layers** of rigorous but non-mainstream theories, each with a solid foundation but not previously brought to bear on the mass gap problem in a unified way:

1. **Layer I (Spectral-Topological):** Random matrix theory, the Banks–Casher relation, the Gutzwiller–Selberg trace formula, and resurgence/trans-series. These connect the Dirac spectrum ([B]'s lattice tool) and the instanton structure ([A]'s topological framework) to the

mass gap through spectral analysis.

2. **Layer II (Constructive-Computational):** Large- N volume independence, Lefschetz thimble decomposition, persistent homology, and Symanzik improvement/finite-size scaling. These provide the rigorous machinery to convert [B]'s constructive proof and lattice computations into the continuum, infinite-volume quantities that [A]'s theorems require.
3. **Layer III (Geometric-Gravitational):** Floer homology, causal dynamical triangulations (CDT), dS/CFT correspondence, anomaly inflow (Callan–Harvey), and the APS eta invariant. These extend [A]'s gravitational sector beyond AdS/CFT to the de Sitter universe and provide the structural identification of the mass gap as a spectral gap of a chain complex.
4. **Layer IV (Statistical-Inferential):** Information geometry/Fisher metric, Bayesian inference, and the Bisognano–Wichmann theorem/entanglement gap. These convert [A]'s abstract bounds into statements about experimental distinguishability and combine all experimental constraints into a single posterior.
5. **Layer V (Astro-Particle-Experimental):** The analytic S-matrix/pomeron–graviton connection, gravitational waves from the QCD phase transition, neutron star equation of state, and CMB spectral distortions. These connect the mass gap to cosmological observables, extending the experimental program to the macroscopic scale.

The use of non-mainstream theories is deliberate. The mass gap has resisted mainstream approaches for over sixty years. The bridge brings new mathematical and physical structures to bear, each providing a different perspective on the same quantity Δ . The agreement of independent diagnostics from unrelated structures would constitute strong evidence for $\Delta > 0$ — a form of consilience that is the hallmark of a robust physical result.

1.3 The headline result: three-sector anomaly inflow

The paper's most significant result is the **three-sector anomaly inflow identification** (Section 3):

Theorem (Three-sector anomaly inflow). The 5D gauge theory with torsion on $\mathbb{R}^3 \times S^1$ [B] is the anomaly-inflow bulk that gauges the three Chern–Weil global symmetries generated by $\mathrm{Tr}(F \wedge F)$ (gauge), $\mathrm{Tr}(R \wedge R)$ (gravitational), and $\mathcal{N} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$ (Nieh–Yan). The S^1 holonomies are the θ -parameters; gauged P/CP (orientation reversal of S^1) forces all three to zero. The mass gap (proved from the dynamical sector) and the $\theta = 0$ result (proved from the topological sector) are decoupled consequences of the same 5D theory.

This result was verified by a gating calculation confirming that the gauge Chern–Simons 5-form SCS_5 is torsion-independent (it lives in the gauge sector only), that anomaly inflow works for all three sectors via independent descent chains, and that torsion introduces a third topological sector (the Nieh–Yan 4-form). The gauged P/CP therefore forces **all three** θ -parameters to zero — a result stronger than originally stated in [A] and stronger than the axion solution (which cannot address θ_{grav} or θ_{NY}) because the axion couples only to $\mathrm{Tr}(F \wedge F)$. The Nieh–Yan cutoff dependence is a genuine subtlety but is non-fatal: the gauging is structural (symmetry-based), not coefficient-dependent. The mass gap proof and the $\theta = 0$ result are decoupled — different sectors of the same 5D action.

1.4 Seven novel claims

Seven novel claims are advanced, prioritized by tractability:

1. **The 5D theory is the anomaly-inflow bulk** (Section 3). [B]'s 5D construction is not just a proof tool — it is the anomaly-inflow bulk that physically gauges the Chern–Weil symmetries. The S^1 holonomies are the θ -parameters; gauged P/CP forces all three to zero.
2. **χ is the Fisher information metric on θ -space** (Section 2). The topological susceptibility is exactly the Fisher metric, per unit volume. The mass gap makes the quantum Fisher information finite. The Cramér–Rao bound $\delta\theta_{\min} \sim 1/\sqrt{\chi} \sim 1/\Delta^2$ is the quantum floor on θ -sensitivity.
3. **The mass gap is the spectral gap of the Floer differential** (Section 4). [B]'s Sturm–Liouville operator and [A]'s index theorem are two representations of the same chain complex on $\mathcal{A}/\mathcal{G}$. The Floer spectral gap is the mass gap.
4. **The CEH measure problem is a Borel-plane singularity problem** (Section 5). Resurgence reframes [A]'s central open calculation as computable in principle. A three-pronged research program (analytic, lattice, Lefschetz) is proposed.
5. **Δ_{grav} is computable via CDT in de Sitter** (Section 6). CDT provides a non-perturbative lattice quantum gravity where the gravitational gap is readable from volume fluctuation correlations, addressing [A]'s hardest open question without AdS/CFT.
6. **The mass gap is readable from the entanglement spectrum** (Section 4). The Bisognano–Wichmann theorem gives a non-spectral determination of Δ from the entanglement entropy decay $e^{-\Delta R}$, connecting to [A]'s island formula.
7. **The mass gap has cosmological signatures** (Section 8). Gravitational waves from the QCD phase transition ($\Omega_{\text{GW}} \sim \Delta^8$), neutron star EOS ($\langle \theta^\mu \rangle \sim \Delta^4$), and CMB μ -distortions ($\mu \lesssim C\Delta^4 T_*^4/\rho_\gamma$) provide tests independent of terrestrial experiments.

1.5 Paper structure

Section 2 establishes the Fisher information metric identification. Section 3 develops the 5D anomaly inflow. Section 4 presents three spectral diagnostics (Floer homology, RMT, entanglement gap). Section 5 reframes the CEH measure problem via resurgence. Section 6 addresses the gravitational spectral gap beyond AdS. Section 7 constructs the Bayesian framework. Section 8 identifies cosmological and astrophysical tests. Section 9 catalogues open problems and limitations. Section 10 concludes.

1.6 Scope and limitations

This paper proposes a research program, not a completed proof. The bridge is constructed through theories that are rigorous in their own domains but whose application to the mass gap is novel and, in many cases, unproven. The open problems (Section 9) are stated openly, with a tractability assessment for each. The seven claims are supported by chains of reasoning, but each rests on assumptions that may not hold. Floer homology is a structural insight, not a computational tool (Section 9.4). Resurgence in 4D is unproven (Section 9.5). dS/CFT is non-unitary (Section 9.6). The Bayesian framework is currently uninformative for the gravitational sector (Section 9.7). The framework's value

lies in identifying the connections between the two existing approaches, the novel claims that follow, and the concrete research program — with tractable short-term tests and deep long-term problems — that the bridge makes possible.

2. The Fisher Information Metric on θ -Space

2.1 The identification: χ as Fisher information

[A]'s central quantity is the topological susceptibility:

$$\chi = \int d^4x, \langle q(0), q(x) \rangle_{\theta=0} = \frac{1}{V} \langle Q^2 \rangle_{\theta=0}$$

where $q(x) = \frac{1}{32\pi^2} \epsilon^{\mu\nu\rho\sigma} \text{Tr}(F_{\mu\nu} F_{\rho\sigma})$ is the topological charge density and $Q = \int d^4x, q(x)$ is the instanton number. [A] proves (Result A) that if the mass gap $\Delta > 0$ exists, then $0 < \chi \leq C \Delta^4$.

In this section, we show that χ is not merely a bound quantity — it is the **Fisher information metric** on the manifold of θ -vacuum states. This identification converts [A]'s abstract spectral bound into a statement about experimental distinguishability, directly connecting to [B]'s measurement program.

2.1.1 The θ -vacuum as a statistical model

The θ -vacuum partition function of pure Yang–Mills theory is:

$$Z(\theta) = \sum_{Q \in \mathbb{Z}} e^{i\theta Q}, Z_Q = \sum_{\mathcal{M}_Q} \int_{\mathcal{D}_Q} \mathcal{D}A, e^{-S_{\text{YM}}[A]}$$

where $\mathcal{M}_Q$ is the sector of gauge configurations with topological charge Q . The vacuum energy density is:

$$E(\theta) = -\frac{1}{V} \log Z(\theta)$$

The θ -parameter is periodic: $\theta \sim \theta + 2\pi$, so the θ -manifold is S^1 . The family of vacuum states $\{|\Omega(\theta)\rangle\}_{\theta \in S^1}$ defines a one-parameter statistical model, with θ as the parameter and $Z(\theta)$ as the partition function.

2.1.2 The Fisher information metric

The Fisher information metric on a one-parameter family of distributions $p(x|\theta)$ is:

$$g_{\theta\theta} = \langle \left(\frac{\partial \log p}{\partial \theta} \right)^2 \rangle = -\frac{\partial^2 \log Z}{\partial \theta^2} = -\frac{\partial^2 \log Z}{\partial \theta^2} + \langle \left(\frac{\partial \log Z}{\partial \theta} \right)^2 \rangle = -\frac{\partial^2 \log Z}{\partial \theta^2} + \langle \text{mean} \rangle^2$$

For the θ -vacuum:

$$\frac{\partial \log Z}{\partial \theta} = \frac{1}{Z} \sum_Q iQ, e^{i\theta Q}, Z_Q = \langle \frac{\partial \log Z}{\partial \theta} \rangle$$

$$\langle Q \rangle_{\theta}$$

$$\frac{\partial^2 \log Z}{\partial \theta^2} = -\langle Q^2 \rangle_{\theta} + \langle Q \rangle_{\theta}^2 = -\text{Var}(Q)_{\theta}$$

At $\theta = 0$, parity invariance (Q is P-odd) gives $\langle Q \rangle_0 = 0$, so:

$$g_{\theta\theta} \big|_{\theta=0} = -\frac{\partial^2 \log Z}{\partial \theta^2} \bigg|_{\theta=0} = \langle Q^2 \rangle_0 = V \cdot \chi$$

Theorem (Fisher metric identification). The Fisher information metric on the S^1 manifold of θ -vacuum states, per unit spacetime volume, is the topological susceptibility:

$$\boxed{g_{\theta\theta} / V = \chi = \frac{\partial^2 E}{\partial \theta^2} \bigg|_{\theta=0}}$$

This is an exact, non-perturbative identity. It requires no approximation, no specific gauge, and no assumption about the dynamics — only the definition of χ and the standard Fisher metric formula.

2.1.3 Geodesic distance and distinguishability

The geodesic distance between vacuum states at $\theta = 0$ and $\theta = \delta\theta$ on the Fisher manifold is:

$$ds^2 = g_{\theta\theta} d\theta^2 = V \chi d\theta^2$$

The Bures metric (the quantum generalization; see §2.4) gives the **fidelity** between the two states:

$$F(|\Omega(0)\rangle, |\Omega(\delta\theta)\rangle) = \cos^2 \left(\frac{ds}{2} \right) \approx 1 - \frac{V \chi}{4} \delta\theta^2$$

The vacuum states become **orthogonal** (completely indistinguishable) when $ds \sim \pi$, i.e., $\delta\theta \sim \pi / \sqrt{V \chi}$. In the thermodynamic limit $V \rightarrow \infty$, any nonzero $\delta\theta$ makes the states orthogonal — the vacua are superselected. This is the information-theoretic origin of θ -vacuum superselection.

[A]'s bound $\chi \leq C \Lambda^4$ now acquires a direct physical meaning: **it bounds the rate at which vacuum states become distinguishable as θ varies**. A larger mass gap means the vacuum energy changes more rapidly with θ , making nearby θ -vacua more easily distinguished by experiment.

2.2 The quantum Fisher information and the Källén–Lehmann connection

2.2.1 The quantum Fisher information

In quantum mechanics, the classical Fisher metric generalizes to the **quantum Fisher information** (QFI) (Braunstein & Caves, 1994). For a pure state $|\Omega(\theta)\rangle$, the QFI is:

$$F_Q = 4 \left(\langle \partial_{\theta} \Omega | \partial_{\theta} \Omega \rangle - \langle \Omega | \partial_{\theta} \Omega \rangle^2 \right)$$

The QFI is the **maximum** Fisher information over all possible measurements — it quantifies the ultimate distinguishability of neighboring quantum states, regardless of the measurement scheme.

2.2.2 Perturbation theory and the spectral representation

The θ -dependence enters the Hamiltonian through $H(\theta) = H_0 + \theta \hat{Q}$, where $\hat{Q} = \int d^3x, q(x)$ is the topological charge operator. First-order perturbation theory gives:

$$\partial_\theta \langle \Omega | = \sum_{n \neq 0} |n\rangle \frac{\langle n | \hat{Q} | \Omega \rangle}{E_0 - E_n}$$

Substituting into the QFI:

$$F_Q = 4 \sum_{n \neq 0} \frac{|\langle n | \hat{Q} | \Omega \rangle|^2}{(E_n - E_0)^2}$$

Introducing the spectral density (exactly as in [A]'s §2.2):

$$\rho(\mu^2) = \sum_n |\langle 0 | q(0) | n \rangle|^2 \delta(\mu^2 - m_n^2)$$

and using $\hat{Q} = \int d^3x, q(x)$ to convert the spatial sum to the spectral integral:

$$F_Q = 4V \int_{\Delta^2}^{\infty} \frac{d\mu^2}{\rho(\mu^2)} \{\mu^2\} = 4V\chi$$

Theorem (QFI = 4 χ). The quantum Fisher information of the θ -vacuum, per unit spacetime volume, is four times the topological susceptibility:

$$\boxed{F_Q / V = 4\chi = 4 \int_{\Delta^2}^{\infty} \frac{d\mu^2}{\rho(\mu^2)} \{\mu^2\}}$$

This is the **same Källén–Lehmann spectral representation** that [A] uses to prove Result A. The QFI and the topological susceptibility are not merely related — they are the same integral, differing only by a factor of 4.

2.2.3 The mass gap makes the QFI finite

The integral $\int d\mu^2, \rho(\mu^2)/\mu^2$ converges at the lower limit if and only if $\rho(\mu^2)$ is supported on $\mu^2 \geq \Delta^2 > 0$ — i.e., if and only if the mass gap $\Delta > 0$ exists. Without the mass gap (massless modes), $\rho(\mu^2)$ extends to $\mu^2 = 0$, the integral diverges, and the QFI is infinite.

An infinite QFI means the vacuum states at different θ are **infinitely distinguishable** — any infinitesimal $\delta\theta$ produces an orthogonal state. A finite QFI means the vacua are **finitely distinguishable** — there is a minimum $\delta\theta$ below which the states cannot be told apart.

The mass gap is the condition that makes the θ -vacuum finitely distinguishable. This is the information-theoretic content of [A]'s Result A. The bound $\chi \leq C\Delta^4$ translates to:

$$F_Q \leq 4C, V\Delta^4$$

The QFI is finite, bounded, and controlled by the mass gap. A theory without a mass gap would have $F_Q = \infty$ — the θ -vacua would be infinitely distinguishable, and the θ -parameter would be infinitely sensitive to perturbations (as in a gapless/conformal theory, where θ is unphysical).

2.2.4 Divergence at quantum phase transitions

The connection between fidelity susceptibility (proportional to QFI) and the spectral gap is well-established in condensed matter physics (You & Li, 2007; Zanardi et al., 2007): the fidelity susceptibility diverges at a quantum phase transition because the gap closes. The QFI is a universal diagnostic:

- **Gapped phase** ($\Delta > 0$): F_Q finite, vacua finitely distinguishable.

- **Critical point** ($\Delta = 0$): $F_Q \rightarrow \infty$, vacua infinitely distinguishable (quantum phase transition).
- **Gapless phase** ($\Delta = 0$, conformal): $F_Q = \infty$, θ unphysical.

In Yang–Mills theory, the confined phase is gapped ($\Delta > 0$) and the deconfined phase is gapless ($\Delta = 0$, at $T = T_c$). The QFI diverges at the deconfinement transition — the θ -parameter becomes unphysical in the deconfined phase because the topological charge is screened by massless gluons. This is consistent with the known result that $\chi \rightarrow 0$ at deconfinement (the topological susceptibility vanishes in the deconfined phase).

2.3 The quantum Cramér–Rao bound and experimental sensitivity

2.3.1 The bound

The quantum Cramér–Rao bound (Braunstein & Caves, 1994) states: for any unbiased estimator $\hat{\theta}$ of θ based on N independent measurements:

$$\text{Var}(\hat{\theta}) \geq \frac{1}{N \cdot F_Q} = \frac{1}{4N, V, \chi}$$

The minimum detectable θ (the "sensitivity") is:

$$\delta\theta_{\min} \sim \frac{1}{2\sqrt{N, V, \chi}} \sim \frac{1}{2\sqrt{V, \Delta^2}}$$

using $\chi \leq C\Delta^4$.

2.3.2 Physical interpretation

This bound says: **the mass gap sets the minimum θ that any experiment can detect**. The larger the gap, the larger χ , and the more sensitive the experiment can be. Specifically:

Regime	χ	F_Q	$\delta\theta_{\min}$	Interpretation
Gapped ($\Delta > 0$)	$\leq C\Delta^4$ (finite)	$4V\chi$ (finite)	$\sim 1/(2\sqrt{V, \Delta^2})$	θ measurable, bounded sensitivity
Critical ($\Delta \rightarrow 0$)	$\rightarrow \infty$ (divergent)	$\rightarrow \infty$	$\rightarrow 0$	θ infinitely sensitive (phase transition)
Gapless ($\Delta = 0$)	0	0	∞	θ unmeasurable, unphysical

2.3.3 Numerical estimate for QCD

For QCD with $\Delta \sim \Lambda_{\text{QCD}} \sim 200, \text{MeV}$ and $\chi \sim \Lambda_{\text{QCD}}^4 \sim (0.2, \text{GeV})^4$:

$$\delta\theta_{\min} \sim \frac{1}{2\sqrt{V, \Lambda_{\text{QCD}}^2}}$$

For an experiment probing a spacetime volume $V \sim L^4$ with $L \sim 1, \text{fm} \sim 5, \text{GeV}^{-1}$:

$$\delta\theta_{\min} \sim \frac{1}{2 \times 5^2 \times 0.04} \sim \frac{1}{2} \sim O(1)$$

This is far above the current neutron EDM bound $|\bar{\theta}| < 10^{-10}$. The discrepancy reflects the fact that the neutron EDM is not a direct measurement of θ through the vacuum energy — it is an indirect measurement through CP-violating observables, which are suppressed by additional factors (quark masses, hadronic matrix elements). The Fisher metric gives the **intrinsic** information-theoretic sensitivity; the actual experimental sensitivity depends on the specific observable and its coupling to θ .

The Fisher bound is nevertheless conceptually important: it tells us that **the mass gap is the fundamental limit on θ -sensitivity**. No experiment can do better than $\delta\theta \sim 1/\sqrt{\chi} \sim 1/\Delta^2$, regardless of the observable. The bound is saturated only by an optimal (Helstrom) measurement, which may not be the one performed in practice.

2.3.4 Connection to [B]'s experimental program

[B] proposes three experimental pathways:

1. **Graviton-pomeron mixing at LHC/FCC** (§3.1): probes θ_{grav} through diffractive scattering. The mixing parameter $\delta \sim \kappa \Lambda_{\text{QCD}}^2 / M_{\text{Pl}}^2$ is related to χ_{grav} (the gravitational topological susceptibility). The Fisher bound gives $\delta\theta_{\text{grav}} \sim 1/\sqrt{\chi_{\text{grav}}} \sim L_{\text{AdS}}^2 / \Delta^2$ (using [A]'s $\chi_{\text{grav}} \sim (\Delta/L)^4$). The LHC sensitivity $\delta \sim 10^{-3}$ is well above this bound, so the experiment is not Fisher-limited — it is limited by detector resolution and luminosity.
2. **Gravitational form factors at EIC** (§3.2): probes the trace anomaly (related to χ through dimensional analysis). The EIC sensitivity is set by the form factor measurement precision, not by the Fisher bound.
3. **QCD metamaterials in cold-atom traps** (§3.3): probes the trace anomaly engineering. The tabletop sensitivity is set by the cold-atom coherence time and trap parameters.

In all three cases, the Fisher bound is far below the experimental sensitivity — meaning the experiments are **not limited by quantum distinguishability** but by classical measurement precision. The Fisher bound becomes relevant only in the regime where the measurement approaches the quantum limit (single-quantum measurements, quantum metrology), which is far beyond current experimental capability for QCD.

The conceptual value of the Fisher bound is that it provides a **theoretical floor** for θ -sensitivity, below which no measurement — however clever — can penetrate. This floor is set by the mass gap.

2.4 The quantum geometric tensor: metric and Berry curvature

2.4.1 The quantum geometric tensor

The **quantum geometric tensor** (QGT; Provost & Valleaux, 1980) on the parameter manifold is:

$$Q_{\theta\theta} = \langle \partial_{\theta} \Omega | (1 - |\Omega\rangle\langle\Omega|) | \partial_{\theta} \Omega \rangle$$

$$\langle \Omega | \text{rangle} = \sum_{n \neq 0} \frac{\langle \Omega | \text{rangle}_n | \hat{Q} | \Omega \rangle}{|\langle \Omega | \text{rangle}|^2} (E_n - E_0)^2$$

The QGT is a complex quantity that decomposes as:

$$Q_{\theta\theta} = g_{\theta\theta} + \frac{i}{2} \text{mathcal{F}}_{\theta\theta}$$

where:

- $g_{\theta\theta} = \text{Re}(Q_{\theta\theta}) = \chi$ (per volume) is the **quantum metric** — the Fisher information metric. This is the symmetric part, already identified in §2.1.
- $\text{mathcal{F}}_{\theta\theta} = -2i \text{Im}(Q_{\theta\theta})$ is the **Berry curvature** — the antisymmetric part.

For a single real parameter θ , the Berry curvature vanishes (it requires at least two parameters). But with multiple θ -parameters (the three-sector case of §3), the QGT becomes a matrix with nontrivial Berry curvature.

2.4.2 Berry curvature and topological charge

With three θ -parameters $(\theta_{\text{gauge}}, \theta_{\text{grav}}, \theta_{\text{NY}})$, the QGT is a 3×3 matrix:

$$Q_{ij} = \sum_{n \neq 0} \frac{\langle \Omega | \hat{Q}_i | n \rangle \langle n | \hat{Q}_j | \Omega \rangle}{|\langle \Omega | \text{rangle}|^2} (E_n - E_0)^2$$

The symmetric part is $g_{ij} = \text{Re}(Q_{ij}) = \chi_{ij}/V$ — the topological susceptibility matrix ([A], §6.5, extended to three sectors in §3.8). The antisymmetric part is the Berry curvature:

$$\text{mathcal{F}}_{ij} = -2i \text{Im}(Q_{ij}) = -2 \sum_{n \neq 0} \text{Im} \left[\frac{\langle \Omega | \hat{Q}_i | n \rangle \langle n | \hat{Q}_j | \Omega \rangle}{|\langle \Omega | \text{rangle}|^2} \right] (E_n - E_0)^2$$

The Berry phase accumulated when the θ -parameters traverse a closed loop γ in the 3D θ -manifold (T^3) is:

$$\gamma_{\text{Berry}} = \oint_{\gamma} \text{mathcal{A}}_i d\theta^i = \int_{\Sigma} \text{mathcal{F}}_{ij} d\theta^i \wedge d\theta^j$$

where Σ is a surface bounded by γ and $\text{mathcal{A}}_i = i \langle \Omega | \partial_{\theta_i} | \Omega \rangle$ is the Berry connection.

2.4.3 The Berry phase IS the topological charge

The Berry connection for the θ -vacuum is:

$$\text{mathcal{A}}_i = i \langle \Omega(\theta) | \partial_{\theta_i} | \Omega(\theta) \rangle = i \sum_{n \neq 0} \frac{\langle \Omega | \hat{Q}_i | n \rangle \langle n | \hat{Q}_i | \Omega \rangle}{|\langle \Omega | \text{rangle}|^2} (E_0 - E_n) \cdot \text{(phase factors)}$$

When θ traverses a full cycle $\theta_i \rightarrow \theta_i + 2\pi$, the Berry phase is:

$$\gamma_{\text{Berry}} = 2\pi \cdot Q_i$$

where Q_i is the topological charge in sector i (gauge, gravitational, or Nieh–Yan). This is because the θ -vacuum is a superposition of topological sectors:

$$|\Omega(\theta)\rangle = \sum_Q e^{i\theta Q} |Q\rangle$$

and a 2π shift in θ produces a phase $e^{i2\pi Q} = 1$ for integer Q , but the Berry phase

tracks the **winding** — the number of times the state winds around the topological sector, which is Q .

This connection between the Berry phase and the topological charge has been established in QCD (Tanizaki & Ünsal, 2014; arXiv:2506.00078, 2025), where the θ -parameter is interpreted as a Berry phase of the QCD vacuum. The θ -vacuum structure — discrete quasivacua separated by domain walls where θ jumps by 2π — is a Berry phase phenomenon.

The Berry phase on θ -space IS the topological charge, and the quantum metric IS the topological susceptibility. The mass gap ensures both are finite:

- The metric $g_{ij} = \chi_{ij}$ is finite because $\chi_{ij} \leq \sqrt{\chi_{ii}\chi_{jj}} \leq C/\Delta^4$ (Result A + Cauchy–Schwarz).
- The Berry curvature $\mathcal{F}_{ij}$ is finite because it involves the same spectral sum with the same gap-protected denominators $(E_n - E_0)^2 \geq \Delta^2$.

2.5 The multi-parameter quantum Cramér–Rao bound

2.5.1 The multi-parameter QFI matrix

With three θ -parameters, the QFI is a 3×3 matrix:

$$F_{ij} = 4V \text{Re}(Q_{ij}) = 4V \chi_{ij}$$

This is exactly the topological susceptibility matrix of §3.8, multiplied by $4V$. The multi-parameter quantum Cramér–Rao bound (Matsumoto, 2002) states:

$$\text{Cov}(\hat{\theta}_i, \hat{\theta}_j) \geq (F^{-1})_{ij} = \frac{1}{4V} (\chi^{-1})_{ij}$$

The simultaneous estimation precision for all three θ -parameters is bounded by the **inverse of the susceptibility matrix**. This is the information-theoretic meaning of the 3×3 susceptibility matrix: it is the Fisher metric on the 3D θ -torus T^3 , and its inverse bounds the estimation covariance.

2.5.2 The gauged P/CP as a Fisher-metric statement

The gauged P/CP mechanism (§3.5) sets $\theta_i = 0$ for all three sectors. In Fisher-metric language, this means:

- The vacuum is at the **origin** of θ -space ($\theta = 0$).
- The Fisher metric at the origin is $\chi_{ij}|_{\theta=0}$, which is positive semi-definite (vacuum stability).
- The gauged P/CP **pins** the vacuum at the origin — the symmetry prevents the vacuum from drifting to $\theta \neq 0$.
- The Fisher metric determines how tightly the vacuum is pinned: the larger χ_{ij} , the stiffer the vacuum against θ -perturbations.

The mass gap sets the stiffness: $\chi_{ij} \leq C/\Delta^4$ means the vacuum is stiff (resistant to θ -perturbations) when the gap is large. Without the gap, $\chi \rightarrow \infty$ (or $\chi \rightarrow 0$, depending on the direction), and the vacuum becomes either infinitely stiff (gapless, conformal) or infinitely soft (deconfined).

2.5.3 Non-commutativity and the incompatibility of simultaneous estimation

In quantum multi-parameter estimation, the QFI matrix may not be simultaneously saturable for all parameters — the **non-commutativity** of the optimal measurements for different parameters introduces a tradeoff (Matsumoto, 2002). The incompatibility is measured by:

$$\mathcal{I}_{ij} = \text{Im}(Q_{ij}) = -\frac{1}{2} \mathcal{F}_{ij}$$

If $\mathcal{F}_{ij} = 0$ for all i, j , all parameters can be simultaneously estimated at the quantum limit. If $\mathcal{F}_{ij} \neq 0$, there is a tradeoff: estimating θ_i more precisely degrades the estimation of θ_j .

In our case, $\mathcal{F}_{ij}$ is the Berry curvature, which is related to the mixed topological charges (e.g., the CEH instanton carries both Q_{gauge} and Q_{grav}), producing nonzero $\mathcal{F}_{\text{gauge,grav}}$. ***The CEH instanton is the source of quantum incompatibility in the simultaneous estimation of θ_{QCD} and θ_{grav} .***

This is a deep connection: the CEH bundling mechanism ([A], §8; §3.7) — which bundles the θ -parameters through a shared instanton — is the same phenomenon that creates Berry curvature and prevents the simultaneous optimal estimation of all three θ -parameters. The stronger the CEH amplitude, the larger the Berry curvature, and the more incompatible the estimation.

2.6 The Fisher metric and the spectral function

2.6.1 The spectral decomposition of the Fisher metric

The Fisher metric has the spectral representation (combining §2.1 and §2.2):

$$\chi = \int_{\Delta^2}^{\infty} \frac{d\mu^2}{\rho(\mu^2)} \mu^2$$

This integral is the **spectral decomposition of the Fisher information**: the Fisher metric is a weighted integral over the spectral density $\rho(\mu^2)$, with weight $1/\mu^2$. The mass gap Δ sets the lower limit, making the integral converge.

The spectral density $\rho(\mu^2)$ is the same object that appears in:

- [A]'s Result A (the Källén–Lehmann representation for χ),
- The QFI formula ($F_Q = 4V\chi$),
- The Banks–Casher relation (Layer I.1 of the bridge: $\rho(0)$ controls the chiral condensate),
- The Gutzwiller–Selberg trace formula (Layer I.2: periodic orbits/instantons contribute to $\rho(\mu^2)$),
- The entanglement gap (Layer IV.3: the entanglement spectrum is a spectral transform of $\rho(\mu^2)$).

The Fisher metric is the **unifying object**: it connects the spectral, topological, and information-theoretic aspects of the mass gap through a single integral.

2.6.2 Direct measurability of the Fisher metric

The spectral density $\rho(\mu^2) = \sum_n |\langle 0 | q | n \rangle|^2 \delta(\mu^2 - m_n^2)$ is, in principle, **directly measurable**: it is the spectral function of the pseudoscalar glueball channel (0^{-} -

χ). In lattice QCD ([B]'s tool), one computes the correlator $\langle q(0), q(x) \rangle$ and extracts $\rho(\mu^2)$ via the maximum entropy method or Bayesian spectral reconstruction.

The fidelity susceptibility (proportional to χ) has been shown to be directly measurable in condensed matter experiments via neutron scattering or angle-resolved photoemission spectroscopy (ARPES) (Liu et al., 2016): "an equality relating the fidelity susceptibility and the spectral function has been derived, making the fidelity susceptibility directly measurable in experiments." The same method, applied to the QCD pseudoscalar channel, would give a direct lattice measurement of the Fisher metric.

This connects the Fisher metric to [B]'s experimental program: the topological susceptibility χ (the Fisher metric) is not a theoretical abstraction — it is a **measurable spectral function**, computable on the lattice and, in principle, extractable from experimental data on glueball production.

2.7 Summary of Section 2

Result	Status	Connection
$\chi = \text{Fisher metric on } \theta\text{-space}$	Exact (§2.1)	[A]'s $\chi \rightarrow$ information geometry
$F_Q = 4V\chi$ (quantum Fisher information)	Exact (§2.2)	QFI $\leftrightarrow$ Källén–Lehmann spectral integral
Mass gap $\rightarrow$ QFI finite $\rightarrow$ vacua finitely distinguishable	Proven (§2.2.3)	[A]'s Result A $\rightarrow$ quantum distinguishability
$\Delta\theta_{\min} \sim 1/\sqrt{\chi} \sim 1/\Delta^2$ (Cramér–Rao)	Derived (§2.3)	Mass gap $\rightarrow$ experimental sensitivity floor
QGT: metric = χ , Berry curvature = topological charge	Established (§2.4)	Berry phase $\leftrightarrow$ Q ; metric $\leftrightarrow$ susceptibility
Multi-parameter: $\text{Cov}(\hat{\theta}) \geq (4V\chi_{ij})^{-1}$	Derived (§2.5)	3×3 susceptibility matrix $\rightarrow$ Fisher matrix
CEH instanton $\rightarrow$ Berry curvature $\rightarrow$ estimation incompatibility	Identified (§2.5.3)	CEH bundling $\leftrightarrow$ quantum non-commutativity
χ directly measurable via spectral function	Established (§2.6)	Lattice QCD $\rightarrow$ Fisher metric measurement

The main theorem

Theorem (Fisher metric bridge). The topological susceptibility χ is the Fisher information metric on the manifold of θ -vacuum states. The quantum Fisher information is $F_Q = 4V\chi$, given by the same Källén–Lehmann spectral integral that [A] uses to prove Result A. The mass gap $\Delta > 0$ is the condition that makes the QFI finite, bounding the distinguishability of θ -vacua from below: $\Delta\theta_{\min} \sim 1/(2\sqrt{V}\Delta^2)$. The quantum geometric tensor

decomposes into the Fisher metric (χ_{ij}) and the Berry curvature (topological charge), unifying the spectral, topological, and information-theoretic structures through a single object.

3. The 5D Anomaly Inflow Identification

3.1 Introduction: Two Results, One Theory

Two results from the companion papers appear, at first sight, to address completely different questions:

- [B] constructs a 5D gauge theory with torsion on $M_5 = \mathbb{R}^{3,1} \times S^1$ and proves the 4D Yang–Mills mass gap $\Delta > 0$ via a convergent polymer expansion and Sturm–Liouville analysis of the KK-reduced spectrum.
- [A] argues that the Chern–Weil (-1) -form global symmetry — generated by the closed 4-forms $\text{Tr}(F \wedge F)$ and $\text{Tr}(R \wedge R)$ — must be gauged by quantum gravity (via the no-global-symmetries conjecture), and that the gauging of discrete P/CP forces $\theta_{\text{QCD}} = \theta_{\text{grav}} = 0$ at the Lagrangian level.

In this section, we demonstrate that these are not merely complementary results — they are **two sectors of the same 5D theory**. The 5D construction of [B] is the anomaly-inflow bulk (in the sense of Callan & Harvey, 1985) that physically realizes the Chern–Weil gauging of [A]. The torsion, far from obstructing this identification, extends it: it introduces a third topological sector — the Nieh–Yan form — and the gauged P/CP mechanism forces all three θ -parameters to zero simultaneously.

The mass gap proof and the $\theta = 0$ result use **different, decoupled sectors** of the 5D action:

Sector	5D Terms	4D Consequence	Paper
Dynamical (kinetic + potential)	$S_{F_{MN}} F^{MN}$, torsion kinetic	Mass gap $\Delta > 0$ (Sturm–Liouville)	[B]
Topological (Chern–Simons)	$S_{CS_5(A)}$, $S_{CS_5(\omega)}$, $S_{CS_5^{\text{NY}}}$	$\theta_{\text{QCD}} =$ $\theta_{\text{grav}} = \theta_{\text{NY}}$ $= 0$ (gauged P/CP)	[A] + this work

The 5D theory does double duty: it proves the mass gap (dynamical sector) and sets $\theta = 0$ (topological sector) through independent mechanisms, unified by the shared 5D geometry.

3.2 The Three Topological 4-Forms

We work in Riemann–Cartan geometry on the 4D boundary $M_4 = \mathbb{R}^{3,1}$, with vierbein e^a , spin connection ω^{ab} , torsion 2-form $T^a = De^a = de^a + \omega^a_b \wedge e^b$, and curvature 2-form $R^{ab} = d\omega^{ab} + \omega^a_c \wedge \omega^{cb}$. The gauge connection A is a Lie-algebra-valued 1-form with field strength $F = dA + A \wedge A$.

Three closed 4-forms govern the topological sectors of the 4D theory:

$$q_{\text{gauge}} = \frac{1}{8\pi^2} \text{Tr}(F \wedge F), \quad q_{\text{grav}} = \frac{1}{24\pi^2} \text{Tr}(R \wedge R), \quad q_{\text{NY}} = \frac{1}{24\pi^2} \text{Tr}(R \wedge T - T \wedge T)$$

$$\frac{1}{8\pi^2} \text{Tr}(R \wedge R), \quad \text{qquad } q_{\text{NY}} = \frac{1}{8\pi^2} \text{mathcal{N}}$$

where the **Nieh–Yan 4-form** (Nieh & Yan, 1982) is:

$$\text{mathcal{N}} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$$

All three share identical structural properties:

Property	$\text{Tr}(F \wedge F)$	$\text{Tr}(R \wedge R)$	$\text{mathcal{N}}$
Closed	$d, \text{Tr}(F \wedge F) = 0$ (Bianchi)	$d, \text{Tr}(R \wedge R) = 0$ (Bianchi)	$d, \text{mathcal{N}} = 0$ (Riemann–Cartan Bianchi)
Total derivative	$d, \text{Tr}(A, dA + \frac{1}{2} A^3) =$	$d, \text{Tr}(\omega, d\omega + \frac{1}{2} \omega^3) =$	$d, U_4, U_4 = \omega^a \wedge e_a \wedge e_b$
P-odd (orientation reversal)	✓	✓	✓
T-odd	✓	✓	✓
Integer-quantized (compact M_4)	$Q \in \mathbb{Z}$	$\sigma \in \mathbb{Z}$	$Q_{\text{NY}} \in \mathbb{Z}$
Invisible perturbatively	✓ (total derivative)	✓ (total derivative)	✓ (total derivative)
Physical through instantons	Gauge instantons	Gravitational instantons	Torsion instantons

The Nieh–Yan form is the torsionful analogue of the Pontryagin density. Its integral:

$$Q_{\text{NY}} = \frac{1}{8\pi^2} \int_{M_4} \text{mathcal{N}}$$

is an integer on compact manifolds (related to the difference of Pontryagin classes for two vierbein conventions; Kreimer & Mielke, 2001). The fact that it is P-odd — changing sign under orientation reversal $e^a \rightarrow -e^a$ — is the crucial property that allows the gauged P/CP mechanism to act on it, exactly as it does on the gauge and gravitational Pontryagin densities.

3.2.1 The combined index theorem with torsion

The Atiyah–Singer index theorem, extended to Riemann–Cartan spacetimes (Chandia & Zanelli, 1997), gives the index of the Dirac operator coupled to both the gauge field and the spacetime geometry with torsion:

$$\text{index}(D) = T(R) \cdot Q_{\text{gauge}} + \dim(R) \cdot \hat{A}(M) + c_{\text{NY}} \cdot Q_{\text{NY}}$$

where:

- $T(R)$ is the Dynkin index of the fermion representation,
- $\hat{A}(M) = -\sigma/8$ is the $\hat{A}$ -genus (gravitational contribution, as in [A], §5.5),
- c_{NY} is the Nieh–Yan coefficient, proportional to ℓ^{-2} where ℓ is the UV cutoff (see §3.6 below).

This is the natural extension of [A]'s combined index theorem (§5.5) to the torsionful case. All three topological charges — gauge, gravitational, and torsion — enter a single integer that counts the chiral zero modes of the Dirac operator. They are not independent: they jointly determine the θ -dependence of the vacuum energy.

3.3 The 5D Chern–Simons Forms and Torsion-Independence

3.3.1 The gauge Chern–Simons 5-form

The gauge Chern–Simons 5-form on M_5 is:

$$CS_5(A) = \text{Tr} \left(A \wedge dA \wedge dA + \frac{3}{2} A \wedge A \wedge A \wedge dA + \frac{3}{5} A \wedge A \wedge A \wedge A \wedge A \right)$$

satisfying $dCS_5(A) = \text{Tr}(F \wedge F)$.

Theorem (torsion-independence of CS_5). The gauge Chern–Simons 5-form $CS_5(A)$ depends only on the gauge connection A and its exterior derivative dA . It is independent of the spacetime metric, the spin connection ω , and the torsion T^a .

Proof. The gauge field strength $F_{MN} = \partial_M A_N - \partial_N A_M + [A_M, A_N]$ is defined using only the exterior derivative d and the gauge Lie bracket. Neither operation involves the spacetime connection. The Chern–Simons form is a polynomial in A , dA , and $A \wedge A$ — all gauge-sector objects. Since Chern–Simons theory is a Schwarz-type topological field theory (metric-independent by construction; Chern & Simons, 1974), no spacetime geometric data — metric, spin connection, or torsion — enters. $\square$

This is the foundational result for the bridge: **the primary requirement for anomaly inflow — the existence of $CS_5(A)$ in the 5D bulk — is satisfied regardless of the torsion content of the 5D theory.**

3.3.2 The gravitational and torsion Chern–Simons 5-forms

By exact analogy:

- The **gravitational** Chern–Simons 5-form $CS_5(\omega) = \text{Tr}(\omega \wedge d\omega \wedge d\omega + \cdots)$ satisfies $dCS_5(\omega) = \text{Tr}(R \wedge R)$. It depends only on the spin connection ω , not on the gauge field A .
- The **torsion** (Nieh–Yan) Chern–Simons 5-form $CS_5^{\text{NY}}(\omega, e)$ satisfies $dCS_5^{\text{NY}} = \mathcal{N}$. Since $\mathcal{N} = dU_4$ with $U_4 = \omega^a \wedge e_a \wedge e_b$, the 5D extension is obtained by the standard Stora–Zumino descent procedure. Its existence is guaranteed by the fact that $\mathcal{N}$ is a closed, exact 4-form.

3.3.3 The three descent chains

Each 4-form descends to a 5D bulk Chern–Simons term through the standard descent:

$$\omega_5^{(0)} \xrightarrow{d} \omega_4^{(1)} \xrightarrow{\delta} \omega_4^{(0)}$$

where $\omega_4^{(0)}$ is the 4D anomaly, $\omega_4^{(1)}$ is the consistent anomaly, and $\omega_5^{(0)}$ is the 5D bulk Chern–Simons form. The three chains are:

Descent chain	$\omega_5^{(0)}$ (bulk)	$\omega_4^{(0)}$ (boundary anomaly)
Gauge	$CS_5(A)$	$\text{Tr}(F \wedge F)$
Gravitational	$CS_5(\omega)$	$\text{Tr}(R \wedge R)$
Torsion (Nieh–Yan)	$CS_5(\text{NY})(\omega, e)$	$\mathcal{N}$

All three are independent — they involve different geometric objects (A , ω , and (ω, e) respectively). The torsion in the 5D theory does not obstruct any descent chain; it **adds a third chain** alongside the two that [A] identified.

3.4 Anomaly Inflow: The Callan–Harvey Mechanism in the 5D Theory

3.4.1 The bulk action

The 5D action of [B], supplemented by the topological sector required for anomaly cancellation, is:

$$S_{5D} = \int_{M_5} d^5x \sqrt{-g} \left(-\frac{1}{4} F_{MN} F^{MN} + \mathcal{L}_{\text{torsion}} \right) + \frac{k_{\text{gauge}}}{24\pi^2} \int_{M_5} CS_5(A) + \frac{k_{\text{grav}}}{24\pi^2} \int_{M_5} CS_5(\omega) + \frac{k_{\text{NY}}}{24\pi^2} \int_{M_5} CS_5(\text{NY})(\omega, e)$$

The first line is the **dynamical sector** ([B]'s original construction): the Yang–Mills kinetic term plus torsion terms. This sector produces the mass gap through the polymer expansion and Sturm–Liouville analysis.

The second line is the **topological sector**: the three Chern–Simons 5-forms. These are metric-independent, do not affect the equations of motion for the dynamical fields (they are topological), and do not contribute to the Sturm–Liouville potential $V(z)$. They are the anomaly-inflow terms.

The mass gap proof (dynamical sector) and the θ -gauging mechanism (topological sector) are decoupled. The 5D theory does double duty through independent mechanisms.

3.4.2 The anomaly inflow

The Callan–Harvey mechanism (1985) states: if the 5D bulk contains a Chern–Simons term CS_5 , then the variation of this term under a gauge (or diffeomorphism) transformation produces a bulk current whose divergence, evaluated on the 4D boundary, exactly cancels the 4D anomaly:

$$\nabla_M J^\mu_{\text{bulk}} + \nabla_\mu J^\mu_{\text{bdy}} = 0$$

For each of the three sectors:

Gauge sector. The bulk gauge Chern–Simons term $\frac{k_{\text{gauge}}}{24\pi^2} \int \text{CS}_5(A)$ produces a bulk current $J^\mu_{\text{gauge}} = \frac{\delta S}{\delta A_\mu} \propto \text{Tr}(F \wedge F)$. On the 4D boundary, the divergence cancels the gauge anomaly $\nabla_\mu J^\mu_5 = \frac{1}{32\pi^2} \text{Tr}(F \wedge F)$.

Gravitational sector. The bulk gravitational Chern–Simons term $\frac{k_{\text{grav}}}{24\pi^2} \int \text{CS}_5(\omega)$ produces a bulk current $J^\mu_{\text{grav}} = \frac{\delta S}{\delta \omega_\mu} \propto \text{Tr}(R \wedge R)$. On the 4D boundary, the divergence cancels the gravitational anomaly $\nabla_\mu J^\mu_5 = \frac{1}{384\pi^2} \text{Tr}(R \wedge R)$.

Torsion (Nieh–Yan) sector. The bulk torsion Chern–Simons term $\frac{k_{\text{NY}}}{24\pi^2} \int \text{CS}_5^{\text{NY}}(\omega, e)$ produces a bulk current $J^\mu_{\text{NY}} = \frac{\delta S}{\delta e_\mu} \propto \mathcal{N}$. On the 4D boundary, the divergence cancels the torsion anomaly $\nabla_\mu J^\mu_5 = \frac{1}{384\pi^2} \mathcal{N}$.

All three mechanisms work identically: the bulk Chern–Simons term is the anomaly-inflow source, the boundary anomaly is the inflow target, and the cancellation is exact (Callan & Harvey, 1985; see also arXiv:1909.08775 for the relation to the η -invariant).

3.4.3 The θ -parameters as S^1 holonomies

Under Kaluza–Klein compactification on S^1 of radius R with coordinate $y \sim y + 2\pi R$, the 5D connections decompose:

$$A = A_\mu dx^\mu + A_5 dy, \quad \omega = \omega_\mu dx^\mu + \omega_5 dy$$

The 5D Chern–Simons terms, when integrated over S^1 , produce 4D θ -terms:

$$\int_{M_5} \text{CS}_5(A) \supset \int_{M_4} \text{Tr}(F \wedge F) \cdot \oint_{S^1} A_5 dy = \theta_{\text{QCD}} \cdot \int_{M_4} \text{Tr}(F \wedge F)$$

where:

$$\theta_{\text{QCD}} = \oint_{S^1} A_5 dy, \quad \theta_{\text{grav}} = \oint_{S^1} \omega_5 dy, \quad \theta_{\text{NY}} = \oint_{S^1} \phi_{\text{NY}}, dy$$

The θ -parameters are geometric quantities — holonomies of the 5D connections around the compact S^1 . They are not free parameters of the 4D theory; they are determined by the 5D vacuum configuration.

This is the key identification: $[A]$'s θ_{QCD} and θ_{grav} (introduced as Lagrangian parameters in the 4D theory) are revealed to be S^1 holonomies in the 5D parent theory. The 5D construction of [B], which was designed to prove the mass gap, simultaneously provides the geometric origin of the θ -parameters.

3.5 Gauged P/CP as Orientation Reversal of S^1

3.5.1 Orientation reversal

The 5D manifold is $M_5 = M_4 \times S^1$. An orientation reversal of the S^1 factor sends $y \rightarrow -y$, reversing the orientation of the compact direction. This is a discrete symmetry of the 5D theory (it preserves M_4 and reverses only the S^1).

Under this reversal, the S^1 holonomy flips sign:

$$\oint_{S^1} \phi, dy \rightarrow \oint_{S^1} \phi, (-dy) = -\oint_{S^1} \phi$$

This applies to **all three** θ -parameters:

$$\theta_{\text{QCD}} \rightarrow -\theta_{\text{QCD}}, \quad \theta_{\text{grav}} \rightarrow -\theta_{\text{grav}}, \quad \theta_{\text{NY}} \rightarrow -\theta_{\text{NY}}$$

The reason is the same in each case: the holonomy is an oriented integral, and reversing the orientation flips its sign.

3.5.2 P-oddness of the 4-forms

Each of the three 4-forms is P-odd — it changes sign under orientation reversal of M_4 :

- $\text{Tr}(F \wedge F) \rightarrow -\text{Tr}(F \wedge F)$ (gauge Pontryagin density; [A], §4.3)
- $\text{Tr}(R \wedge R) \rightarrow -\text{Tr}(R \wedge R)$ (gravitational Pontryagin density; [A], §5.3)
- $\mathcal{N} \rightarrow -\mathcal{N}$ (Nieh–Yan form; §3.2 above)

The S^1 orientation reversal induces an M_4 orientation reversal (because the total 5D orientation is preserved). The combined effect is: each θ -parameter flips sign, and each 4-form flips sign, so the product $\theta \cdot q$ (the θ -term in the action) is invariant — but only if $\theta \rightarrow -\theta$ and $q \rightarrow -q$ simultaneously.

3.5.3 The gauging argument

If the S^1 orientation reversal is a **gauged** discrete symmetry (not merely a global one), then by Elitzur's theorem the vacuum must preserve it. The partition function must be invariant under $\theta \rightarrow -\theta$:

$$Z(\theta) = Z(-\theta) \quad \text{for all } \theta$$

Since $Z(\theta) = \sum_Q e^{i\theta Q}$, Z_Q (sum over topological sectors), the invariance requires $e^{i\theta Q} = e^{-i\theta Q}$ for all Q , which forces:

$$\boxed{\theta_{\text{QCD}} = 0, \quad \theta_{\text{grav}} = 0, \quad \theta_{\text{NY}} = 0}$$

This is [A]'s gauged P/CP argument (§4.3), now extended to **three** topological sectors. The torsion in the 5D theory introduces the Nieh–Yan sector, and the gauged P/CP mechanism extends to cover it automatically — the orientation reversal flips θ_{NY} just as it flips θ_{QCD} and θ_{grav} .

3.5.4 The no-global-symmetries conjecture

[A] (§4.2) argues that the Chern–Weil (-1) -form global symmetry (generated by $\text{Tr}(F \wedge F)$ and $\text{Tr}(R \wedge R)$) must be gauged by quantum gravity, because quantum gravity has no exact global symmetries (the swampland conjecture; Heidenreich et al., 2021). The 5D anomaly-inflow mechanism provides a **concrete realization** of this gauging:

- The S^1 orientation reversal is the gauge transformation of the discrete P symmetry.

- The Chern–Simons 5-forms are the gauge fields of the (-1) -form global symmetries (the bulk $\mathcal{CS}_5$ is the gauge field; the boundary θ is the holonomy).
- The anomaly inflow ensures that the gauging is consistent — the 4D anomaly is cancelled by the 5D bulk current.

In the torsionful theory, the Nieh–Yan (-1) -form global symmetry (generated by $\mathcal{N}$) is also gauged, by the same mechanism. The 5D theory gauges **all three** (-1) -form symmetries simultaneously.

3.6 The Nieh–Yan Cutoff Dependence

3.6.1 The subtlety

Unlike the gauge and gravitational anomaly coefficients, which are quantized (integer Chern–Simons levels k_{gauge} , $k_{\text{grav}} \in \mathbb{Z}$), the Nieh–Yan anomaly coefficient c_{NY} carries a mass dimension:

$$c_{\text{NY}} \propto \ell^{-2}$$

where ℓ is the UV cutoff. This has been a subject of debate (Chandia & Zanelli, 1997; Kreimer & Mielke, 2001; arXiv:2409.06766, 2024). The current understanding is:

- On manifolds **without boundary**, if the Dirac index is well-defined, the Nieh–Yan invariant vanishes (arXiv:2409.06766).
- On manifolds **with boundary** or with singular torsion (defects), the Nieh–Yan term contributes, with a coefficient proportional to ℓ^{-2} .
- The torsion contribution is physically meaningful when the torsion is **dynamical** (integrated over in the path integral), not merely a background field.

3.6.2 Why the gauged P/CP is unaffected

The gauged P/CP mechanism (§3.5) is **structural, not dynamical**: it sets $\theta_{\text{NY}} = 0$ by symmetry (orientation reversal), not by computing the anomaly coefficient. The argument is:

1. $\mathcal{N}$ is P-odd (changes sign under orientation reversal).
2. The S^1 holonomy θ_{NY} flips sign under orientation reversal.
3. If orientation reversal is gauged, the vacuum must preserve it, forcing $\theta_{\text{NY}} = 0$.

None of these steps involves the coefficient c_{NY} . The coefficient affects the **magnitude** of the Nieh–Yan contribution to the vacuum energy $E(\theta_{\text{NY}}) = \frac{1}{2} \chi_{\text{NY}} \theta_{\text{NY}}^2 + \dots$, but the vacuum energy at $\theta_{\text{NY}} = 0$ is simply $E(0) = 0$ regardless of χ_{NY} .

3.6.3 The 5D theory provides a natural cutoff

In the 5D theory compactified on S^1 of radius R , the natural UV cutoff is the KK scale:

$$\ell^{-1} \sim R^{-1} \sim \Lambda_{\text{UV}}^{\text{KK}}$$

The Nieh–Yan coefficient is $c_{\text{NY}} \sim R^{-2}$. In the 4D limit $R \rightarrow 0$, this diverges — but the torsion itself scales with R (the torsion components along S^1 vanish as $R \rightarrow 0$). The product:

$$c_{\text{NY}} \cdot \int \mathcal{N} \sim R^{-2} \cdot (R \cdot \Lambda_{\text{QCD}})^2 \sim \Lambda_{\text{QCD}}^2$$

is finite in the 4D limit. The cutoff dependence is absorbed into the 4D torsion coupling, which is renormalized along with the other 4D couplings.

3.6.4 The torsion is dynamical

[B]'s 5D theory treats the torsion as a **dynamical field** (it appears in the action with kinetic terms). When torsion is dynamical, the Nieh–Yan term is a well-defined contribution to the effective action — the cutoff dependence is absorbed into the torsion field renormalization, just as the gauge anomaly coefficient is absorbed into the Chern–Simons level renormalization. This is in contrast to the case of background torsion, where the cutoff dependence is genuinely ambiguous.

3.6.5 Summary

The Nieh–Yan cutoff dependence is a genuine subtlety that should be discussed openly. It does not affect the structural result ($\theta_{\text{NY}} = 0$ by gauged P/CP), but it does affect the quantitative prediction of the Nieh–Yan contribution to the vacuum energy. We treat the Nieh–Yan susceptibility χ_{NY} as an order-of-magnitude estimate, not a precision prediction.

3.7 The CEH Instanton: Three-Sector Bundling

3.7.1 The colored Eguchi–Hanson instanton

[A] (§8) constructs the colored Eguchi–Hanson (CEH) instanton: a Ricci-flat, self-dual gravitational instanton whose $SU(2)$ self-dual spin connection is embedded into the color group $SU(N)$. It carries:

$$Q_{\text{gauge}} \neq 0, \quad Q_{\text{grav}} = \pm 1$$

with classical action $S_{\text{CEH}} = 8\pi^2/g^2$ (gauge-scale, not Planck-suppressed, because the Einstein–Hilbert action vanishes on Ricci-flat backgrounds).

3.7.2 Extension to the torsionful case

In the 5D theory with torsion, the CEH instanton acquires a third topological charge:

$$Q_{\text{NY}} = \frac{1}{8\pi^2} \int_{M_4} \mathcal{N} = \frac{1}{8\pi^2} \int_{M_4} \left(T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b \right)$$

Since the CEH instanton has $R_{\alpha\beta\mu\nu} \neq 0$ (it is self-dual but not flat) and the torsion is generically nonzero on the instanton background (the torsion is an independent field), $Q_{\text{NY}} \neq 0$ in general.

The CEH instanton now **bundles all three θ -parameters**:

$$\bar{\theta}_{\text{CEH}} = \theta_{\text{QCD}} + \theta_{\text{grav}} + \theta_{\text{NY}} + \arg \det M$$

3.7.3 The gauged P/CP kills all three

With gauged P/CP setting $\theta_{\text{QCD}} = \theta_{\text{grav}} = \theta_{\text{NY}} = 0$:

$$\bar{\theta}_{\text{CEH}} = 0 + 0 + 0 + \arg \det M$$

The CEH contribution to the vacuum energy is:

$$E_{\text{CEH}} \sim K_{\text{CEH}} \cos(\bar{\theta}_{\text{CEH}}) = K_{\text{CEH}} \cos(\arg \det M)$$

This is CP-conserving (no θ -dependent CP violation), **regardless of the CEH amplitude** K_{CEH} — whether it is gauge-scale or Planck-suppressed. This is exactly [A]’s robustness argument (§9.2), now extended to three sectors.

3.7.4 The axion quality problem is now three-fold

The QCD axion couples to $\text{Tr}(F \wedge F)$, shifting $\theta_{\text{QCD}} \rightarrow \theta_{\text{QCD}} + a/f_a$. It does **not** couple to $\text{Tr}(R \wedge R)$ or $\mathcal{N}$, so it cannot shift θ_{grav} or θ_{NY} . If the CEH amplitude is gauge-scale, the axion would need to cancel **three** independent θ -parameters simultaneously — which it cannot do with a single field. The axion quality problem, identified in [A] (§8.6) for two sectors, is now **three-fold**.

The gauged P/CP, by contrast, kills all three at once — it is the only single-mechanism solution that works in all scenarios.

3.8 The Topological Susceptibility Matrix

3.8.1 The 3×3 susceptibility

The 4D vacuum energy depends on all three θ -parameters:

$$E(\theta_{\text{QCD}}, \theta_{\text{grav}}, \theta_{\text{NY}}) - E(0,0,0) = \frac{1}{2} \begin{pmatrix} \theta_{\text{QCD}} & \theta_{\text{grav}} & \theta_{\text{NY}} \end{pmatrix} \begin{pmatrix} \chi_{\text{QCD}} & \chi_{\text{cross}}^{(12)} & \chi_{\text{cross}}^{(13)} \\ \chi_{\text{cross}}^{(12)} & \chi_{\text{grav}} & \chi_{\text{cross}}^{(23)} \\ \chi_{\text{cross}}^{(13)} & \chi_{\text{cross}}^{(23)} & \chi_{\text{NY}} \end{pmatrix} \begin{pmatrix} \theta_{\text{QCD}} \\ \theta_{\text{grav}} \\ \theta_{\text{NY}} \end{pmatrix} + O(\theta^4)$$

where $\chi_{ij} = \int d^4x, \langle q_i(0), q_j(x) \rangle_{\theta=0}$ is the 3×3 topological susceptibility matrix.

3.8.2 Bounds

By Result A ([A], §2), extended to each sector: if the mass gap $\Delta > 0$ controls the spectral density in each sector, then:

$$0 < \chi_{ii} \leq C_i \cdot \Delta^4 \quad \text{for each } i \in \{\text{gauge, grav, NY}\}$$

(The gravitational and Nieh–Yan sectors require a spectral gap in the pseudoscalar sector; see [A], §§6–7.)

By the Cauchy–Schwarz inequality ([A], §6.5, extended):

$$|\chi_{ij}| \leq \sqrt{\chi_{ii} \chi_{jj}} \quad \text{for all } i, j$$

Vacuum stability (positive semi-definiteness of the energy) requires:

$$\det(\chi_{ij}) \geq 0$$

3.8.3 The Fisher information geometry

As discussed in the bridge (Layer IV.1), the susceptibility matrix χ_{ij} is the **Fisher information metric** on the 3-dimensional θ -manifold (T^3 for three periodic θ -parameters). Its eigenvalues determine the experimental distinguishability of vacuum states at nearby θ -values:

$$\delta\theta_{\min} \sim \frac{1}{\sqrt{\lambda_{\min}(\chi)}}$$

where $\lambda_{\min}(\chi)$ is the smallest eigenvalue. The bound $\chi_{ii} \leq C_i \Delta^4$ bounds the distinguishability from above: if the mass gap is large, vacuum states at different θ are easily distinguishable. The gauged P/CP sets $\theta = 0$, so the Fisher metric determines how tightly the experimental constraints (neutron EDM, collider data) confirm $\theta = 0$.

3.9 Summary of Section 3

Claim	Status	Mechanism
$CS_5(A)$ is torsion-independent	Proven (§3.3)	CS_5 uses only gauge connection; torsion is spacetime-sector
Torsion introduces a third topological sector (Nieh–Yan)	Established (§3.2)	$\mathcal{N} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$; closed, exact, P-odd
Anomaly inflow works for all three sectors	Established (§3.4)	Three independent descent chains; Callan–Harvey mechanism
θ -parameters are S^1 holonomies	Derived (§3.4.3)	$\theta_i = \oint_{S^1} \phi_i dy$; geometric, not free
Gauged P/CP forces $\theta_{\text{QCD}} = \theta_{\text{grav}} = \theta_{\text{NY}} = 0$	Proven (§3.5)	Orientation reversal flips all three; Elitzur's theorem
Nieh–Yan cutoff dependence is non-fatal	Discussed (§3.6)	Gauging is structural; 5D provides natural cutoff; torsion is dynamical
CEH instanton bundles three θ -parameters	Derived (§3.7)	$Q_{\text{gauge}}, Q_{\text{grav}}, Q_{\text{NY}} \neq 0$ on CEH background

Axion quality problem is three-fold	Derived (§3.7.4)	Axion couples only to $\text{Tr}(F \wedge F)$, not to $\text{Tr}(R \wedge R)$ or $\mathcal{N}$
Mass gap and $\theta = 0$ are decoupled	Established (§3.4.1)	Different sectors of same 5D action
5D theory does double duty	Main result	Proves $\Delta > 0$ (dynamical) and sets $\theta = 0$ (topological)

The main theorem

Theorem (Three-sector anomaly inflow). The 5D gauge theory with torsion on $\mathbb{R}^{3,1} \times S^1$ ([B]) is the anomaly-inflow bulk that gauges the three Chern–Weil (-1) -form global symmetries generated by $\text{Tr}(F \wedge F)$, $\text{Tr}(R \wedge R)$, and $\mathcal{N}$. The S^1 holonomies are the θ -parameters; gauged P/CP (orientation reversal of S^1) forces all three to zero. The mass gap (proved from the dynamical sector) and the $\theta = 0$ result (proved from the topological sector) are decoupled consequences of the same 5D theory.

4. Spectral Diagnostics: Floer Homology, Random Matrix Theory, and the Entanglement Gap

4.1 Introduction: Three spectral windows on the mass gap

The previous two sections established two complementary identifications. §2 showed that the topological susceptibility χ is the Fisher information metric on θ -space, with the mass gap Δ controlling the finiteness and magnitude of the QFI. §3 showed that the 5D gauge theory with torsion is the anomaly-inflow bulk for three topological sectors, with the mass gap (proved from the dynamical sector) and the $\theta = 0$ result (proved from the topological sector) being decoupled consequences of the same theory.

Both sections rely on $\Delta > 0$ but neither *determines* Δ from an independent observable. In this section, we exhibit three independent **spectral diagnostics** — each rooted in a different mathematical structure — that read off the mass gap from quantities computable on the lattice or measurable in principle. Each diagnostic connects a different face of [A]'s topological framework to [B]'s constructive-experimental program:

1. **Floer homology** (§4.1): the mass gap is the spectral gap of the Floer differential on $\mathcal{A} \wedge \mathcal{G}$ — a structural identification linking [A]'s index theorem to [B]'s Sturm–Liouville operator.
2. **Random matrix theory** (§4.2): the Dirac eigenvalue statistics on the lattice encode the spectral gap controlling Result A — a quantitative diagnostic using [B]'s lattice tools.

3. **Bisognano–Wichmann theorem** (§4.3): the mass gap controls entanglement entropy decay — a non-spectral determination of Δ connecting to $[A]$'s island formula.

These three diagnostics are independent: they probe the mass gap through (i) the topology of gauge-field space, (ii) the spectrum of the Dirac operator, and (iii) the entanglement structure of the vacuum. Their agreement would constitute strong evidence for $\Delta > 0$ from three unrelated mathematical structures.

4.2 Floer homology and the mass gap as spectral gap

4.2.1 The Floer chain complex of the Yang–Mills functional

Floer homology (Floer, 1988) is an infinite-dimensional analogue of Morse homology, constructed on the space of gauge connections modulo gauge transformations $\mathcal{A}/\mathcal{G}$. The **critical points** of the Yang–Mills functional $S_{\text{YM}}[A] = \frac{1}{2g^2} \int \text{Tr}(F \wedge *F)$ on a 3-manifold Y (the spatial slice) are the flat connections (minima, $F = 0$) and the instantons (saddles, $*F = \pm F$). The **Floer differential** $\partial: CF_k \rightarrow CF_{k-1}$ counts gradient-flow lines — anti-self-dual instantons on $Y \times \mathbb{R}$ connecting critical points of adjacent index:

$$\partial: CF_k \rightarrow CF_{k-1}, \quad \partial[a] = \sum_{[b]: \mu([b]) = k-1} \#(\mathcal{M}([a], [b])),$$

where $\mathcal{M}([a], [b])$ is the moduli space of anti-self-dual connections on $Y \times \mathbb{R}$ asymptotic to $[a]$ at $t \rightarrow -\infty$ and $[b]$ at $t \rightarrow +\infty$, μ is the **Maslov index** (the Floer grading), and $\#(\mathcal{M})$ counts solutions modulo the $\mathbb{R}$ -translation. The Floer homology groups $HF_*(Y) = \ker \partial / \text{im } \partial$ are topological invariants of Y .

Atiyah (1988) and Donaldson (1983) established that Floer homology for the Yang–Mills functional is well-defined: the instanton moduli spaces are compact (after Uhlenbeck compactification) in the relevant directions, the differential satisfies $\partial^2 = 0$ (when the moduli spaces have the expected dimension, i.e., when no exotic phenomena obstruct the gluing), and the resulting homology is independent of the metric on Y .

4.2.2 The spectral gap of the Floer differential

The Floer differential ∂ is an operator on the chain complex CF_* . Its **spectral gap** — the smallest nonzero eigenvalue of the associated "Floer Laplacian" $\partial^* \partial + \partial \partial^*$ — measures the energy barrier between gauge-inequivalent topological sectors. This spectral gap is precisely the mass gap Δ .

To see this, consider the gradient flow equation on $\mathcal{A}/\mathcal{G}$:

$$\frac{\partial A}{\partial t} = -\frac{\delta S_{\text{YM}}}{\delta A} = -*_A F$$

This is the **anti-self-dual Yang–Mills equation** $F^+ = 0$ on $Y \times \mathbb{R}$. Linearizing around a critical point $[A_*]$ (a flat connection or an instanton), the fluctuation operator is:

$$\hat{\mathcal{O}}_{\text{Floer}} = -*_{A_*} D_{A_*} + \text{curvature terms}$$

This is a self-adjoint elliptic operator on $\mathcal{A}/\mathcal{G}$ (after gauge-fixing). Its spectrum $\{\lambda_k\}$ determines the **linearized gradient flow** near $[A_*]$. The spectral gap — the

smallest λ_k for $k \neq 0$ — is the energy cost of the lowest excitation above the vacuum. This is the mass gap Δ .

Key point: The Floer spectral gap is *not* the gap of a single operator on a single space. It is the gap of the *entire chain complex* — the minimum over all critical points of the smallest nonzero eigenvalue of the linearized flow. This is a global quantity on $\mathcal{A} \wedge \mathcal{G}$, not a local one.

4.2.3 Identification with [B]'s Sturm–Liouville operator

[B] (arXiv:2506.00284, §2.2.1) proves the mass gap by analyzing a Sturm–Liouville operator in the 5D theory:

$$\hat{O} = -\frac{d^2}{dz^2} + V(z), \quad z \in [0, L]$$

where z is the coordinate on the compactified fifth dimension S^1 and $V(z)$ is the effective potential arising from the gauge-field dynamics. The Sturm–Liouville analysis shows that $\hat{O}$ has a discrete spectrum μ_n^2 with $\mu_0^2 = 0$ (the zero mode) and $\mu_1^2 = \Delta^2 > 0$ (the first excited state). [B] proves $\mu_1^2 > 0$ using a combination of:

- **Polymer expansion** (convergence of the cluster expansion at strong coupling, §2.3),
- **Sturm–Liouville oscillation theory** (counting eigenvalues via nodal points, §2.2.2),
- **Lattice regularization** (discretizing $\hat{O}$ and proving the gap persists in the continuum limit, §2.4).

The **identification** is:

$$\hat{\mathcal{O}} \xrightarrow{\text{Floer}} \hat{O}_{\text{SL}}$$

Both operators are self-adjoint, elliptic, and arise from linearizing the same physical problem — the stability of the Yang–Mills vacuum against topological fluctuations. The Floer operator acts on the infinite-dimensional space $\mathcal{A} \wedge \mathcal{G}$; the Sturm–Liouville operator acts on the 1D interval $[0, L]$ (the compactified fifth dimension). The Sturm–Liouville operator is a **dimensional reduction** of the Floer operator: the 5D theory's compactification S^1 reduces the infinite-dimensional Floer problem to a 1D spectral problem.

More precisely, the 5D compactification decomposes the gauge field as $A(x, z) = \sum_n \phi_n(x) \psi_n(z)$, where ψ_n are eigenfunctions of $\hat{O}_{\text{SL}}$. The lowest mode ψ_0 is the 4D gauge field; the first excited mode ψ_1 has mass $\mu_1 = \Delta$. The Floer spectral gap — the energy barrier on $\mathcal{A} \wedge \mathcal{G}$ — is realized as the Sturm–Liouville gap — the mass of the first Kaluza–Klein mode.

4.2.4 Identification with [A]'s index theorem

[A]'s central object is the Atiyah–Singer index theorem (§5.5):

$$\text{index}(D) = \dim \ker D - \dim \text{coker } D = \text{Tr}(R) \cdot Q_{\text{gauge}} + \dim(R) \cdot \hat{A}(M) - \frac{\eta}{2}$$

where D is the Dirac operator coupled to the gauge and gravitational backgrounds. The index counts the net number of instantons (topological charge Q_{gauge}) weighted by the gravitational background ($\hat{A}(M)$) and the boundary contribution ($\eta/2$).

The Floer chain complex CF_* has an Euler characteristic:

$$\chi(HF) = \sum_k (-1)^k \dim HF_k = \text{index}(D)$$

This is the **Atiyah–Singer index theorem in Floer-theoretic form**: the Euler characteristic of the Floer complex is the index of the Dirac operator. [A] computes $\text{index}(D)$; Floer homology organizes the same data into a chain complex. The two are computing different invariants of the *same* object:

- **[A] (index theorem)**: counts the net number of zero modes of D — a single integer.
- **[B] (Sturm–Liouville)**: computes the full spectrum $\{\mu_n^2\}$ of $\hat{O}$ — the spectral gap $\mu_1^2 = \Delta^2$ is the first excited state.
- **Floer homology**: organizes the critical points of S_{YM} into a chain complex with a spectral gap — the mass gap Δ is the spectral gap of the complex.

The three are related by:

$$\underbrace{\text{index}(D)}_{\text{[A]}} = \underbrace{\chi(HF)}_{\text{Floer}} \quad \text{and} \quad \underbrace{\Delta^2 = \mu_1^2}_{\text{[B]}} = \underbrace{\text{spectral gap of } \partial^{\dagger}}_{\text{Floer}}$$

4.2.5 The three-graded chain complex in the torsionful case

In the 5D theory with torsion (§3), the topological structure has three sectors: gauge ($\text{Tr}(F \wedge F)$), gravitational ($\text{Tr}(R \wedge R)$), and Nieh–Yan ($\mathcal{N}$).

Correspondingly, the Floer chain complex acquires a **three-fold grading**:

$$SCF_k = \bigoplus_{(Q_{\text{gauge}}, \sigma, Q_{\text{NY}})} CF_k^{(Q_{\text{gauge}}, \sigma, Q_{\text{NY}})}$$

where $Q_{\text{gauge}} \in \mathbb{Z}$ is the gauge topological charge, $\sigma \in \mathbb{Z}$ is the gravitational instanton number (the signature $\sigma(M)$), and $Q_{\text{NY}} \in \mathbb{Z}$ is the Nieh–Yan charge. The three gradings correspond to the three descent chains of §3.4. The Floer differential respects this decomposition (the three sectors are independent at the level of the chain complex), and the spectral gap is the minimum over all three sectors:

$$\Delta^2 = \min(\Delta_{\text{gauge}}^2, \Delta_{\text{grav}}^2, \Delta_{\text{NY}}^2)$$

This three-fold decomposition is the Floer-theoretic avatar of the three-sector anomaly inflow of §3. The mass gap in each sector is controlled by the corresponding Chern–Weil form, and the total mass gap is the minimum. The gauged P/CP (§3.5) forces all three θ -parameters to zero, but the mass gap is a dynamical quantity independent of the θ -parameters (§3.4.1, decoupling).

4.2.6 Structural insight, not computational tool

We emphasize that Floer homology, in its current mathematical development, is a **structural insight** rather than a computational tool for extracting Δ numerically. The Floer chain complex is well-defined for specific 3-manifolds Y (e.g., homology 3-spheres), but the Yang–Mills mass gap problem concerns $Y = \mathbb{R}^3$ (or T^3 for finite-volume regularization), where the Floer theory is less developed. The value of the Floer perspective is conceptual:

- It **unifies** [A]’s index theorem and [B]’s Sturm–Liouville analysis as two representations of the same chain complex.
- It **explains** why the mass gap is a topological invariant (it is the spectral gap of a chain complex whose Euler characteristic is the index).

- It **predicts** that the mass gap is the same in all three sectors (gauge, gravitational, Nieh–Yan) — a consequence of the three-fold grading.

Recent work (arXiv:1312.1350) has explored the connection between Floer homology, the Yang–Mills mass gap, the glueball spectrum, and the conformal window in large- N QCD, providing a roadmap for making the Floer perspective more computational.

4.2.7 The Donaldson–Witten topological field theory

The bridge from Floer homology to physics is mediated by the **Donaldson–Witten topological Yang–Mills theory** (Witten, 1988). This is a topological twist of 4D $N=2$ supersymmetric Yang–Mills theory, whose observables are the Donaldson invariants of the underlying 4-manifold. The key structural result is:

- The **Donaldson invariants** (polynomial functions on the instanton moduli space) are the physical observables of the TQFT.
- The **Floer homology** of the 3-manifold Y is the **Hilbert space** of the TQFT quantized on $Y \times \mathbb{R}$.
- The **Hamiltonian** of the TQFT is the Floer differential (up to the BRST exact terms): $H_{\text{TQFT}} = \{Q_{\text{BRST}}, \cdot\}$, and the physical states are $H\!F^*(Y) = \ker \partial_{\text{im}}, \partial_{\text{partial}}$.

The spectral gap of H_{TQFT} is the mass gap Δ . In the TQFT, the gap is exact (the theory is topological), but in the *physical* (non-twisted) Yang–Mills theory, the gap is the same quantity — the energy of the lowest excitation above the vacuum. The Donaldson–Witten TQFT provides the bridge from the mathematical Floer homology to the physical Yang–Mills spectrum.

4.3 Random matrix theory and the Banks–Casher diagnostic

4.3.1 The Banks–Casher relation

The **Banks–Casher relation** (Banks & Casher, 1980) connects the topological susceptibility to the Dirac eigenvalue density:

$$\chi = \frac{\Sigma^2}{N_f} \lim_{Q \rightarrow 0} \frac{1}{Q} \frac{d}{dQ} \Sigma^2 = \frac{\pi}{\rho(0)}$$

where Σ is the chiral condensate, $\rho(\lambda)$ is the spectral density of the Dirac operator at $\lambda = 0$ (per unit volume), and N_f is the number of light quark flavors. In the quenched approximation ($N_f = 0$, pure Yang–Mills), the relation simplifies to:

$$\chi = \pi \rho_{\text{quenched}}(0)$$

where $\rho_{\text{quenched}}(0)$ is the Dirac spectral density at the origin in the pure gauge theory.

Key observation: In a confining theory with a mass gap $\Delta > 0$, the Dirac operator in the **quenched** (pure YM) theory has a **spectral gap**: the eigenvalues are bounded away from zero, $\rho_{\text{quenched}}(\lambda) = 0$ for $|\lambda| < \Delta_{\text{Dirac}}$, where Δ_{Dirac} is the Dirac spectral gap. The Banks–Casher relation then gives $\chi = \pi \cdot 0 = 0$ in the strict quenched limit — but this is the *unquenched* relation. In the quenched theory, χ is computed directly from the gauge-field topological susceptibility ($[A]$), not through the Dirac

spectrum.

The bridge is more subtle and goes through the **unquenched** theory with dynamical quarks, where the Dirac spectrum *does* have eigenvalues near zero (the chiral condensate is nonzero), and the Banks–Casher relation holds. The mass gap Δ (the glueball mass) controls the *ultraviolet* behavior of $\rho(\lambda)$, while the *infrared* behavior (near $\lambda = 0$) is controlled by the chiral condensate. The two scales are related but distinct.

4.3.2 The connection to the Källén–Lehmann spectral density

[A]’s Result A relates χ to the Källén–Lehmann spectral density $\rho(\mu^2)$ of the topological charge correlator (§2.2):

$$\chi = \int_{\Delta^2}^{\infty} \frac{d\mu^2 \rho(\mu^2)}{\mu^2}$$

The Banks–Casher relation relates χ to the Dirac spectral density $\rho_{\text{Dirac}}(\lambda)$ at $\lambda = 0$. These are two *different* spectral densities:

Quantity	Spectral density	Variable	Domain
[A] / 2 (Källén–Lehmann)	$\rho(\mu^2)$	μ^2 (mass squared)	$\mu^2 \geq \Delta^2$
Banks–Casher	$\rho_{\text{Dirac}}(\lambda)$	λ (Dirac eigenvalue)	$\lambda \in \mathbb{R}$

The **connection** between them goes through the index theorem. The index of the Dirac operator $\text{index}(D) = n_+ - n_-$ counts the difference of zero modes with positive and negative chirality, and equals the topological charge Q_{gauge} . The spectral density of *nonzero* Dirac eigenvalues encodes the gauge-field topology through the index. More precisely, the **density of near-zero modes** $\rho_{\text{Dirac}}(\lambda)$ for $\lambda \sim 0$ is related to the topological susceptibility through:

$$\chi = \pi \rho_{\text{Dirac}}(0) \quad \text{(Banks–Casher, unquenched)}$$

and the **mass gap** Δ controls the **spacing** of the low-lying eigenvalues. In a gapped theory, the Dirac eigenvalues near zero are discrete and well-separated (the Banks–Casher density $\rho_{\text{Dirac}}(0)$ is nonzero but finite). In a gapless theory, the eigenvalues are dense (the density diverges, or the spectrum is continuous).

4.3.3 Random matrix theory: universal predictions for the Dirac spectrum

Random matrix theory (RMT), specifically the **chiral Gaussian Unitary Ensemble** (chGUE), provides **universal predictions** for the low-lying Dirac eigenvalue spectrum in the microscopic limit (eigenvalues of order $1/V$) (Verbaarschot & Wettig, 2000; Damgaard et al., 2010). The chGUE is an ensemble of $N \times N$ anti-Hermitian matrices with chiral structure, and its eigenvalue statistics are:

- **Universal:** the microscopic spectral density $\rho_s(\zeta)$ (where $\zeta = \lambda V \Sigma$ is the rescaled eigenvalue) depends only on the global symmetry pattern (the Dyson index), not on the dynamics.
- **Topological:** the RMT prediction depends on the topological charge Q — in the sector $Q =$

k , the lowest eigenvalue distribution is $P_k(\zeta_{\min})$, which differs from the $Q = 0$ sector.

- **Mass-gap sensitive:** the **eigenvalue spacing** $\lambda_{n+1} - \lambda_n$ near zero is set by $1/(V\Sigma)$, and the **departure from RMT predictions** at larger eigenvalues signals the mass gap scale Δ . Specifically, the RMT prediction holds for $\lambda \ll \Delta$ (the microscopic regime) and breaks down for $\lambda \gtrsim \Delta$ (the mesoscopic regime), where the dynamics take over.

4.3.4 The diagnostic: measuring Δ from Dirac eigenvalue statistics

The diagnostic works as follows:

1. **Generate gauge-field configurations** on the lattice at a given lattice spacing a and volume $V = L^4$ ([B]'s lattice framework, §2.4).
2. **Compute the Dirac operator spectrum** on each configuration, in each topological sector Q .
3. **Compare the low-lying eigenvalue statistics** to the RMT (chGUE) prediction for the given Q .
4. **Identify the crossover scale** λ_{cross} where the eigenvalue statistics depart from RMT. This crossover scale is related to the mass gap: $\lambda_{\text{cross}} \sim \Delta_{\text{Dirac}} \sim \Delta$ (up to normalization).

The crossover scale is the **mesoscopic–microscopic boundary**: below it, the eigenvalues are universal (RMT); above it, they are dynamical (mass-gap controlled). The mass gap is thus readable from the Dirac spectrum as the scale where universality breaks down.

4.3.5 Connection to the Fisher metric and §2

The Fisher metric identification (§2) gives $\chi = F_Q/(4V)$, and the Banks–Casher relation gives $\chi = \pi, \rho_{\text{Dirac}}(0)$. Combining:

$$F_Q = 4V\pi, \rho_{\text{Dirac}}(0)$$

The quantum Fisher information is proportional to the Dirac spectral density at the origin. The mass gap makes F_Q finite (§2.2.3) by ensuring $\rho(\mu^2)$ is supported on $\mu^2 \geq \Delta^2$. The Banks–Casher relation connects this to the Dirac spectrum: a finite $\rho_{\text{Dirac}}(0)$ (nonzero chiral condensate, gapped theory) gives a finite χ and finite F_Q .

The full chain is:

$$\Delta > 0 \quad \xrightarrow{\text{[A], Result A}} \quad \chi \leq C\Delta^4 \quad \xrightarrow{\text{§2, Fisher}} \quad F_Q = 4V\chi \quad \xrightarrow{\text{Banks–Casher}} \quad \rho_{\text{Dirac}}(0) = \chi/\pi$$

The mass gap, the Fisher information, the topological susceptibility, and the Dirac spectral density are four manifestations of the same quantity, measurable on the lattice through [B]'s tools.

4.3.6 Reanalyzing existing lattice data

A practical consequence: **existing lattice QCD data** can be reanalyzed to extract the mass gap diagnostic without new simulations. The ingredients are:

- **Published Dirac spectra:** lattice collaborations routinely compute the Dirac eigenvalue

spectrum for topological charge measurements. These spectra contain the crossover scale λ_{cross} .

- **RMT comparison:** the chGUE predictions are known analytically for arbitrary Q (Damgaard & Fukaya, 2011). The comparison is a straightforward fit.
- **Extraction of Δ :** the crossover scale $\lambda_{\text{cross}} \sim \Delta$ provides a non-spectroscopic determination of the mass gap, independent of the glueball mass extraction (which requires computing correlation functions of operators with the correct quantum numbers).

The RMT diagnostic is **complementary** to the standard glueball mass extraction:

- **Glueball spectroscopy:** computes Δ from the exponential decay of a correlation function $C(t) \sim e^{-\Delta t}$. Requires good overlap with the glueball operator, which is notoriously difficult (large noise at large t).
- **RMT diagnostic:** computes Δ from the Dirac eigenvalue statistics. No glueball operator needed; the Dirac spectrum is a gauge-invariant quantity. The crossover scale is a statistical property of the spectrum, not a correlation function.

Both methods should agree if the mass gap is indeed a physical quantity (not a lattice artifact). Their agreement would be strong evidence for $\Delta > 0$ from two independent lattice observables.

4.4 The Bisognano–Wichmann theorem and the entanglement gap

4.4.1 The Bisognano–Wichmann theorem

The **Bisognano–Wichmann theorem** (Bisognano & Wichmann, 1975, 1976) is a rigorous result in algebraic quantum field theory (AQFT). It states:

Theorem (Bisognano–Wichmann). For a relativistic QFT satisfying the Wightman axioms, the reduced density matrix of the vacuum state restricted to a **half-space** R (e.g., $x_1 > 0$) is:

$$\rho_R = \frac{e^{-2\pi K_R}}{Z_R}$$

where K_R is the **modular Hamiltonian** — the generator of the Lorentz boost flow that preserves the half-space R — and $Z_R = \text{Tr}(e^{-2\pi K_R})$ is the normalization.

For a Lorentz-invariant vacuum, K_R is **exactly** the boost operator:

$$K_R = \int_{x_1 > 0} d^3x, \left(x_1 T^{00}(x) + x_2 T^{01}(x) \right)$$

This is a **non-approximate** result: the modular Hamiltonian is known exactly, not just for free theories or conformal theories, but for any Lorentz-invariant QFT satisfying the Wightman axioms. The entanglement Hamiltonian (the logarithm of the reduced density matrix) is the boost operator.

4.4.2 The entanglement entropy and the mass gap

The **entanglement entropy** of the half-space R is:

$$S_R = -\text{Tr}(\rho_R \log \rho_R) = 2\pi \langle K_R \rangle + \log Z_R$$

For a CFT (gapless theory), the entanglement entropy follows the **Bekenstein–Hawking law**: S_R

$\sim \text{Area}(\partial R)/(4\epsilon^2)$, where ϵ is a UV cutoff. The entropy is **volume-law** in the UV and divergent.

For a **gapped** theory ($\Delta > 0$), the entanglement entropy of a half-space still has a UV-divergent area-law contribution, but the **universal** part (the finite, cutoff-independent piece) is controlled by the mass gap. Specifically, the corrections to the area law decay exponentially (Calabrese & Cardy, 2004; Casini & Huerta, 2009):

$$S_R = S_{\text{area}} + \frac{c}{\Delta^2} e^{-\Delta R_{\text{eff}}} \cdot (\text{geometric factors})$$

where R_{eff} is the effective size of the entanglement cut (the distance from the boundary of R to the "edge" of the system), and c is a dimensionless constant.

For a **strip** geometry of width R in a gapped QFT, the entanglement entropy is (Casini & Huerta, 2009):

$$S(R) = \frac{c_1}{\epsilon} L^{d-2} - \frac{c_2}{\Delta} L^{d-2} e^{-\Delta R} + O(e^{-2\Delta R})$$

The **leading correction** to the area law is $e^{-\Delta R}$, and the **exponential decay rate** is the mass gap Δ . This is the **entanglement gap diagnostic**: the mass gap is readable from the exponential decay of the entanglement entropy correction.

4.4.3 Non-spectral determination of Δ

The key feature of the entanglement gap diagnostic is that it determines Δ **without computing the spectrum** of any operator. The entanglement entropy $S(R)$ is a property of the **ground state** (the vacuum), not of the excitation spectrum. The diagnostic works as follows:

1. **Prepare the vacuum** on the lattice (or in the cold-atom simulator, [B]'s §3.3).
2. **Compute the reduced density matrix** ρ_R for a strip region of width R , for several values of R .
3. **Compute the entanglement entropy** $S(R) = -\text{Tr}(\rho_R \log \rho_R)$.
4. **Fit** the R -dependence of the subleading term: $S(R) - S_{\text{area}} \propto e^{-\Delta R}$.
5. **Extract** Δ from the exponential decay rate.

This is a **ground-state property**, not a spectroscopic measurement. It does not require creating excitations or measuring correlation functions at large time separations. The entanglement entropy is computed from the vacuum wave function, which is available on the lattice (or in the cold-atom system) as the ground state of the Hamiltonian.

4.4.4 Connection to [A]'s island formula

[A] (§7.5) employs the **island formula** (Penington, 2020; Almheiri et al., 2019) to compute the entanglement entropy in the AdS/CFT context. The island formula states:

$$S(R) = \min_{\text{islands}} \left[\frac{\text{Area}(\partial I)}{4G_N} + S_{\text{bulk}}(R \cup I) \right]$$

where I is the "island" region in the bulk, and the minimization is over all possible islands. The island formula reproduces the Page curve of black hole evaporation and resolves the entropy paradox.

The connection to the Bisognano–Wichmann theorem is:

- **Boundary (CFT) side:** the Bisognano–Wichmann theorem gives the exact modular Hamiltonian for the half-space, and the entanglement entropy is $S = 2\pi \langle K_R \rangle + \log Z_R$. The exponential decay $e^{-\Delta R}$ is a boundary statement.
- **Bulk (gravity) side:** the island formula gives the entanglement entropy in terms of the bulk geometry. The bulk spectral gap $\Delta_{\text{grav}} = \Delta/L$ ([A]'s AdS/CFT identification) controls the bulk entanglement decay. The island boundary ∂I is the bulk analog of the entanglement cut ∂R .
- **The mass gap is the quantity that makes both sides well-defined.** Without a gap ($\Delta = 0$), the entanglement entropy has no exponential decay (the correction is power-law, as in a CFT), and the island formula reduces to the RT/HRT formula without the island contribution. The gap is what makes the island formula nontrivial.

The combined picture:

$$\Delta > 0 \quad \xrightarrow{\text{Bisognano–Wichmann (boundary)}} \quad S(R) - S_{\text{area}} \sim e^{-\Delta R} \quad \xrightarrow{\text{island formula (bulk)}} \quad \Delta_{\text{grav}} = \Delta/L$$

The mass gap controls the entanglement decay on both sides of the AdS/CFT duality, and the two sides are related by the standard $\Delta_{\text{grav}} = \Delta/L$ rescaling.

4.4.5 The modular Hamiltonian and the eta invariant

The modular Hamiltonian K_R (the boost operator) has a **spectrum** — the eigenvalues of K_R on the vacuum. For a gapped theory, this spectrum has a gap (the "modular gap"). The **spectral asymmetry** of K_R — the imbalance between positive and negative eigenvalues — is measured by the **eta invariant** $\eta(K_R)$, which is the APS boundary term in [A]'s index theorem (§5.5):

$$\text{index}(D) = T(R) \cdot Q_{\text{gauge}} + \dim(R) \cdot \hat{A}(M) - \frac{\eta(K_R)}{2}$$

The eta invariant of the modular Hamiltonian on the entanglement cut ∂R is the **boundary correction** to the index theorem. In the 5D theory with torsion (§3), the index theorem has three sectors, and the eta invariant receives contributions from all three:

$$\eta(K_R) = \eta_{\text{gauge}}(K_R) + \eta_{\text{grav}}(K_R) + \eta_{\text{NY}}(K_R)$$

The entanglement gap is a **combined gauge + gravitational + Nieh–Yan** quantity. The modular Hamiltonian's spectral asymmetry, measured by the three-sector eta invariant, is a topological invariant of the entanglement cut, and the mass gap ensures it is finite (the spectral asymmetry would diverge in a gapless theory where the modular spectrum is continuous).

4.4.6 The entanglement spectrum and the Floer connection

The **entanglement spectrum** — the eigenvalues of K_R (or equivalently, $-\log \rho_R$) — is a physical observable in condensed matter (Li & Haldane, 2008). In a gapped QFT, the entanglement spectrum is discrete (the modular Hamiltonian has a gap), and the lowest entanglement eigenvalue is set by Δ .

The connection to Floer homology (§4.2) is through the **Donaldson–Witten TQFT** (§4.2.7). In the TQFT, the Hilbert space on Y is $\mathcal{H}F_*(Y)$, and the Hamiltonian is the Floer differential. The entanglement spectrum of the TQFT vacuum (restricted to a subregion) is the spectrum of the Floer

Laplacian restricted to the subregion — the same operator whose spectral gap is Δ . The entanglement gap and the Floer spectral gap are the **same quantity**, viewed from two perspectives:

- **Floer (§4.2):** the spectral gap of $\partial^\dagger$ on $\mathcal{A}/\mathcal{G}$.
- **Entanglement (§4.4):** the gap in the entanglement spectrum of K_R .

This is a deep connection: the mass gap is simultaneously a topological invariant (Floer), a spectral property (Sturm–Liouville / RMT), and an entanglement property (Bisognano–Wichmann). The three diagnostics of §4 are not independent — they are three windows on the same quantity.


4.5 Summary of Section 4

The three diagnostics and their connections

Diagnostic	Mathematical structure	[A] connection	[B] connection	What it measures
Floer homology (§4.2)	Chain complex on $\mathcal{A}/\mathcal{G}$	Index theorem = $\chi(\text{HF})$	Sturm–Liouville = spectral gap of $\partial^\dagger$	Structural: Δ = spectral gap of the Floer complex
RMT / Banks–Casher (§4.3)	Dirac eigenvalue statistics	$\chi = \lim_{\mu \rightarrow 0} \frac{1}{\rho(\mu^2)}$	Lattice Dirac spectrum	Quantitative: Δ = crossover scale of Dirac eigenvalue statistics
Entanglement gap (§4.4)	Bisognano–Wichmann modular Hamiltonian	Island formula; η invariant	Lattice ground state; cold-atom simulator	Non-spectral: Δ = exponential decay rate of entanglement entropy

The unifying picture

The three diagnostics form a closed triangle:



All three vertices compute the same quantity Δ :

1. **Floer** gives the structural reason: Δ is the spectral gap of the chain complex whose Euler characteristic is the index.

2. **RMT** gives the quantitative diagnostic: Δ is readable from the Dirac eigenvalue statistics on the lattice.
3. **Entanglement** gives the non-spectral diagnostic: Δ is the exponential decay rate of the entanglement entropy correction.

The Fisher metric (§2) and the anomaly inflow (§3) provide the framework: χ (the Fisher metric) is the bridge between the topological ([A]) and the spectral ([B]), and the 5D theory with torsion unifies the three topological sectors. The three diagnostics of §4 provide the **observational handles** — the ways to *measure* Δ from three independent physical quantities.

Consistency check

If the mass gap $\Delta > 0$ is a genuine physical quantity (not a lattice artifact or a mathematical fiction), then the three diagnostics must agree:

- **Floor spectral gap** = Δ (by definition).
- **RMT crossover scale** $\lambda_{\text{cross}} \sim \Delta$ (from the Dirac spectrum).
- **Entanglement decay rate** $\sim \Delta$ (from the entanglement entropy).

The agreement of these three independently computed quantities would constitute strong evidence for the mass gap, from three unrelated mathematical structures (topology, random matrix theory, and entanglement theory). Conversely, a disagreement would signal that one of the diagnostics is misidentified or that the mass gap is not a simple scalar (e.g., it could be direction-dependent on $\mathcal{A} \wedge \mathcal{G}$, or sector-dependent in the three-sector theory).

5. The CEH Instanton and Resurgence: A Research Program

5.1 The central open calculation of [A]

[A] identifies, in §8.7, the **one-loop fluctuation determinant around the colored Eguchi–Hanson (CEH) instanton** as the decisive open calculation in the combined gauge + gravitational theory. This determinant determines whether the CEH amplitude K_{CEH} is:

- **Gauge-scale** ($K_{\text{CEH}} \sim \Lambda_{\text{QCD}}^4$): the CEH instanton contributes at the QCD scale, making the θ -term physically relevant. In this scenario, the axion solution to Strong CP fails (the axion does not couple to the gravitational sector), and the gauged P/CP mechanism is the only way to set $\theta = 0$.
- **Planck-suppressed** ($K_{\text{CEH}} \sim \Lambda_{\text{QCD}}^4 / M_{\text{Pl}}^n$): the CEH instanton is screened by quantum gravity, and the axion solution suffices to set $\theta_{\text{QCD}} = 0$ (though not $\theta_{\text{grav}} = 0$ or $\theta_{\text{NY}} = 0$).

[A]'s §8.9 presents this as a **decision table**: the axion vs. gauged P/CP debate hinges on the CEH one-loop measure, which is currently uncomputed.

In §3.7, we showed that the CEH instanton **bundles all three θ -parameters** (θ_{QCD} , θ_{grav} , θ_{NY}) simultaneously, making the three-sector problem inseparable. The CEH measure problem is therefore not merely the gauge vs. gravity question of [A] — it is a **three-sector** question. The gauged P/CP mechanism (§3.5) resolves it structurally by forcing all three θ -parameters to zero, but the *dynamical* question — whether the CEH saddle contributes at gauge scale — remains open.

This section reframes the CEH measure problem as a **Borel-plane singularity analysis** using resurgence theory, and proposes a concrete research program to make it computable. We also address, via the same framework, the **5D loophole** ([B]'s use of a 5D theory to prove a 4D result) and the **Lefschetz thimble decomposition** as a lattice-computable realization.

5.2 Resurgence and trans-series: the framework

5.2.1 Trans-series structure

Resurgence theory (Écalle, 1981; revived in QFT by Ünsal, Mariño, Dumitrescu, 2010s) provides a unified framework for perturbative and non-perturbative contributions through **trans-series** — formal series that include both powers and exponentials. For a generic observable $f(g^2)$ in a theory with instanton action S_0 :

$$f(g^2) = \sum_{k=0}^{\infty} a_k g^{2k} + \sigma e^{-S_0/g^2} \sum_{k=0}^{\infty} b_k g^{2k} + \sigma^2 e^{-2S_0/g^2} \sum_{k=0}^{\infty} c_k g^{2k} + \cdots$$

The first term is the **perturbative series** (around the trivial vacuum). The second term is the **one-instanton** contribution (around the instanton saddle, action S_0). The third is the **two-instanton** (instanton–anti-instanton) contribution. The trans-series parameter σ is related to the instanton density and is fixed by boundary conditions or renormalization conditions.

The key result of resurgence is the **large-order/low-order duality**: the perturbative coefficients a_k at large order $k \gg 1$ are controlled by the instanton fluctuations b_k at low order:

$$a_k \sim \frac{S_0^{-k}}{2\pi i} \sum_{\text{saddles}} \frac{b_0(\sigma)}{(k-1)!} \left(k!, S_0^{-k} \right) \quad \text{as } k \rightarrow \infty$$

This means that **Borel ambiguities in the perturbative series are exactly cancelled by instanton contributions**, and vice versa. The trans-series, when properly Borel-resummed along the correct lateral direction, gives an unambiguous result. This cancellation has been verified rigorously in 2D models: $\mathbb{CP}^{N-1}$, the principal chiral model, and the Gross–Neveu model (Dunne & Ünsal, 2014; Mariño, 2012).

5.2.2 The Borel plane

The **Borel transform** of the perturbative series $f_{\text{pert}}(g^2) = \sum a_k g^{2k}$ is:

$$\hat{f}(t) = \sum_{k=0}^{\infty} \frac{a_k}{k!} t^k$$

The Borel-resummed function is:

$$f_{\text{Borel}}(g^2) = \int_0^{\infty} dt e^{-t/g^2} \hat{f}(t)$$

The Borel plane (t -plane) has **singularities** at the locations of the non-perturbative saddles:

- **Borel singularity at $t = S_0$** : the one-instanton saddle. The nature of the singularity (pole,

branch point) encodes the fluctuation determinant.

- **Borel singularity at $t = 2S_0$** : the two-instanton saddle.
- **Borel singularity at $t = nS_0$** : the n -instanton saddle.

If the Borel transform $\hat{f}(t)$ is analytic along the positive real axis (no singularities on the integration contour), the Borel resummation is unambiguous. If there are singularities on the contour (the **Borel ambiguity**), the lateral Borel resummations $f^{\pm}$ (above and below the singularity) differ by an amount proportional to the instanton contribution, and the trans-series cancellation restores ambiguity-freeness.

5.2.3 Resurgence in 4D Yang–Mills: the state of the art

The application of resurgence to 4D Yang–Mills is still in its infancy:

- **2D models** ($\mathbb{CP}^{N-1}$, PCM, Gross–Neveu): trans-series computed, Borel-plane structure mapped, large-order/low-order duality verified (Dunne & Ünsal, 2014; Cherman et al., 2014). These models have instantons and anti-instantons with well-defined fluctuation determinants.
- **4D $\mathcal{N}=2$ SYM** (Seiberg–Witten theory): the instanton calculus is exact (Nekrasov's partition function), and the Borel-plane structure is controlled by the Seiberg–Witten curve. Resurgence has been verified in the $\mathcal{N}=2$ context (Mariño, 2012).
- **4D $\mathcal{N}=1$ and pure YM**: no exact instanton calculus is available. The perturbative series for the vacuum energy, the glueball mass, or the topological susceptibility is known only to a few orders, and the instanton calculus is non-supersymmetric (no localization). The trans-series for pure 4D YM is **unknown**.

This is a genuine obstacle, but not a fatal one. The trans-series structure is constrained by the Borel-plane singularity pattern, which is determined by the **saddle structure of the path integral** — the instantons and their generalizations. Even without computing the full trans-series, one can:

1. **Identify the saddles** (instantons, instanton–anti-instanton pairs, CEH instantons, gravitational instantons).
2. **Locate the Borel singularities** (at $t = S_{\text{saddle}}$ for each saddle).
3. **Determine the nature of each singularity** (pole, branch point, essential singularity) from the fluctuation determinant around the saddle.
4. **Check for ambiguity cancellation** (lateral Borel resummations differ by instanton terms).

Steps 1–2 are known for YM. Step 3 is the open calculation for the CEH instanton. Step 4 is the resurgence conjecture.

5.3 The CEH instanton as a Borel-plane singularity

5.3.1 The saddle structure

The CEH instanton ($[A]$, §8) is a solution of the combined Yang–Mills + gravity equations that simultaneously carries:

- **Gauge topological charge** $Q_{\text{gauge}} = \pm 1$ (from the colored Eguchi–Hanson

gauge bundle),

- **Gravitational topological charge** $\sigma = \pm 1$ (from the Eguchi–Hanson gravitational instanton, the signature),
- **Nieh–Yan charge** $Q_{\text{NY}} = \pm 1$ (from the torsion, in the 5D theory with torsion of §3).

Its action is:

$$S_{\text{CEH}} = \frac{8\pi^2}{g^2} + \frac{\pi^2}{\kappa^2} + S_{\text{NY}}$$

where the first term is the gauge instanton action, the second is the gravitational instanton action, and the third is the Nieh–Yan contribution (which depends on the torsion coupling). In the limit $\kappa \rightarrow 0$ (gravity decoupled), the CEH reduces to the ordinary gauge instanton with action $8\pi^2/g^2$.

5.3.2 Borel-plane singularities of the combined theory

The combined gauge + gravitational + Nieh–Yan theory has the following saddle structure in the path integral:

Saddle	Action	Borel singularity	What it encodes
Trivial vacuum	$S = 0$	$t = 0$ (origin)	Perturbative series
Gauge instanton	$S = 8\pi^2/g^2$	$t = 8\pi^2$	Gauge topological sector
Gravitational instanton	$S = \pi^2/\kappa^2$	$t = \pi^2/\kappa^2 \sim M_{\text{Pl}}^2$	Gravitational topological sector
CEH instanton	$S = 8\pi^2/g^2 + \pi^2/\kappa^2 + S_{\text{NY}}$	$t = S_{\text{CEH}}$	Three-sector bundling
Instanton–anti-instanton	$S = 2 \times 8\pi^2/g^2$	$t = 16\pi^2$	Two-instanton sector

The **CEH Borel singularity** is at $t = S_{\text{CEH}}$. Its nature — whether it is a simple pole (non-degenerate saddle, analytic fluctuation determinant) or an essential singularity (degenerate saddle, non-analytic determinant) — determines whether K_{CEH} is gauge-scale or Planck-suppressed.

5.3.3 The gauge-scale vs. Planck-suppressed question on the Borel plane

[A]'s open question — "is K_{CEH} gauge-scale or Planck-suppressed?" — translates to a Borel-plane question:

Gauge-scale scenario: The CEH fluctuation determinant is analytic at the saddle. The Borel singularity at $t = S_{\text{CEH}}$ is a **simple pole** (or a branch point with finite residue). The one-loop correction $\sum_k b_k g^{2k}$ is a well-defined asymptotic series, and the trans-series term $e^{-S_{\text{CEH}}} \sum_k b_k g^{2k}$ contributes at gauge scale. In this case, $K_{\text{CEH}} \sim$

Λ_{QCD}^4 and the CEH instanton is physically relevant.

Planck-suppressed scenario: The CEH fluctuation determinant has a singularity at the Planck scale (e.g., a UV divergence in the one-loop determinant that requires a Planck-scale counterterm). The Borel singularity at $t = S_{\text{CEH}}$ acquires an **additional singularity** at $t = M_{\text{Pl}}^2 \rho^2$ (where ρ is the instanton size), which screens the gauge-scale contribution. The trans-series term is effectively $e^{-S_{\text{CEH}}} \times (\text{Planck suppression})$, and $K_{\text{CEH}} \sim \Lambda_{\text{QCD}}^4 / M_{\text{Pl}}^n$.

The **Borel-plane singularity structure** thus encodes the answer to [A]'s central question. The trans-series for the combined theory has two relevant singularities:

$$\hat{f}(t) \sim \frac{R_{\text{gauge}}}{t - S_{\text{gauge}}} + \frac{R_{\text{CEH}}}{t - S_{\text{CEH}}} + \frac{R_{\text{Planck}}}{t - M_{\text{Pl}}^2 \rho^2} + \cdots$$

where R_{gauge} , R_{CEH} , R_{Planck} are the residues (fluctuation determinants). The question is whether R_{CEH} is of order 1 (gauge-scale) or of order $1/M_{\text{Pl}}^n$ (Planck-suppressed). This is a **computable** question: one computes the one-loop fluctuation determinant around the CEH saddle and reads off the residue.

5.3.4 The three-sector generalization

In the three-sector theory (§3), the Borel plane has **three** sets of singularities, one for each topological sector:

$$\hat{f}(t) \sim \frac{R_{\text{gauge}}}{t - S_{\text{gauge}}} + \frac{R_{\text{grav}}}{t - S_{\text{grav}}} + \frac{R_{\text{NY}}}{t - S_{\text{NY}}} + \frac{R_{\text{CEH}}}{t - S_{\text{CEH}}} + \cdots$$

The CEH singularity is special: it is the **only** singularity that bundles all three sectors. Its residue R_{CEH} encodes the combined gauge + gravitational + Nieh–Yan fluctuation determinant. The three-sector θ -parameters are:

$$\theta_i = \oint_{S^1} \phi_i dy, \quad i \in \{\text{gauge, grav, NY}\}$$

and the CEH saddle contributes $e^{i(\theta_{\text{gauge}} Q_{\text{gauge}} + \theta_{\text{grav}} \sigma + \theta_{\text{NY}} Q_{\text{NY}})} \times R_{\text{CEH}}$ to the vacuum energy. The gauged P/CP (§3.5) sets all $\theta_i = 0$, which eliminates the CEH contribution to the vacuum energy — but does **not** eliminate the CEH saddle from the path integral. The saddle still contributes to the trans-series fluctuations; it is the θ -dependent *phase* that is set to zero, not the θ -independent *amplitude* R_{CEH} .

This is the decoupling (§3.4.1) in Borel-plane language: the mass gap (dynamical sector) and the $\theta = 0$ result (topological sector) are independent. The CEH amplitude R_{CEH} is a dynamical quantity (it depends on the fluctuation determinant, which depends on the mass gap through the spectral density); the CEH phase (the θ -dependence) is a topological quantity (it depends on the θ -parameters, which are set to zero by gauged P/CP).

5.4 The 5D loophole and Borel summability

5.4.1 The 5D loophole

[B]'s constructive proof (arXiv:2506.00284, §2.2.3) uses a 5D gauge theory on $\mathbb{R}^{3,1}$

$\times S^1$ and takes the compactification radius $R \rightarrow 0$ to recover 4D Yang–Mills. The "5D loophole" is the concern that the 5D $\rightarrow$ 4D limit may not be smooth — the compactification introduces a tower of Kaluza–Klein (KK) modes whose contributions may not vanish as $R \rightarrow 0$, potentially invalidating the 4D mass gap proof.

5.4.2 The trans-series structure of the compactified theory

The 5D theory compactified on S^1 of radius R has a spectrum of KK modes with masses $m_n = n/R$ for $n = 0, 1, 2, \dots$. The $n = 0$ mode is the 4D gauge field; the $n \neq 0$ modes are the KK tower. In the trans-series for an observable $f(g^2, R)$:

$$f(g^2, R) = f_0(g^2) + \sum_{n=1}^{\infty} f_n(g^2) e^{-n/(R\Lambda)} + \dots$$

where $f_0(g^2)$ is the 4D (zero-mode) contribution and the f_n are exponentially suppressed KK corrections. The exponential suppression $e^{-n/(R\Lambda)}$ arises because the KK modes have mass n/R and their contributions to low-energy observables are suppressed by $e^{-m_n L}$ where $L \sim 1/\Lambda$ is the IR scale.

5.4.3 Borel summability of the 5D $\rightarrow$ 4D limit

The question is whether the KK corrections are **Borel-summable** — i.e., whether the 4D limit $R \rightarrow 0$ is unambiguous. The Borel transform of the R -dependence has singularities at:

- $t = 1/(R\Lambda)^2$: the one-KK-mode singularity.
- $t = n/(R\Lambda)^2$: the n -KK-mode singularity.

As $R \rightarrow 0$, these singularities recede to $t \rightarrow \infty$, and the Borel plane becomes increasingly free of singularities on the positive real axis. In the limit $R \rightarrow 0$, the Borel transform is entire (no singularities), and the Borel resummation is unambiguous.

This is the **resurgence argument for the 5D loophole**: the 5D $\rightarrow$ 4D compactification is Borel-summable because the KK-mode singularities retreat to infinity. The 4D limit is unambiguous, and [B]'s mass gap proof (if valid in 5D) survives the compactification.

The subtlety is that the 5D instantons (instantons in the 5D theory) have a different action than the 4D instantons. The 5D instanton action is $S_0^{(5D)} = 8\pi^2/g_5^2$ where $g_5^2 = g^2 R$ (the 5D coupling is dimensionful). The 4D instanton action $S_0 = 8\pi^2/g^2$ is recovered in the $R \rightarrow 0$ limit. The Borel singularity at $t = S_0^{(5D)} = 8\pi^2/(g^2 R)$ also retreats to infinity as $R \rightarrow 0$, confirming that the 4D Borel plane has the correct singularity at $t = 8\pi^2/g^2$.

5.4.4 Large- N volume independence as a cross-check

The Eguchi–Kawai (1982) volume independence provides an independent cross-check. At $N = \infty$, the 5D theory compactified on S^1 of radius R is equivalent to the 4D theory for $R \rightarrow 0$, **provided the S^1 holonomy preserves the center symmetry** (Kovner & Rosenstein, 1989; Ünsal, 2008). This means:

- The 4D limit is **exact** at $N = \infty$ (not merely Borel-summable).
- At finite N , the corrections are $\sim e^{-N}$ (the string tension contribution), which is a different suppression than the KK suppression $e^{-1/(R\Lambda)^2}$.

The two arguments are complementary:

- **Resurgence** (§5.4.3): the $5D \rightarrow 4D$ limit is Borel-summable at any N (the KK singularities retreat to infinity).
- **Volume independence** (Eguchi–Kawai): the $5D \rightarrow 4D$ limit is exact at $N = \infty$ (the corrections vanish identically).

Both support the smoothness of the $5D \rightarrow 4D$ limit, addressing the 5D loophole from two directions.

5.5 Lefschetz thimble decomposition: the lattice-computable realization

5.5.1 The Picard–Lefschetz decomposition

The **Lefschetz thimble method** (Cristoforetti et al., 2013; Witten, 2010) complexifies the path integral variables and decomposes the integration contour into steepest-descent cycles (thimbles), each associated with a saddle point. On each thimble, the imaginary part of the action is constant, so the sign problem disappears. The full integral is:

$$Z = \sum_{\sigma} n_{\sigma} \int_{\mathcal{J}_{\sigma}} \mathcal{D}A e^{-S[A]}$$

where $\mathcal{J}_{\sigma}$ are the thimbles (steepest-descent cycles), n_{σ} are the **intersection numbers** (Stokes phenomenon coefficients), and the sum is over all relevant saddles.

5.5.2 Topological sectors as thimbles

The central difficulty in computing χ on the lattice is that the topological charge Q is not a smooth function of the gauge field — it is integer-valued and jumps across gauge orbits. In the Lefschetz framework, each topological sector $Q = k$ corresponds to a set of thimbles:

$$Z_Q = \sum_{\sigma: Q_{\sigma} = Q} n_{\sigma} \int_{\mathcal{J}_{\sigma}} \mathcal{D}A e^{-S[A]}$$

where Q_{σ} is the topological charge of the saddle associated with thimble $\mathcal{J}_{\sigma}$. The topological susceptibility becomes:

$$\chi = \frac{1}{V} \left\langle \left(\sum_{\sigma} n_{\sigma} Q_{\sigma} \right)^2 \right\rangle$$

This makes the topological sectors **explicitly computable**: each thimble is a smooth integral (no sign problem), and the Stokes coefficients n_{σ} encode the topology. The integer-valued topological charge, which is discontinuous on the real path integral, becomes a smooth label on each thimble.

5.5.3 Connection to [A]'s Result A and the mass gap

[A]'s Result A uses the spectral representation $\chi = \int_0^{\infty} d\mu^2 \rho(\mu^2) / \mu^2$ and the mass gap to bound the integral. The Lefschetz decomposition makes $\rho(\mu^2)$ **computable on the lattice**: each thimble contributes a piece of the spectral density, and the thimble structure (which saddles contribute) is determined by the mass gap (which suppresses large instantons, i.e., saddles with large action are exponentially suppressed by e^{-S}).

The two approaches are complementary:

- **[A]**: gives the *bound* $\chi \leq C \Delta^4$ from the spectral representation and the mass gap.
- **Lefschetz method**: gives the *computation* of χ from the thimble decomposition, where

each thimble's contribution is controlled by the saddle action (which is gap-controlled).

5.5.4 The CEH saddle on the lattice

The CEH instanton is a saddle of the combined gauge + gravitational path integral. In the Lefschetz framework, it corresponds to a thimble $\mathcal{J}_{\text{CEH}}$ with:

- Topological charges $Q_{\text{gauge}} = \pm 1$, $\sigma = \pm 1$, $Q_{\text{NY}} = \pm 1$ (the three-sector charges).
- Action $S_{\text{CEH}} = 8\pi^2/g^2 + \pi^2/\kappa^2 + S_{\text{NY}}$.
- Fluctuation determinant R_{CEH} (the one-loop measure, [A]'s §8.7 open calculation).

The **lattice computation** of R_{CEH} proceeds as follows:

1. **Discretize the CEH geometry:** the Eguchi–Hanson gravitational instanton is a smooth Riemannian 4-manifold. In Regge calculus ([B]'s tool, §2.4, and the bridge document's Layer III.2), it can be approximated by a triangulated manifold, and the Pontryagin number $\sigma = \pm 1$ is computed as a sum over deficit angles. The colored EH instanton (with gauge bundle) is a lattice geometry with both gauge and gravitational connections defined on the edges.
2. **Compute the fluctuation determinant:** on the triangulated CEH geometry, the one-loop fluctuation determinant $\det(\mathcal{O}_{\text{CEH}})$ is computed by diagonalizing the fluctuation operator on the lattice. This is a **finite-dimensional** computation (the lattice has finitely many sites), and the determinant is a product of eigenvalues.
3. **Extrapolate to the continuum:** the lattice determinant is extrapolated to the continuum ($a \rightarrow 0$) using finite-size scaling. The question is whether the extrapolation converges to a gauge-scale value ($R_{\text{CEH}} \sim 1$) or a Planck-suppressed value ($R_{\text{CEH}} \sim 1/M_{\text{Pl}}^n$).

This is a **concrete lattice computation**, not merely a formal question. The ingredients are all available:

- The CEH geometry (Eguchi–Hanson + gauge bundle) is known.
- The Regge discretization of gravitational instantons is established (Hamada, 2010s).
- The fluctuation determinant computation is standard lattice technology (diagonalize a matrix).
- The continuum extrapolation is standard (finite-size scaling, [B]'s §2.4).

The challenge is the **combined gauge + gravitational lattice**: most lattice computations treat either the gauge sector (lattice QCD) or the gravitational sector (Regge calculus / CDT) separately. A combined lattice with both gauge and gravitational connections on the same triangulation is technically challenging but not conceptually impossible. The bridge document's Layer III.2 (CDT) and Layer II.3 (persistent homology / TDA) provide the tools.

5.5.5 The Stokes phenomenon and the three-sector bundling

The **Stokes phenomenon** — the jumping of the intersection numbers n_{σ} as the coupling constants are varied — encodes the topological structure. In the three-sector theory, the Stokes coefficients jump when the relative strengths of the gauge, gravitational, and Nieh–Yan couplings change. The CEH thimble has nontrivial Stokes coefficients that couple all three sectors:

$$n_{\text{CEH}} = n_{\text{CEH}}(g^2, \kappa^2, \alpha_{\text{NY}})$$

where α_{NY} is the torsion coupling. The **three-sector bundling** (§3.7) is a Stokes phenomenon: the CEH thimble cannot be decomposed into separate gauge, gravitational, and Nieh–Yan thimbles. This is the Lefschetz-theoretic avatar of the CEH bundling — the inseparability of the three θ -parameters is encoded in the thimble structure.

5.6 The Gutzwiller–Selberg trace formula: instantons as periodic orbits

5.6.1 The trace formula

The **Gutzwiller trace formula** (Gutzwiller, 1971) relates the quantum energy spectrum of a chaotic system to the classical periodic orbits:

$$\rho(E) = \bar{\rho}(E) + \sum_p \sum_{k=1}^{\infty} A_{p,k} \cos\left(k S_p(E) - \frac{\pi}{2} \mu_{p,k}\right)$$

where $\rho(E)$ is the spectral density, $\bar{\rho}$ is the smooth (Weyl) part, the sum is over periodic orbits p and their repetitions k , S_p is the classical action, and $\mu_{p,k}$ is the Maslov index. The **Selberg trace formula** is the exact version for constant-curvature manifolds (Selberg, 1956).

5.6.2 Instantons as periodic orbits

In Euclidean Yang–Mills theory, instantons are saddle points of the Euclidean action — they are the "periodic orbits" of the gauge theory in imaginary time. The mass gap is the lowest spectral line. The trace formula connects these:

- **Instantons $\leftrightarrow$ periodic orbits:** each instanton of charge Q and size ρ contributes a term to the spectral density $\rho(\mu^2)$ in Result A. The instanton action $S = 8\pi^2/g^2$ plays the role of the classical action S_p .
- **Spectral gap $\leftrightarrow$ orbit structure:** in chaotic systems, the spectral gap (level repulsion) is controlled by the shortest periodic orbit. In Yang–Mills, the shortest "orbit" is the minimal-action instanton with $\rho \sim 1/\Delta$, giving $\lambda_{\min} \sim \Delta$ — exactly $[A]$'s averaged spectral bound (§3.3c).

5.6.3 The mass gap as the shortest orbit

The Gutzwiller formula predicts that the spectral density near the ground state is:

$$\rho(\mu^2) \sim \sum_{\text{instantons } p} A_p \delta(\mu^2 - \Delta_p^2)$$

where Δ_p is the mass associated with instanton p . The **mass gap** Δ is the lowest Δ_p — the mass of the lightest instanton excitation. In the instanton gas picture, the lightest excitation is the glueball, which is the fluctuation around the vacuum (the trivial "orbit"), not around an instanton. The trace formula reframes this: the glueball mass is the spectral gap, and the instantons are the non-trivial orbits whose repetitions generate the subleading oscillatory terms.

The connection to §4.2 (Floer homology) is direct: the Floer differential counts gradient-flow lines (instantons), and the Gutzwiller formula sums over periodic orbits (instanton repetitions). The Floer

spectral gap (the mass gap) is the Gutzwiller spectral gap. The two are the same quantity, expressed in different mathematical languages: Floer in the topological language (chain complex), Gutzwiller in the semiclassical language (periodic orbit sum).

5.6.4 The CEH instanton in the trace formula

The CEH instanton is a **combined gauge + gravitational periodic orbit** in the trace formula. Its contribution to the spectral density is:

$$\rho_{\text{CEH}}(\mu^2) \sim A_{\text{CEH}} \delta(\mu^2 - \Delta_{\text{CEH}}^2)$$

where Δ_{CEH} is the mass of the CEH excitation and $A_{\text{CEH}} = R_{\text{CEH}}$ is the one-loop amplitude. The gauge-scale vs. Planck-suppressed question is: is $A_{\text{CEH}} \sim 1$ (gauge-scale) or $A_{\text{CEH}} \sim 1/M_{\text{Pl}}^n$ (Planck-suppressed)?

The trace formula provides a **spectral** signature: if $A_{\text{CEH}} \sim 1$, the CEH contributes a visible spectral line at $\mu^2 = \Delta_{\text{CEH}}^2$. If $A_{\text{CEH}} \sim 1/M_{\text{Pl}}^n$, the CEH spectral line is invisible (Planck-suppressed). The search for the CEH spectral line is, in principle, a **spectroscopic** question — though the CEH excitation (a combined gauge + gravitational excitation) is not a standard particle and may not couple to conventional detectors.

5.7 The research program

5.7.1 Three-pronged approach

We propose a three-pronged research program to resolve the CEH measure problem:

Prong 1: Borel-plane analysis (analytic). Compute the Borel-plane singularity structure of the combined gauge + gravitational + Nieh–Yan theory. The steps are:

1. Compute the perturbative series $f_{\text{pert}}(g^2) = \sum a_k g^{2k}$ to high order (this requires perturbative YM + gravity calculations, feasible at least for the vacuum energy).
2. Identify the Borel singularities from the large-order behavior of a_k (the large-order growth encodes the nearest singularity).
3. Determine the residue R_{CEH} from the large-order/low-order duality.
4. Check for Planck-scale screening singularity.

This is an **analytic** calculation, feasible in principle but requiring substantial perturbative computation.

Prong 2: Lattice computation (numerical). Compute the CEH fluctuation determinant on a combined gauge + gravitational lattice. The steps are:

1. Construct the triangulated CEH geometry (Eguchi–Hanson + gauge bundle).
2. Compute the fluctuation operator on the lattice.
3. Diagonalize and compute the determinant.
4. Extrapolate to the continuum.

This is a **numerical** calculation, requiring a combined gauge + gravitational lattice code (not currently available, but the ingredients — lattice QCD and Regge calculus — are mature).

Prong 3: Lefschetz thimble computation (semianalytic). Decompose the combined path integral into thimbles, identify the CEH thimble, and compute its contribution. The steps are:

1. Complexify the combined gauge + gravitational path integral.
2. Identify the saddle points (trivial vacuum, gauge instantons, gravitational instantons, CEH instanton).
3. Construct the thimbles and compute the Stokes coefficients.
4. Compute the CEH thimble integral (the fluctuation determinant R_{CEH}).

This is a **semianalytic** calculation, combining the thimble decomposition (rigorous) with the lattice computation of each thimble integral (numerical).

5.7.2 Timeline and tractability

Prong	Method	Status	Tractability	Key obstacle
1. Borel-plane	High-order perturbation theory	Unstarted for 4D YM + gravity	Medium (2D analogs work)	4D perturbative series unknown to high order
2. Lattice	Combined gauge + gravity lattice	Unstarted	Medium (ingredients mature)	Combined lattice code does not exist
3. Lefschetz	Thimble decomposition	Early stage (2D models only)	Low–medium	4D thimble structure unknown; computational cost

None of the three prongs is currently feasible without substantial development, but all three are **well-posed** research programs with clear milestones. The 2D analogs (Prong 1 in $\mathbb{CP}^{N-1}$, Prong 3 in 2D YM) provide proofs of concept. The key novelty is the **combined gauge + gravitational saddle** (the CEH instanton), which has no analog in 2D models.

5.7.3 The role of the mass gap

The mass gap Δ enters the CEH measure problem in two ways:

1. **As a scale:** the CEH fluctuation determinant R_{CEH} depends on the mass gap through the spectral density $\rho(\mu^2)$ (which is supported on $\mu^2 \geq \Delta^2$). The mass gap sets the IR scale of the fluctuation determinant.
2. **As a regulator:** the mass gap makes the Borel-plane singularities well-separated. Without the gap ($\Delta = 0$), the instanton size ρ is unbounded, and the Borel singularities merge into a cut (continuous spectrum). With the gap ($\Delta > 0$), the instanton size is bounded $\rho \lesssim 1/\Delta$, and the Borel singularities are discrete. The mass gap makes the trans-series well-defined.

This is a second manifestation of the **decoupling** (§3.4.1): the mass gap (dynamical) and the $\theta = 0$ result (topological) are independent, but the mass gap is **necessary for the well-posedness of the CEH measure problem**. Without the gap, the Borel-plane analysis breaks down (the singularities are

not discrete), and the CEH measure problem is ill-defined. With the gap, the problem is well-posed and, in principle, computable.

5.8 Summary of Section 5

Key results

Result	Status	Connection
CEH measure problem = Borel-plane singularity analysis	Framework established (§5.3)	[A] §8.7 reframed as computable
Gauge-scale vs. Planck-suppressed = residue of CEH Borel singularity	Question sharpened (§5.3.3)	[A] §8.9 decision table
Three-sector bundling = single CEH Borel singularity	Derived (§5.3.4)	§3.7 CEH three-sector bundling
5D loophole = Borel-summability of KK corrections	Addressed (§5.4)	[B] §2.2.3
Lefschetz thimbles = lattice-computable topological sectors	Framework (§5.5)	[B] lattice tools
CEH lattice computation = Regge + gauge lattice	Proposed (§5.5.4)	Bridge Layer III.2
Instantons as periodic orbits = Gutzwiller trace formula	Structural (§5.6)	Bridge Layer I.2
Three-pronged research program	Proposed (§5.7)	Analytic + numerical + semianalytic

The main thesis of Section 5

[A]'s central open calculation — the CEH one-loop fluctuation determinant — is **not a dead end** but a **well-posed research program**. Resurgence theory reframes it as a Borel-plane singularity analysis, the Lefschetz thimble method provides a lattice-computable realization, and the Gutzwiller trace formula provides a spectral signature. The three-pronged research program (analytic Borel-plane, lattice Regge+gauge, Lefschetz thimbles) offers three independent paths to the same answer.

The mass gap $\Delta > 0$ plays a dual role: it is the dynamical quantity proved by [B] (decoupled from the $\theta = 0$ result), and it is the **regulator that makes the CEH measure problem well-posed** (by making the Borel-plane singularities discrete). Without the mass gap, the CEH measure problem is ill-defined; with the mass gap, it is computable in principle.

6. The Gravitational Spectral Gap Beyond AdS

6.1 [A]'s open question: does $\Delta_{\text{grav}} > 0$ exist outside AdS?

[A] derives the gravitational spectral gap $\Delta_{\text{grav}} \sim \Delta / L_{\text{AdS}}$ (§7) using the AdS/CFT correspondence: the mass gap Δ of the boundary CFT (Yang–Mills) translates to the bulk gravitational spectral gap via the AdS radius L_{AdS} . This is a rigorous derivation **within the AdS/CFT framework**, but it leaves open the most serious question of [A] (§10.1, §10.2):

Does the gravitational spectral gap $\Delta_{\text{grav}} > 0$ exist **outside** AdS/CFT — i.e., in de Sitter space (our universe) or in flat space?

In de Sitter space, no mechanism analogous to AdS/CFT is known to derive Δ_{grav} from first principles. The AdS derivation relies on:

1. The **boundary CFT** being well-defined (unitary, with a positive spectral density).
2. The **holographic dictionary** mapping boundary spectral data to bulk observables.
3. The **AdS geometry** providing a confining box (the IR wall) that discretizes the spectrum.

None of these ingredients survive straightforwardly in de Sitter space. The dS/CFT correspondence (§6.3) provides a partial analog, but the boundary theory is non-unitary, and the holographic dictionary is less established. Outside any holographic framework, the question is wide open.

This section addresses [A]'s open question from three directions:

1. **Causal Dynamical Triangulations (CDT)** (§6.2): a non-perturbative lattice quantum gravity that produces an emergent de Sitter geometry, where Δ_{grav} is **computable** from the lattice data.
2. **dS/CFT correspondence** (§6.3): a holographic analog of AdS/CFT for de Sitter space, where $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$ (with H the Hubble parameter), modulo the non-unitarity subtlety.
3. **Ashtekar gravity as gauge theory** (§6.4): [A]'s identification of gravity as a gauge theory (Ashtekar variables) extends the mass gap from the gauge sector to the gravitational sector structurally.

The three approaches are complementary: CDT provides a computational answer in de Sitter, dS/CFT provides a holographic answer in de Sitter, and Ashtekar gravity provides a structural answer independent of the spacetime signature.

6.2 Causal Dynamical Triangulations: the gravitational gap in de Sitter

6.2.1 The CDT framework

Causal Dynamical Triangulations (CDT; Ambjørn, Jurkiewicz, Loll, since 1998) is a lattice

regularization of the gravitational path integral using piecewise flat simplicial manifolds with a **causal structure** — a preferred foliation by spatial slices. The path integral sums over all triangulations weighted by the Regge action:

$$Z_{\text{CDT}} = \sum_{T \in \mathcal{T}} \frac{1}{|C_T|} e^{-S_{\text{Regge}}[T]}$$

where $\mathcal{T}$ is the set of causally triangulated manifolds, $|C_T|$ is the symmetry factor (automorphism group order), and $S_{\text{Regge}}[T]$ is the Regge action (the lattice discretization of the Einstein–Hilbert action with cosmological constant).

CDT is the **only lattice gravity approach** that has been shown to produce an emergent 4D de Sitter geometry in the large-volume limit. The key results are:

- **Emergent de Sitter geometry:** the expectation value of the spatial volume $\langle V(\tau) \rangle$ as a function of proper time τ follows the de Sitter profile $V(\tau) \propto \cosh^3(H\tau)$, where H is the emergent Hubble parameter.
- **Spectral dimension flow:** the spectral dimension d_S (defined from the return probability of a diffusion process on the triangulation) runs from $d_S \approx 2$ at Planck scales to $d_S \approx 4$ at large scales. This **dimensional reduction** is a signature of quantum gravity effects.
- **Phase structure:** the CDT phase diagram has a **de Sitter phase** and a **branched-polymer phase** (and a crumpled phase at strong coupling), separated by phase transitions. The de Sitter phase is the physically relevant one.

6.2.2 The spectral dimension and the gravitational gap

The **spectral dimension** d_S is defined via the diffusion process on the triangulation:

$$\langle P(\sigma) \rangle \sim \sigma^{-d_S/2} \quad \text{as } \sigma \rightarrow \infty$$

where $P(\sigma)$ is the return probability of a diffusion process after σ steps, and σ is the diffusion time (proportional to the squared diffusion length).

The spectral dimension is related to the **graviton propagator** $G(r)$ by:

$$G(r) \sim r^{-(d_S - 2)} \quad \text{at distance } r$$

In 4D classical gravity, $d_S = 4$ and $G(r) \sim 1/r^2$ (massless graviton). In CDT, the spectral dimension flows:

- **At large scales** ($r \gg \ell_{\text{Planck}}$): $d_S \rightarrow 4$, and $G(r) \sim 1/r^2$ — the graviton is effectively massless.
- **At Planck scales** ($r \sim \ell_{\text{Planck}}$): $d_S \rightarrow 2$, and $G(r) \sim 1/r^0 = \text{const}$ — the graviton propagator is **modified**.

The dimensional reduction from 4 to 2 at short distances is a signature of a **spectral gap** in the graviton propagator. More precisely, the running spectral dimension implies that the graviton propagator acquires a momentum-dependent effective mass:

$$G(p) \sim \frac{1}{p^2 + m_{\text{eff}}^2(p)}$$

where $m_{\text{eff}}^2(p) \rightarrow 0$ at low momentum (large distances) and $m_{\text{eff}}^2(p) \rightarrow \Lambda_{\text{Planck}}^2$ at high momentum (short distances). This is a **dynamical mass generation** for the graviton, not a hard mass gap in the traditional sense.

6.2.3 The de Sitter phase and volume fluctuations

In the CDT de Sitter phase, the spatial volume fluctuations around the de Sitter profile are **suppressed**:

$$\langle \delta V(\tau) \rangle = V(\tau) - \langle V(\tau) \rangle, \quad \langle (\delta V)^2 \rangle \ll \langle V \rangle^2$$

The suppression of volume fluctuations is a **lattice signature of a spectral gap**. In a gapless theory (massless modes), the fluctuations would be of the same order as the mean (as in a conformal theory, where $\langle (\delta V)^2 \rangle \sim \langle V \rangle^2$). The suppression in the de Sitter phase indicates that the fluctuations are controlled by a massive mode — the **gravitational gap**.

More precisely, the volume fluctuation correlator decays exponentially:

$$\langle \delta V(\tau), \delta V(\tau') \rangle \sim e^{-\Delta_{\text{grav}}^{\text{CDT}} |\tau - \tau'|}$$

where $\Delta_{\text{grav}}^{\text{CDT}}$ is the gravitational spectral gap in CDT. The exponential decay is the lattice signature of a gap — exactly the same signature as the entanglement gap diagnostic of §4.4 ($e^{-\Delta R}$). In CDT, the "distance" is the proper time τ , and the gap is the gravitational gap $\Delta_{\text{grav}}^{\text{CDT}}$.

6.2.4 Computing $\Delta_{\text{grav}}^{\text{CDT}}$ on the CDT lattice

The gravitational spectral gap in CDT is **computable** from the lattice data:

1. **Generate CDT configurations** in the de Sitter phase (Ambjørn et al. have done this since the 2000s).
2. **Measure the volume fluctuations** $\delta V(\tau)$ and their correlator $\langle \delta V(\tau), \delta V(\tau') \rangle$.
3. **Fit the exponential decay rate** to extract $\Delta_{\text{grav}}^{\text{CDT}}$.
4. **Extrapolate** to the continuum limit (large triangulation size, small lattice spacing).

This is a **direct lattice computation** of $\Delta_{\text{grav}}^{\text{CDT}}$, independent of AdS/CFT . It works in de Sitter space (the emergent geometry of CDT), not in AdS. The computation is analogous to the glueball mass extraction in lattice QCD: measure a correlator, fit the exponential decay, extract the gap.

6.2.5 The gravitational topological susceptibility on CDT

[A]'s gravitational topological susceptibility χ_{grav} (§6) can in principle be computed on CDT backgrounds by measuring the fluctuations of the Pontryagin density $\text{Tr}(R \wedge R)$ on the triangulation:

$$\chi_{\text{grav}} = \frac{1}{V} \langle Q_{\text{grav}}^2 \rangle, \quad Q_{\text{grav}} = \frac{1}{8\pi^2} \int \text{Tr}(R \wedge R)$$

On the CDT lattice, the Pontryagin density is computed from the deficit angles at the hinges of the triangulation (the Regge calculus discretization). The gravitational topological charge Q_{grav} is the sum of deficit angles (weighted by the appropriate coefficients), and χ_{grav} is its variance per unit volume.

This computation is **analogous** to the gauge topological susceptibility computation in lattice QCD ([B]'s tool), but for the gravitational sector. It has not been performed — CDT studies have focused on the geometric properties (volume profile, spectral dimension) and the phase structure, not on the

topological fluctuations. But the ingredients are available: the Regge calculus definition of the Pontryagin number is standard, and the CDT configurations are generated.

If both $\Delta_{\text{grav}}^{\text{CDT}}$ and $\chi_{\text{grav}}^{\text{CDT}}$ can be computed on the same CDT configurations, one can test [A]'s gravitational analogue of Result A:

$$\chi_{\text{grav}} \leq C_{\text{grav}}, (\Delta_{\text{grav}})^4$$

This would be the **gravitational sector test** of [A]'s bound, computed entirely within CDT, independent of AdS/CFT.

6.3 dS/CFT correspondence: the holographic gap in de Sitter

6.3.1 The dS/CFT proposal

The dS/CFT correspondence (Strominger, 2001; Maldacena, 2003) proposes that de Sitter space in $d+1$ dimensions is dual to a Euclidean CFT in d dimensions living at **future infinity** $\mathcal{I}^+$ (or past infinity $\mathcal{I}^-$). The correspondence is motivated by the same holographic principle as AdS/CFT, but with crucial differences:

Feature	AdS/CFT	dS/CFT
Boundary	Timelike (∂AdS)	Spacelike ($\mathcal{I}^+$, future null infinity)
Boundary theory	Unitary CFT	Non-unitary Euclidean CFT
S-matrix	Well-defined (boundary correlators)	No conventional S-matrix (de Sitter has no asymptotic states)
Spectral representation	Positive spectral density (Källén–Lehmann)	Not guaranteed (non-unitarity)
Tests	Extensive (thousands of papers)	Limited (3D examples, CMB correlators)

The non-unitarity of the boundary CFT is the central obstacle. It arises because the de Sitter isometry group $SO(1, d+1)$ acts on the boundary as the **Euclidean conformal group** $SO(1, d+1)$, which has non-unitary representations (the boundary theory is Euclidean, not Lorentzian). This means the Källén–Lehmann spectral representation (which requires a positive spectral density) may fail.

6.3.2 The dS gravitational gap

[A]'s AdS derivation gives $\Delta_{\text{grav}}^{\text{AdS}} \sim \Delta / L_{\text{AdS}}$. In dS/CFT, the analogous statement would be:

$$\Delta_{\text{grav}}^{\text{dS}} \sim \frac{\Delta}{L_{\text{dS}}} = \Delta \cdot H$$

where $H = 1/L_{\text{dS}}$ is the Hubble parameter. The boundary mass gap Δ (Yang–Mills) translates to the bulk gravitational gap via the de Sitter radius, exactly as in AdS but with $L_{\text{AdS}} \rightarrow L_{\text{dS}} = 1/H$.

The derivation parallels the AdS case:

- The boundary CFT (Yang–Mills) has a mass gap $\Delta > 0$.
- The bulk graviton is the dual of the boundary stress tensor.
- The boundary stress tensor has a spectral gap Δ (set by the mass gap).
- The bulk graviton spectral gap is $\Delta_{\text{grav}} = \Delta / L_{\text{dS}} = \Delta \cdot H$.

6.3.3 The non-unitarity subtlety

The derivation above is **formally** correct but faces a subtlety: the Källén–Lehmann spectral representation ([A]’s §2, used to prove Result A: $\chi \leq C\Delta^4$) requires a **positive spectral density** $\rho(\mu^2) \geq 0$. In a unitary theory, this is guaranteed by the positivity of the Hilbert space norm. In a non-unitary theory (the dS/CFT boundary CFT), the spectral density may not be positive, and the Källén–Lehmann argument may fail.

This means:

- **In AdS/CFT:** the boundary CFT is unitary, $\rho(\mu^2) \geq 0$, and [A]’s Result A holds rigorously. The gravitational bound $\chi_{\text{grav}} \leq C_{\text{grav}} (\Delta_{\text{grav}})^4$ follows.
- **In dS/CFT:** the boundary CFT is non-unitary, $\rho(\mu^2)$ may not be positive, and [A]’s Result A may not hold. The gravitational bound $\chi_{\text{grav}} \leq C_{\text{grav}} (\Delta_{\text{grav}})^4$ is **conjectural**.

The non-unitarity is a genuine obstacle, but it may not be fatal:

1. **The mass gap itself** (the existence of $\Delta > 0$) does not require positivity of ρ . The gap is a statement about the **support** of ρ (no eigenvalues below Δ), not about its sign. Even if ρ is not positive, a gap in the support implies a gap in the spectrum.
2. **The bound** $\chi \leq C\Delta^4$ uses the positivity of ρ to bound the integral $\int d\mu^2 \rho(\mu^2)/\mu^2$. Without positivity, the integral may not be bounded. However, if ρ has a definite sign (even if negative), the bound may hold with a modified constant.
3. **The physical observables** (the gravitational gap Δ_{grav} , the topological susceptibility χ_{grav}) are bulk quantities, not boundary quantities. The non-unitarity of the boundary theory may not directly affect the bulk observables. The bulk de Sitter geometry has a positive-definite Hilbert space (the de Sitter group $SO(1, d+1)$ has unitary representations), and the bulk spectral representation may be well-defined even if the boundary one is not.

The resolution of this subtlety is an open problem. The dS/CFT derivation of $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$ is a **conjecture**, not a theorem, until the non-unitarity issue is resolved. The CDT computation (§6.2) provides an independent check: if CDT finds $\Delta_{\text{grav}}^{\text{CDT}} > 0$ in the de Sitter phase, it confirms the dS/CFT prediction without relying on the non-unitary boundary theory.

6.3.4 Cosmological observables and the amplification question

If $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$, then the gravitational θ -term in our

universe has a physical effect of order:

$$\chi_{\text{grav}} \sim (\Delta_{\text{grav}})^4 \sim (\Delta \cdot H)^4$$

With $\Delta \sim \Lambda_{\text{QCD}} \sim 200, \text{MeV}$ and $H \sim 10^{-33}, \text{eV}$ (the current Hubble parameter):

$$\chi_{\text{grav}} \sim (200 \times 10^6 \times 10^{-33})^4 \sim (2 \times 10^{-25})^4 \sim 10^{-100}, \text{eV}^4$$

This is **incredibly small** — far below any observable. The gravitational θ -sector is observationally irrelevant in de Sitter at face value.

However, the **CEH bundling** ([A], §8; §3.7; §5.3) could **amplify** this. The CEH instanton bundles all three θ -parameters, and its amplitude R_{CEH} may be gauge-scale (not Planck-suppressed). If $R_{\text{CEH}} \sim \Lambda_{\text{QCD}}^4$, then the CEH contribution to the vacuum energy is $\sim \Lambda_{\text{QCD}}^4 \cdot e^{i(\theta_{\text{QCD}} + \theta_{\text{grav}} + \theta_{\text{NY}})}$, which is gauge-scale regardless of H . The de Sitter suppression $\sim (\Delta \cdot H)^4$ applies to the *direct* gravitational θ -effect; the CEH-bundled effect is a *combined* gauge + gravitational effect that is not suppressed by H .

The dS/CFT framework allows one to compute this amplification factor: the CEH instanton, viewed as a bulk saddle in de Sitter, contributes to the boundary CFT partition function as a non-perturbative saddle. Its amplitude in the boundary theory is R_{CEH} , which is the same quantity computed in §5.3 (the Borel-plane residue). The dS/CFT correspondence maps this to a bulk gravitational effect, which is the amplification of the gravitational θ -sector through CEH bundling.

This is the **cosmological significance** of the CEH measure problem (§5): the answer to "is R_{CEH} gauge-scale or Planck-suppressed?" determines whether the gravitational θ -sector is observationally relevant in our universe. If gauge-scale, the CEH amplifies the gravitational θ -effect from $\sim (\Delta H)^4$ to $\sim \Delta^4$, making it potentially observable (though still constrained by the gauged P/CP, which sets $\theta_{\text{grav}} = 0$). If Planck-suppressed, the gravitational θ -sector is permanently unobservable.

6.4 Ashtekar gravity: the structural identification of gravity as gauge theory

6.4.1 The Ashtekar variables

[A] (§4) employs the **Ashtekar formulation** of general relativity, which rewrites gravity as a gauge theory. The Ashtekar variables (A^i_a, E^a_i) are a complex $SU(2)$ gauge connection (the **Ashtekar connection**) and its conjugate electric field (the **densitized triad**), where $i = 1, 2, 3$ is the $SU(2)$ gauge index and $a = 1, 2, 3$ is the spatial index.

The Ashtekar connection is:

$$A^i_a = \Gamma^i_a + i, \gamma, K^i_a$$

where Γ^i_a is the spin connection (the Levi-Civita connection on the triad) and K^i_a is the extrinsic curvature (related to the time derivative of the triad). The **Immirzi parameter** γ is a real constant (in the original Ashtekar formulation, $\gamma = i$; in the real formulation, γ is real).

In these variables, the Einstein–Hilbert action becomes:

$$S_{\text{EH}} = \int dt, d^3x, E^a_i \dot{A}^i_a - \Lambda S_{\text{constraint}}[A, E]$$

where $S_{\text{constraint}}$ is a sum of constraints (Gauss, diffeomorphism, Hamiltonian) that generate the gauge symmetries of gravity. The theory is a **gauge theory** with gauge group $SU(2)$ (or $SL(2, \mathbb{C})$ in the complex formulation).

6.4.2 The structural mass gap

The Ashtekar identification has a profound consequence for the mass gap: **if gravity is a gauge theory, and the gauge theory has a mass gap, then gravity has a gravitational gap.**

The argument is structural:

1. The Ashtekar connection A^i_a is a gauge connection, exactly like the Yang–Mills connection A_μ^a .
2. The Yang–Mills mass gap $\Delta > 0$ ([B]'s result) is a statement about the spectrum of the gauge-theory Hamiltonian.
3. The gravitational Hamiltonian, in Ashtekar variables, is the same type of operator (a gauge-theory Hamiltonian with $SU(2)$ gauge group).
4. If the Yang–Mills gauge theory with gauge group $SU(N)$ has a mass gap, and the Ashtekar gravity is a gauge theory with gauge group $SU(2)$, then the Ashtekar gravity **may** have a gravitational gap $\Delta_{\text{grav}} > 0$.

This is a **structural** argument, not a proof. The caveats are:

- **Different gauge groups:** Yang–Mills is $SU(3)$ (QCD); Ashtekar gravity is $SU(2)$. The mass gap is known (conjecturally) for $SU(N)$ with $N \geq 2$, but the gravitational case has additional constraints (diffeomorphism invariance, the Hamiltonian constraint).
- **Diffeomorphism invariance:** the gravitational Hamiltonian is a constraint ($H = 0$ on physical states), not a standard Hamiltonian with a spectrum. The "gravitational gap" is not the gap above the vacuum of H (which is zero by constraint), but the gap in the **physical excitations** (gravitons) above the vacuum.
- **Background independence:** the Ashtekar variables are background-independent (no fixed spacetime metric), while the Yang–Mills mass gap is defined on a fixed background ($\mathbb{R}^{3,1}$). The background-independent mass gap is a more subtle quantity.

Despite these caveats, the structural identification is powerful: it provides a **reason** to expect $\Delta_{\text{grav}} > 0$ independent of AdS/CFT or dS/CFT. The gravitational gap is the Ashtekar-gauge-theory analog of the Yang–Mills mass gap.

6.4.3 The three-sector unification

In the 5D theory with torsion (§3), the Ashtekar identification extends the three-sector framework to gravity:

Sector	Chern–Weil form	θ -parameter	Ashtekar identification
Gauge	$\text{Tr}(F \wedge F)$	θ_{QC} Yang–Mills $SU(3)$	

	$D\}}\}$	connection
Gravitational	$\text{Tr}(R \wedge R)$	Ashtekar $SU(2)$ connection
Nieh–Yan	$\mathcal{N} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$	Torsion in the Ashtekar connection

The Ashtekar connection $A^i_a = \Gamma^i_a + i\gamma K^i_a$ includes both the spin connection Γ (which enters $\text{Tr}(R \wedge R)$) and the extrinsic curvature K (which, combined with torsion, enters $\mathcal{N}$). The three sectors are three aspects of the **same Ashtekar gauge theory**:

- The **gauge sector** is the Yang–Mills $SU(3)$ connection (independent of gravity).
- The **gravitational sector** is the Ashtekar $SU(2)$ connection (the spin connection part).
- The **Nieh–Yan sector** is the torsion in the Ashtekar connection (the K part, with torsion).

The mass gap in the gauge sector ([B]'s result) and the structural gravitational gap (the Ashtekar identification) are two faces of the same gauge-theory mass gap, now in a unified framework with three sectors.

6.4.4 Loop quantum gravity and the area spectrum

The Ashtekar formulation is the starting point of **Loop Quantum Gravity** (LQG; Rovelli, Smolin, Ashtekar, since the 1990s). In LQG, the spatial geometry is quantized: the area and volume operators have **discrete spectra**:

$$\hat{A}[\mathcal{S}], \psi_n = 8\pi\gamma\ell_{\text{Pl}}^2\sqrt{j(j+1)}, \psi_n, \text{quad } j = 0, \frac{1}{2}, 1, \frac{3}{2}, \dots$$

where $\mathcal{S}$ is a surface, ℓ_{Pl} is the Planck length, and j is the $SU(2)$ spin quantum number. The **area gap** — the smallest nonzero eigenvalue of the area operator — is:

$$\Delta A_{\min} = 8\pi\gamma\ell_{\text{Pl}}^2\sqrt{\frac{3}{4}} = 4\pi\gamma\ell_{\text{Pl}}^2\sqrt{3}$$

The area gap is the **geometric** analog of the mass gap. It is a discrete spectral gap in the geometry, arising from the quantization of the Ashtekar connection. The gravitational spectral gap Δ_{grav} (the gap in the graviton spectrum) and the area gap $\Delta A_{\min}$ (the gap in the geometry spectrum) are related but distinct quantities:

- Δ_{grav} is the gap in the **excitation spectrum** (the graviton mass).
- $\Delta A_{\min}$ is the gap in the **geometric spectrum** (the minimum area).

Both arise from the Ashtekar gauge-theory structure. The area gap is rigorously derived in LQG (it is a theorem). The gravitational spectral gap is conjectural (it depends on the dynamics of the graviton, which LQG has not fully computed). But the structural identification suggests that both gaps have the same origin: the gauge-theory mass gap applied to the Ashtekar connection.

6.5 The eta invariant and the boundary gravitational gap

6.5.1 The APS index theorem with boundary

The **Atiyah–Patodi–Singer (APS) index theorem** (Atiyah, Patodi & Singer, 1975) extends the Atiyah–Singer index theorem to manifolds with boundary, introducing the **eta invariant** $\eta(D)$ — the spectral asymmetry of the Dirac operator:

$$\text{index}(D) = \int_M \hat{A}(M) \wedge \text{ch}(E) - \frac{\eta(D) + h}{2}$$

where $h = \dim \ker D$ is the dimension of the kernel. The eta invariant measures how asymmetric the spectrum is:

$$\eta(D) = \sum_{\lambda \neq 0} \text{sign}(\lambda) \quad (\text{regularized})$$

where the sum is over the nonzero eigenvalues λ of D .

6.5.2 The eta invariant as the boundary gravitational gap

On a manifold with boundary (e.g., a finite-volume lattice, or a region with an entanglement cut as in §4.4), the index theorem acquires the $\eta/2$ boundary term. The eta invariant is the **boundary correction** to the index, and it is the **integral of the Chern–Simons 3-form** over the boundary:

$$\eta(D) \sim \int_{\partial M} CS_3[A]$$

This connects to the anomaly inflow mechanism (§3.4): the bulk CS_5 (the 5D Chern–Simons 5-form) flows to the boundary CS_3 (the 4D Chern–Simons 3-form), whose integral is the eta invariant. The entire chain — from the 5D bulk to the 4D θ -term to the boundary spectral asymmetry — is unified.

The eta invariant measures the **spectral asymmetry** of the Dirac operator on the boundary. In a gapped theory ($\Delta > 0$), the Dirac spectrum on the boundary has a gap, and the eta invariant is well-defined (the spectral asymmetry is finite). In a gapless theory ($\Delta = 0$), the Dirac spectrum on the boundary is continuous, and the eta invariant may diverge (the spectral asymmetry is ill-defined).

The **gravitational eta invariant** $\eta(D_{\text{grav}})$ — the eta invariant of the Dirac operator coupled to the gravitational background (the Ashtekar connection) — measures the spectral asymmetry of the gravitational Dirac operator. Its finiteness is controlled by the **gravitational gap** $\Delta_{\text{grav}} > 0$:

- If $\Delta_{\text{grav}} > 0$ (gravitational gap exists): the gravitational Dirac spectrum is discrete, $\eta(D_{\text{grav}})$ is finite, and the boundary index theorem is well-defined.
- If $\Delta_{\text{grav}} = 0$ (no gravitational gap): the gravitational Dirac spectrum is continuous, $\eta(D_{\text{grav}})$ may diverge, and the boundary index theorem breaks down.

This provides a **diagnostic** for the gravitational gap: compute the eta invariant on a manifold with boundary (e.g., a finite-volume CDT triangulation with boundary), and check whether it is finite. A finite eta invariant indicates a gravitational gap; a divergent one indicates the absence of a gap.

6.5.3 The three-sector eta invariant

In the three-sector theory (§3), the Dirac operator couples to all three backgrounds (gauge, gravitational, Nieh–Yan), and the eta invariant receives contributions from all three:

$$\eta(D) = \eta_{\text{gauge}}(D) + \eta_{\text{grav}}(D) + \eta_{\text{NY}}(D)$$

Each sector's eta invariant is controlled by the corresponding gap:

- η_{gauge} : controlled by the gauge mass gap $\Delta > 0$ ([B]'s result).
- η_{grav} : controlled by the gravitational gap $\Delta_{\text{grav}} > 0$ (the open question).
- η_{NY} : controlled by the Nieh–Yan gap $\Delta_{\text{NY}} > 0$ (related to the torsion scale).

The finiteness of the total eta invariant requires all three gaps to be positive. The gauge gap is established ([B]); the gravitational gap is the open question of this section; the Nieh–Yan gap is a new quantity introduced by the torsion (§3), related to the torsion scale.

6.5.4 Connection to the entanglement gap

The eta invariant on the entanglement cut (§4.4.5) is the modular Hamiltonian's spectral asymmetry. The three-sector eta invariant provides the **boundary correction** to the three-sector index theorem on the entanglement cut. The entanglement gap (§4.4) is the gap in the modular Hamiltonian spectrum, and the eta invariant measures its spectral asymmetry. The two are related:

- **Entanglement gap** (§4.4): the gap in the modular Hamiltonian K_R spectrum, controlled by Δ .
- **Eta invariant** (§6.5): the spectral asymmetry of K_R , controlled by the same Δ (and, in the three-sector case, by all three gaps).

The finiteness of the eta invariant is a **necessary condition** for the entanglement gap to be well-defined: a divergent eta invariant would mean the modular spectrum is continuous (no gap). The gravitational gap $\Delta_{\text{grav}} > 0$ is necessary for the gravitational part of the eta invariant to be finite, and hence for the three-sector entanglement gap to be well-defined.

6.6 Summary of Section 6

Three approaches to the gravitational gap beyond AdS

Approach	Where it works	Method	Status	What it provides
CDT (§6.2)	de Sitter	Lattice quantum gravity	Active program; gap computable from volume fluctuations	Direct computation of $\Delta_{\text{grav}}^{\text{CDT}}$ in de Sitter
dS/CFT (§6.3)	de Sitter	Holographic	Proposed; non-unitarity subtlety	Conjectural $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$; amplification via CEH
Ashtekar	Any	Structural	Established (Ashtekar	Structural reason for

gravity (§6.4)	signature		variables); LQG area gap rigorous	$\Delta_{\text{grav}} > 0$; three-sector unification
Eta invariant (§6.5)	Manifold with boundary	APS index theorem	Rigorous theorem	Diagnostic: finite $\eta \leftrightarrow$ gravitational gap

The key results

1. **CDT makes Δ_{grav} computable in de Sitter** (§6.2.4). The volume fluctuation correlator decays as $e^{-\Delta_{\text{grav}}|\tau|}$, and the gravitational topological susceptibility χ_{grav} can be measured on the same configurations, enabling a direct test of $\chi_{\text{grav}} \leq C_{\text{grav}}(\Delta_{\text{grav}})^4$.
2. **dS/CFT gives $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$, modulo non-unitarity** (§6.3.2–6.3.3). The non-unitarity of the boundary CFT may obstruct the Källén–Lehmann argument, but the gap itself (a statement about spectral support) may survive. CDT provides an independent check.
3. **The CEH bundling amplifies the gravitational θ -effect** (§6.3.4). Without CEH, the gravitational θ -effect is $\sim (\Delta H)^4 \sim 10^{-100} \text{eV}^4$ (unobservable). With CEH at gauge scale, the bundled effect is $\sim \Delta^4$ (potentially observable, though set to zero by gauged P/CP).
4. **Ashtekar gravity provides a structural reason for $\Delta_{\text{grav}} > 0$** (§6.4.2). Gravity as a gauge theory (Ashtekar $\text{SU}(2)$) inherits the gauge-theory mass gap structurally. The LQG area gap is a rigorous theorem; the gravitational spectral gap is conjectural but structurally motivated.
5. **The eta invariant is a diagnostic for the gravitational gap** (§6.5.2). On a manifold with boundary (e.g., finite-volume CDT), a finite eta invariant indicates a gravitational gap; a divergent one indicates its absence. The three-sector eta invariant requires all three gaps to be positive.

The main thesis of Section 6

[A]'s most serious open question — does $\Delta_{\text{grav}} > 0$ exist outside AdS/CFT? — is addressed by three independent approaches. CDT provides a **computational** answer in de Sitter (the gap is readable from lattice data). dS/CFT provides a **holographic** answer (conjectural, pending non-unitarity resolution). Ashtekar gravity provides a **structural** answer (gravity as gauge theory inherits the mass gap). The three approaches are complementary: CDT is numerical, dS/CFT is holographic, Ashtekar is structural. Their agreement would constitute strong evidence for $\Delta_{\text{grav}} > 0$ in our universe, from three independent lines of reasoning.

The eta invariant (§6.5) ties the gravitational gap to the boundary index theorem and the entanglement gap (§4.4), completing the circle: the gravitational gap controls the eta invariant, which controls the entanglement gap, which is a diagnostic for the mass gap. The entire chain — from the gauge mass gap to the gravitational gap to the entanglement gap — is unified in the three-sector Ashtekar framework.

7. The Bayesian Framework for Combined $\bar{\theta}$ Constraints

7.1 Introduction: from theoretical predictions to statistical inference

The previous sections established the theoretical framework: the Fisher metric identification (§2), the three-sector anomaly inflow (§3), the spectral diagnostics (§4), the CEH/resurgence research program (§5), and the gravitational gap (§6). These sections produce a coherent set of theoretical predictions:

- $\theta_{\text{QCD}} = \theta_{\text{grav}} = \theta_{\text{NY}} = 0$ (gauged P/CP, §3.5).
- $\chi \leq C \Delta^4$ ([A]'s Result A, reinterpreted as Fisher metric bound, §2).
- $\chi_{\text{grav}} \leq C_{\text{grav}} (\Delta_{\text{grav}})^4$ (gravitational analogue, §6).
- $\Delta_{\text{grav}}^{\text{dS}} \sim \Delta \cdot H$ (dS/CFT, §6.3, conjectural).
- The CEH measure problem determines whether the gravitational θ -effect is gauge-scale or Planck-suppressed (§5).

These predictions must be tested against experimental data. The data comes from multiple, independent channels:

1. **Neutron EDM:** the strongest direct constraint on $\bar{\theta}$.
2. **Collider experiments** (LHC/FCC): graviton-pomeron mixing constrains θ_{grav} .
3. **EIC gravitational form factors:** constrain the trace anomaly, related to χ .
4. **Lattice QCD:** direct computation of χ and Δ .
5. **Astrophysical observations:** neutron stars, gravitational waves, CMB.

No single channel is sufficient to test the full theoretical framework. The Bayesian framework combines all channels into a single posterior distribution, quantifying the evidence for or against each theoretical prediction. This section constructs the combined Bayesian framework, which is the statistical bridge between [A] and [B].

7.2 The effective CP-violating parameter and its constraints

7.2.1 The effective $\bar{\theta}$

The effective CP-violating parameter in QCD is:

$$\bar{\theta} = \theta_{\text{QCD}} + \arg \det M_q$$

where θ_{QCD} is the QCD θ -angle and $\arg \det M_q$ is the phase of the quark mass matrix. In the three-sector theory (§3), there are three effective parameters:

$$\bar{\theta}_{\text{QCD}} = \theta_{\text{QCD}} + \arg \det M_q, \quad \bar{\theta}_{\text{grav}} =$$

$$\theta_{\text{grav}}, \quad \bar{\theta}_{\text{NY}} = \theta_{\text{NY}}$$

The gauged P/CP mechanism (§3.5) predicts $\bar{\theta}_{\text{QCD}} = \bar{\theta}_{\text{grav}} = \bar{\theta}_{\text{NY}} = 0$. The axion solution (if it worked for all three sectors, which it does not — §3.7.4) would set $\bar{\theta}_{\text{QCD}} = 0$ but leave $\bar{\theta}_{\text{grav}}$ and $\bar{\theta}_{\text{NY}}$ unconstrained.

7.2.2 The experimental constraints

Neutron EDM. The neutron electric dipole moment (EDM) is the strongest direct constraint on $\bar{\theta}_{\text{QCD}}$:

$$d_n \approx 2.4 \times 10^{-16} \text{ e}\cdot\text{cm}$$

The current experimental bound (Abel et al., 2020) is:

$$|d_n| < 1.8 \times 10^{-26} \text{ e}\cdot\text{cm} \quad (90\%, \text{CL})$$

This gives:

$$|\bar{\theta}_{\text{QCD}}| < \frac{1.8 \times 10^{-26}}{2.4 \times 10^{-16}} \approx 7.5 \times 10^{-11} \approx 10^{-10}$$

The neutron EDM constrains only $\bar{\theta}_{\text{QCD}}$ (the gauge sector). It does not directly constrain $\bar{\theta}_{\text{grav}}$ or $\bar{\theta}_{\text{NY}}$, because the neutron EDM is generated by the QCD θ -term, not the gravitational or Nieh–Yan θ -terms (which couple to different operators).

Collider experiments (LHC/FCC). [B] (§3.1) proposes searching for graviton-pomeron mixing at the LHC and future FCC. The mixing parameter is:

$$\delta \sim \frac{\kappa, \Lambda_{\text{QCD}}^2}{M_{\text{Pl}}^2}$$

where κ is the gravitational coupling. The mixing is generated by the gravitational θ_{grav} -term $\theta_{\text{grav}}, \text{Tr}(R \wedge R)/(8\pi^2)$, which mixes the graviton and the pomeron (the QCD Regge trajectory) through the CEH instanton. The LHC sensitivity is $\delta \sim 10^{-3}$, corresponding to:

$$|\theta_{\text{grav}}| \lesssim \frac{\delta, M_{\text{Pl}}^2}{\kappa, \Lambda_{\text{QCD}}^2} \sim 10^{-3} \times \frac{M_{\text{Pl}}^2}{\Lambda_{\text{QCD}}^2} \sim 10^{33}$$

This is an extremely weak bound (the gravitational θ -term is very hard to constrain directly). However, if the CEH amplification (§6.3.4) brings the gravitational θ -effect to gauge scale, the LHC sensitivity to δ translates to a much stronger bound on the *bundled* θ -parameter. The Bayesian framework (§7.4) handles this through the CEH likelihood.

EIC gravitational form factors. [B] (§3.2) proposes measuring gravitational form factors at the Electron-Ion Collider (EIC). The gravitational form factors encode the trace anomaly $\theta^\mu_\mu = 4\epsilon - 3p$, which is related to χ through dimensional analysis:

$$\theta^\mu_\mu \sim \Lambda_{\text{QCD}}^4 \sim \chi$$

The EIC sensitivity to the trace anomaly constrains χ (and hence, through $\chi \leq C\Delta^4$, constrains Δ).

Lattice QCD. [B]'s lattice framework (§2.4) provides direct computations of:

- χ (from the topological charge fluctuations on the lattice),

- Δ (from the glueball mass extraction or the RMT diagnostic of §4.3),
- $\rho(\mu^2)$ (the spectral density, from the Källén–Lehmann representation).

These are theoretical predictions, not experimental data, but they enter the Bayesian framework as a "theory likelihood" — the probability of the lattice data given the theoretical model.

Astrophysical observations. Neutron star properties (mass-radius relation, tidal deformability) constrain the QCD equation of state, which depends on the trace anomaly and (indirectly) on Δ . Gravitational waves from the QCD phase transition (§8.2) constrain Δ cosmologically. CMB distortions (§8.3) constrain χ_{grav} .

7.2.3 The three-sector parameter space

The full parameter space of the three-sector theory is:

$$\boldsymbol{\theta} = (\bar{\theta}_{\text{QCD}}, \bar{\theta}_{\text{grav}}, \bar{\theta}_{\text{NY}}, \Delta, \Delta_{\text{grav}}, \chi, \chi_{\text{grav}}, \chi_{\text{NY}}, \chi_{\text{cross}}, K_{\text{CEH}})$$

where:

- $\bar{\theta}_i$ are the three effective θ -parameters (predicted to be zero by gauged P/CP).
- $\Delta, \Delta_{\text{grav}}$ are the gauge and gravitational mass gaps.
- $\chi, \chi_{\text{grav}}, \chi_{\text{NY}}$ are the three-sector topological susceptibilities.
- χ_{cross} is the cross-susceptibility (the off-diagonal elements of the 3×3 susceptibility matrix, §3.8).
- K_{CEH} is the CEH amplitude (gauge-scale or Planck-suppressed, §5.3).

The theoretical predictions constrain this space:

- Gauged P/CP: $\bar{\theta}_i = 0$ (a point prediction, or a delta-function prior).
- Result A: $\chi \leq C \Delta^4$ (an inequality constraint).
- dS/CFT: $\Delta_{\text{grav}} \sim \Delta \cdot H$ (a relation, conjectural).
- Cauchy–Schwarz: $\chi_{\text{cross}}^2 \leq \chi, \chi_{\text{grav}}$ ([A], §6.5).
- CEH: $K_{\text{CEH}} \in \{\text{gauge-scale, Planck-suppressed}\}$ (a binary hypothesis).

The Bayesian framework combines the theoretical predictions (as priors or model hypotheses) with the experimental data (as likelihoods) to produce posterior distributions.

7.3 The Bayesian framework

7.3.1 Bayes' theorem for the three-sector theory

The posterior distribution for the parameters $\boldsymbol{\theta}$ given the combined data $\mathcal{D}$ is:

$$\mathbb{P}(\boldsymbol{\theta} | \mathcal{D}) = \frac{\mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta})}{\int \mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}}$$

where:

- $\mathcal{L}(\mathcal{D} | \boldsymbol{\theta})$ is the **likelihood** — the probability of the data given the parameters.
- $\pi(\boldsymbol{\theta})$ is the **prior** — the probability distribution before seeing the data.
- $\mathcal{Z} = \int d\boldsymbol{\theta} \mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta})$ is the **evidence** (marginal likelihood) — the normalization.

7.3.2 The likelihood: combining independent channels

The likelihood factorizes over independent experimental channels (the data are independent given the parameters):

$$\mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) = \mathcal{L}\{\text{EDM}\}(\bar{\theta}_{\text{QCD}}) \times \mathcal{L}\{\text{LHC}\}(\bar{\theta}_{\text{grav}}), K_{\text{CEH}} \times \mathcal{L}\{\text{EIC}\}(\chi, \Delta) \times \mathcal{L}\{\text{lattice}\}(\chi, \Delta, \rho) \times \mathcal{L}\{\text{astro}\}(\Delta, \chi_{\text{grav}})$$

Each likelihood is a Gaussian (or a bounded Gaussian, for upper-limit constraints) centered on the measured value with the experimental uncertainty:

Neutron EDM likelihood:

$$\mathcal{L}\{\text{EDM}\}(\bar{\theta}_{\text{QCD}}) = \exp\left[-\frac{1}{2}\left(\frac{d_n(\bar{\theta}_{\text{QCD}}) - d_n^{\text{obs}}}{\sigma_{d_n}}\right)^2\right]$$

where $d_n(\bar{\theta}) = 2.4 \times 10^{-16} \text{ e}\cdot\text{cm}$ is the theoretical prediction, $d_n^{\text{obs}} = 0$ (no detection; the bound is an upper limit), and $\sigma_{d_n} = 1.8 \times 10^{-26} \text{ e}\cdot\text{cm}$ is the experimental uncertainty. This is a one-sided constraint: $|\bar{\theta}_{\text{QCD}}| < 10^{-10}$.

LHC likelihood (graviton-pomeron mixing):

$$\mathcal{L}\{\text{LHC}\}(\bar{\theta}_{\text{grav}}, K_{\text{CEH}}) = \exp\left[-\frac{1}{2}\left(\frac{\delta(\bar{\theta}_{\text{grav}}, K_{\text{CEH}}) - \delta^{\text{obs}}}{\sigma_{\delta}}\right)^2\right]$$

where $\delta(\bar{\theta}_{\text{grav}}, K_{\text{CEH}})$ is the predicted mixing parameter. If the CEH amplitude is gauge-scale ($K_{\text{CEH}} \sim \Lambda_{\text{QCD}}^4$), the mixing is $\delta \sim \kappa \Lambda_{\text{QCD}}^2 / M_{\text{Pl}}^2 \times \bar{\theta}_{\text{grav}}$, and the LHC bound $\delta < 10^{-3}$ constrains $\bar{\theta}_{\text{grav}}$. If the CEH amplitude is Planck-suppressed, the mixing is negligibly small, and the LHC gives no constraint. The likelihood thus **depends on the CEH hypothesis**.

EIC likelihood (trace anomaly / gravitational form factors):

$$\mathcal{L}\{\text{EIC}\}(\chi, \Delta) = \exp\left[-\frac{1}{2}\left(\frac{\theta^{\mu\mu}(\chi, \Delta) - \theta^{\text{obs}}}{\sigma_{\text{EIC}}}\right)^2\right]$$

where $\theta^{\mu\mu} \sim \Lambda_{\text{QCD}}^4 \sim \chi$ is the trace anomaly, and θ^{obs} is the measured value (future EIC data).

Lattice likelihood (theory likelihood):

$$\mathcal{L}_{\text{lattice}}(\chi, \Delta, \rho) = \exp\left[-\frac{1}{2}\left(\frac{\rho - \rho^{\text{obs}}}{\sigma_{\rho}}\right)^2\right]$$

$$\sum_i \left(\frac{\chi_i^{\{\text{lattice}\}}}{C(\Delta)} - \chi(\Delta) \right)^2$$

where $\chi_i^{\{\text{lattice}\}}$ are the lattice measurements of χ at different lattice spacings/volumes, and $\chi(\Delta) \leq C(\Delta)^4$ is the theoretical prediction. The lattice likelihood is a "theory likelihood" — it measures the agreement between the lattice data and the theoretical model (Result A).

Astrophysical likelihood:

$$\mathcal{L}^{\{\text{astro}\}}(\Delta, \chi^{\{\text{grav}\}}) = \mathcal{L}^{\{\text{NS}\}}(\Delta) \times \mathcal{L}^{\{\text{GW}\}}(\Delta) \times \mathcal{L}^{\{\text{CMB}\}}(\chi^{\{\text{grav}\}})$$

where the three factors are neutron star, gravitational wave, and CMB constraints (detailed in §8).

7.3.3 The prior: encoding theoretical predictions

The prior $\pi(\boldsymbol{\theta})$ encodes the theoretical predictions. We consider three prior scenarios:

Prior 1: Flat (uninformative). No theoretical preference. The prior on $\bar{\theta}_i$ is flat on $[0, 2\pi)$ (the periodicity range). The prior on Δ is flat on $[0, \infty)$ (or with a physical lower bound $\Delta > 0$). The prior on $K_{\{\text{CEH}\}}$ is a binary mixture: $\pi(K_{\{\text{CEH}\}}) = \frac{1}{2} \delta(K_{\{\text{CEH}\}} - \Lambda^4) + \frac{1}{2} \delta(K_{\{\text{CEH}\}} - \Lambda^4/M_{\{\text{Pl}\}}^n)$. This prior tests the theoretical predictions without bias.

Prior 2: Gauged P/CP (informed). The prior on $\bar{\theta}_i$ is a delta function at zero: $\pi(\bar{\theta}_i) = \delta(\bar{\theta}_i)$, reflecting the gauged P/CP prediction. The prior on Δ is informed by the lattice data: $\pi(\Delta) \propto \mathcal{L}^{\{\text{lattice}\}}(\Delta)$. The prior on $K_{\{\text{CEH}\}}$ is the same binary mixture. This prior tests the CEH hypothesis *given* the gauged P/CP prediction.

Prior 3: Axion (competing model). The prior on $\bar{\theta}_{\{\text{QCD}\}}$ is a delta function at zero (the axion sets $\bar{\theta}_{\{\text{QCD}\}} = 0$), but the priors on $\bar{\theta}_{\{\text{grav}\}}$ and $\bar{\theta}_{\{\text{NY}\}}$ are flat (the axion does not couple to these sectors, §3.7.4). The prior on $K_{\{\text{CEH}\}}$ is a delta function at the Planck-suppressed value (the axion scenario assumes the CEH is Planck-suppressed). This prior tests the axion model as a competing hypothesis.

The comparison of the three priors (via Bayes factors, §7.4) quantifies the evidence for the gauged P/CP model over the axion model.

7.3.4 The posterior and its interpretation

The posterior distribution $p(\boldsymbol{\theta} | \mathcal{D})$ combines the likelihood and the prior:

$$p(\boldsymbol{\theta} | \mathcal{D}) \propto \mathcal{L}^{\{\text{EDM}\}} \times \mathcal{L}^{\{\text{LHC}\}} \times \mathcal{L}^{\{\text{EIC}\}} \times \mathcal{L}^{\{\text{lattice}\}} \times \mathcal{L}^{\{\text{astro}\}} \times \pi(\boldsymbol{\theta})$$

The posterior has several key features:

The gauge sector ($\bar{\theta}_{\{\text{QCD}\}}$). The neutron EDM likelihood dominates: $|\bar{\theta}_{\{\text{QCD}\}}| < 10^{-10}$. With a flat prior, the posterior peaks at $\bar{\theta}_{\{\text{QCD}\}} = 0$ (the EDM bound is consistent with zero). With the gauged P/CP prior, the posterior is a delta function at zero (the prediction is confirmed). With the axion prior, the posterior

is also at zero (the axion sets $\bar{\theta}_{\text{QCD}} = 0$). **The gauge sector does not distinguish the models** — all three priors give $\bar{\theta}_{\text{QCD}} = 0$.

The gravitational sector ($\bar{\theta}_{\text{grav}}$). The LHC likelihood constrains $\bar{\theta}_{\text{grav}}$ only if K_{CEH} is gauge-scale (§7.3.2). If K_{CEH} is gauge-scale, the LHC bound $\delta < 10^{-3}$ gives $|\bar{\theta}_{\text{grav}}| < 10^{33}$ (very weak). If K_{CEH} is Planck-suppressed, the LHC gives no constraint. **The gravitational sector is currently unconstrained** — the Bayesian framework shows this explicitly.

The CEH parameter (K_{CEH}). The CEH hypothesis is tested by comparing the posterior probabilities of the gauge-scale and Planck-suppressed scenarios. If the EDM bound is violated (i.e., if the CEH generates a nonzero $\bar{\theta}_{\text{QCD}}$ at gauge scale), the posterior would show tension with the EDM data, favoring the Planck-suppressed scenario. If the EDM bound is satisfied (as it is), the data is consistent with both scenarios (the CEH could be gauge-scale but with $\theta = 0$ set by gauged P/CP, or Planck-suppressed with θ unconstrained). **The CEH hypothesis is currently undetermined** — the Bayesian framework quantifies this indeterminacy.

7.4 Model selection: Bayes factors for the gauged P/CP vs. axion

7.4.1 The Bayes factor

The **Bayes factor** is the ratio of evidences for two competing models:

$$B_{12} = \frac{\mathcal{Z}_1}{\mathcal{Z}_2} = \frac{p(\mathcal{D} | M_1)}{p(\mathcal{D} | M_2)}$$

where $\mathcal{Z}_i = \int d\theta \mathcal{L}(\mathcal{D} | \theta, M_i) \pi(\theta | M_i)$ is the evidence for model M_i . The Bayes factor quantifies the evidence for model 1 over model 2, independent of the priors (it marginalizes over the parameter space).

We compute the Bayes factor for three model comparisons:

7.4.2 Gauged P/CP vs. axion (gauge sector only)

Model $M_{\text{P/CP}}$ (gauged P/CP): $\bar{\theta}_{\text{QCD}} = 0$ (delta-function prior). The likelihood is $\mathcal{L}_{\text{EDM}}(0) = \exp[0] = 1$ (the EDM data is consistent with $\bar{\theta} = 0$). The evidence is:

$$\mathcal{Z}_{\text{P/CP}} = \mathcal{L}_{\text{EDM}}(0) \times \mathcal{L}_{\text{LHC}}(0, K_{\text{CEH}}) \times \dots = 1 \times (\text{other factors})$$

Model M_{axion} (axion): $\bar{\theta}_{\text{QCD}} = 0$ (the axion sets it to zero), but $\bar{\theta}_{\text{grav}}$ is free. The likelihood is the same as the gauged P/CP for the gauge sector. The evidence is:

$$\mathcal{Z}_{\text{axion}} = \mathcal{L}_{\text{EDM}}(0) \times \mathcal{L}_{\text{LHC}}(\text{free}, \bar{\theta}_{\text{grav}}, K_{\text{CEH}}^{\text{gauge}}) \times \dots$$

The Bayes factor is:

$$B_{\text{P/CP, axion}} = \frac{\mathcal{Z}_{\text{P/CP}}}{\mathcal{Z}_{\text{axion}}} = \frac{\mathcal{L}_{\text{LHC}}(0, K_{\text{CEH}}^{\text{gauge}})}{\int d\bar{\theta}_{\text{grav}} \mathcal{L}_{\text{LHC}}(\bar{\theta}_{\text{grav}}, K_{\text{CEH}}^{\text{gauge}})}$$

$$K_{\{\text{CEH}\}}^{\{\text{Planck}\}} / (2\pi)$$

The numerator is the LHC likelihood at $\bar{\theta}_{\{\text{grav}\}} = 0$ with gauge-scale CEH. The denominator is the LHC likelihood averaged over $\bar{\theta}_{\{\text{grav}\}}$ with Planck-suppressed CEH (which gives no constraint, so the average is ~ 1). The Bayes factor is:

$$B_{\{\text{P/CP, axion}\}} \sim \mathcal{L}_{\{\text{LHC}\}}(0, K_{\{\text{CEH}\}}^{\{\text{gauge}\}}) \sim 1$$

(both models predict $\bar{\theta}_{\{\text{QCD}\}} = 0$, and the LHC does not distinguish them because the gravitational sector is unconstrained). **The current data does not distinguish the gauged P/CP from the axion** — the Bayes factor is of order 1.

This is an important negative result: the Bayesian framework shows that the current experimental data, while consistent with $\bar{\theta} = 0$, does not distinguish the *mechanism* (gauged P/CP vs. axion). Distinguishing the models requires data that constrains the **gravitational sector** — either a detection of graviton-pomeron mixing at the LHC (favoring gauge-scale CEH and hence gauged P/CP) or a null result at much higher sensitivity (favoring Planck-suppressed CEH and hence the axion).

7.4.3 Gauge-scale CEH vs. Planck-suppressed CEH

Model $M_{\{\text{gauge-scale}\}}$: $K_{\{\text{CEH}\}} \sim \Lambda_{\{\text{QCD}\}}^4$. The gravitational θ -effect is gauge-scale, and the gauged P/CP is needed to set $\theta = 0$.

Model $M_{\{\text{Planck}\}}$: $K_{\{\text{CEH}\}} \sim \Lambda_{\{\text{QCD}\}}^4 / M_{\{\text{Pl}\}}^n$. The gravitational θ -effect is Planck-suppressed, and the axion suffices.

The Bayes factor is:

$$B_{\{\text{gauge, Planck}\}} = \frac{\mathcal{Z}_{\{\text{gauge-scale}\}}}{\mathcal{Z}_{\{\text{Planck}\}}} = \frac{p(\mathcal{D} | M_{\{\text{gauge-scale}\}})}{p(\mathcal{D} | M_{\{\text{Planck}\}})}$$

With the current data:

- The EDM bound is satisfied by both models (gauged P/CP sets $\theta = 0$ in the gauge-scale model; the axion sets $\theta_{\{\text{QCD}\}} = 0$ in the Planck model).
- The LHC gives no constraint on either model (the gravitational sector is too weakly coupled).
- The lattice data constrains χ and Δ , but not $K_{\{\text{CEH}\}}$.

Therefore:

$$B_{\{\text{gauge, Planck}\}} \sim 1$$

The current data does not distinguish the CEH hypotheses. **This is the quantitative version of [A]'s §8.9 decision table:** the CEH measure problem (§5) is not resolved by the current data. The Bayesian framework makes this precise — the Bayes factor is of order 1, and distinguishing the hypotheses requires new data (LHC at higher luminosity, or a CEH-specific observable).

7.4.4 Future projections

The Bayes factor can be projected for future experimental scenarios:

Scenario 1: LHC detects graviton-pomeron mixing at $\delta \sim 10^{-3}$. This would favor the gauge-scale CEH model (the mixing is too large for Planck suppression). The Bayes factor becomes:

$$B_{\{\text{gauge, Planck}\}} \sim \frac{\mathcal{L}_{\{\text{LHC}\}}(\delta^{\{\text{obs}\}})}{K_{\{\text{CEH}\}}^{\{\text{gauge}\}} \mathcal{L}_{\{\text{LHC}\}}(\delta^{\{\text{obs}\}})}$$

$$K\{\text{CEH}\}^{\{\text{Planck}\}} \gg 1$$

The evidence favors the gauge-scale model, supporting the gauged P/CP mechanism.

Scenario 2: LHC null result at $\delta < 10^{-5}$. This would push the bound on θ_{grav} lower (if the CEH is gauge-scale) or be consistent with Planck suppression. The Bayes factor remains ~ 1 unless the sensitivity reaches the Planck-suppressed scale (which is far beyond any foreseeable experiment).

Scenario 3: Next-generation EDM experiment (n2EDM) reaches $d_n < 10^{-28}$ e·cm. This would give $|\bar{\theta}_{\text{QCD}}| < 10^{-12}$. If the CEH is gauge-scale and generates $\bar{\theta}_{\text{QCD}} \neq 0$ (i.e., if the gauged P/CP is not exact or has corrections), the EDM bound could be violated, favoring the Planck model. If the bound is satisfied, the gauged P/CP is confirmed to higher precision, but the CEH hypothesis remains undetermined.

7.5 The Fisher metric as the prior on θ -space

7.5.1 The Fisher-informed prior

The Fisher metric identification (§2) provides a natural, **theory-informed prior** on θ -space. The Fisher metric $g_{ij} = \chi_{ij}/V$ is the natural Riemannian metric on the θ -manifold T^3 (the 3-torus of the three θ -parameters). The natural prior on a Riemannian manifold is the **Jeffreys prior**:

$$\pi_{\text{Jeffreys}}(\boldsymbol{\theta}) \propto \sqrt{\det g_{ij}(\boldsymbol{\theta})} = \sqrt{\det(\chi_{ij}/V)}$$

The Jeffreys prior is the **non-informative prior** on a Riemannian manifold — it is invariant under reparametrization and assigns equal probability to equal-volume regions of the manifold. For the θ -space with the Fisher metric, the Jeffreys prior is:

$$\pi_{\text{Jeffreys}}(\boldsymbol{\theta}) \propto \sqrt{\det \chi_{ij}}$$

This prior is **non-flat**: it assigns higher probability to regions where the Fisher metric is large (high χ) and lower probability to regions where the Fisher metric is small (low χ). The physical interpretation is: the Jeffreys prior assigns probability proportional to the **distinguishability** of the θ -vacua. Regions where the vacua are highly distinguishable (high χ) are assigned more prior probability, because the data can constrain them more tightly.

7.5.2 The mass gap and the Jeffreys prior

The mass gap enters the Jeffreys prior through the bound $\chi_{ij} \leq C\Delta^4$. The Jeffreys prior is bounded:

$$\pi_{\text{Jeffreys}} \propto \sqrt{\det \chi_{ij}} \leq (C\Delta^4)^{3/2} = C^{3/2} \Delta^6$$

The mass gap **regularizes the Jeffreys prior**: without the gap ($\Delta = 0$), $\chi \rightarrow \infty$ (or $\chi \rightarrow 0$, depending on the direction), and the Jeffreys prior is ill-defined. With the gap ($\Delta > 0$), χ is finite, and the Jeffreys prior is well-defined.

This is the Bayesian version of the information-theoretic statement of §2: **the mass gap makes the θ -vacuum finitely distinguishable**, and this translates to the **Jeffreys prior being well-defined**. The mass gap is the condition for the Bayesian framework to be applicable.

7.5.3 The posterior with the Fisher-informed prior

With the Jeffreys prior, the posterior is:

$$p(\boldsymbol{\theta} | \mathcal{D}) \propto \mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) \times \sqrt{\det \chi_{ij}(\boldsymbol{\theta})}$$

The Jeffreys prior modifies the posterior in regions where χ varies significantly. Near $\boldsymbol{\theta} = 0$ (where the gauged P/CP predicts the vacuum), χ_{ij} is at its maximum (the vacuum is stiffest), and the Jeffreys prior enhances the posterior probability. Far from $\boldsymbol{\theta} = 0$, χ_{ij} decreases (the vacuum becomes softer), and the Jeffreys prior suppresses the posterior.

The effect is small if χ is approximately constant on θ -space (which it is, to leading order, since χ is the second derivative of $E(\theta)$ and the higher derivatives are suppressed). But the principle is important: the Fisher metric provides a **natural, theory-informed prior** that is not flat but is determined by the physics (the topological susceptibility matrix).

7.6 The Cramér–Rao bound as a Bayesian convergence rate

7.6.1 The Fisher information and the posterior width

The Fisher information matrix (§2) determines the **minimum width** of the posterior distribution. The Cramér–Rao bound (§2.3) gives:

$$\text{Cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{F}^{-1} = \frac{1}{4V} \boldsymbol{\chi}^{-1}$$

where $\mathbf{F} = 4V \boldsymbol{\chi}$ is the QFI matrix (§2.5). The posterior covariance cannot be smaller than $\mathbf{F}^{-1}$ — the Fisher information sets the **resolution limit** of the Bayesian inference.

In the Bayesian framework, this means:

$$\text{Cov}(\boldsymbol{\theta} | \mathcal{D}) \gtrsim \frac{1}{4V} \boldsymbol{\chi}^{-1}$$

The posterior width is bounded below by the inverse Fisher information. The mass gap enters through $\chi \leq C \Delta^4$: the larger the gap, the larger χ , the smaller the minimum posterior width, and the more precisely $\boldsymbol{\theta}$ can be determined.

7.6.2 The mass gap as the Bayesian resolution limit

Combining the Cramér–Rao bound with the Fisher metric identification:

$$\Delta\theta_{\min} \sim \frac{1}{\sqrt{F_Q}} \sim \frac{1}{2\sqrt{V\chi}} \sim \frac{1}{2\sqrt{V}\Delta^2}$$

This is the **Bayesian resolution limit**: no amount of data can determine θ to better precision than $\Delta\theta_{\min}$, because the Fisher information is finite (bounded by the mass gap). This is the information-theoretic floor on the Bayesian inference.

The current EDM bound ($|\bar{\theta}| < 10^{-10}$) is far above this floor ($\Delta\theta_{\min} \sim 1$ for $V \sim (5 \text{ GeV}^{-1})^4$, §2.3.3). The EDM bound is not Fisher-limited — it is limited by the experimental sensitivity (the EDM measurement precision), not by the quantum distinguishability of the θ -vacua.

The Fisher floor becomes relevant only when the experimental sensitivity approaches $\delta\theta_{\min} \sim 1/\sqrt{\chi} \sim 1/\Delta^2$. For $\Delta \sim \Lambda_{\text{QCD}} \sim 200$ MeV, this is $\delta\theta_{\min} \sim 1/(0.2, \text{GeV})^2 \sim 25, \text{GeV}^{-2} \sim O(1)$ (in natural units). The EDM bound is 10 orders of magnitude below this floor, meaning the Bayesian inference is far from the quantum limit.

7.7 Summary of Section 7

The Bayesian framework

Component	What it does	Key formula	Mass gap role
Likelihood (§7.3.2)	Combines EDM + LHC + EIC + lattice + astro	$\mathcal{L} = \prod_k \mathcal{L}_k$	Lattice constrains $\chi \leq C\Delta^4$
Prior (§7.3.3)	Encodes theoretical predictions	$\pi(\boldsymbol{\theta})$ (flat, P/CP, or axion)	Gauged P/CP: $\theta = 0$
Posterior (§7.3.4)	Combined inference	$p(\boldsymbol{\theta})$	$\mathcal{D}) \propto \mathcal{L} \times \pi$
Bayes factor (§7.4)	Model selection	$B = \mathcal{Z}_1 / \mathcal{Z}_2$	Distinguishes gauged P/CP vs. axion
Jeffreys prior (§7.5)	Fisher-informed prior	$\pi \propto \sqrt{\det \chi_{ij}}$	Mass gap makes prior well-defined
Cramér–Rao bound (§7.6)	Resolution limit	$\delta\theta_{\min} \sim 1/\sqrt{\chi}$	Mass gap sets the floor

Key results

1. **The current data does not distinguish the gauged P/CP from the axion** (§7.4.2). The Bayes factor $B_{\{\text{P/CP, axion}\}} \sim 1$: both models predict $\bar{\theta}_{\text{QCD}} = 0$, and the gravitational sector is unconstrained. Distinguishing them requires data that constrains $\bar{\theta}_{\text{grav}}$ — the LHC graviton-pomeron mixing or a CEH-specific observable.
2. **The CEH hypothesis is currently undetermined** (§7.4.3). The Bayes factor $B_{\{\text{gauge, Planck}\}} \sim 1$: the EDM bound is satisfied by both the gauge-scale and Planck-suppressed scenarios. This is the quantitative version of [A]’s §8.9 decision table.
3. **The Fisher metric provides a natural Jeffreys prior** on θ -space (§7.5), and the mass gap makes this prior well-defined. The prior is non-flat: it assigns probability proportional to the distinguishability of the θ -vacua, which is controlled by χ (and hence by Δ).
4. **The Cramér–Rao bound is the Bayesian resolution limit** (§7.6): $\delta\theta_{\min} \sim 1/\sqrt{\chi} \sim 1/\Delta^2$. The current EDM bound (10^{-10}) is far above this limit

(~ 1), meaning the Bayesian inference is far from the quantum floor.

5. **Future experimental scenarios** (§7.4.4) can shift the Bayes factors: an LHC detection of graviton-pomeron mixing would favor the gauge-scale CEH (and hence the gauged P/CP); a next-generation EDM at 10^{-28} e·cm would test the gauged P/CP to higher precision.

The main thesis of Section 7

The Bayesian framework is the statistical bridge between [A] and [B]. It combines the theoretical predictions (gauged P/CP, $\chi \leq C\Delta^4$, CEH hypothesis) with the experimental data (EDM, LHC, EIC, lattice, astrophysical) into a single posterior distribution. The current data is consistent with the gauged P/CP prediction ($\theta = 0$) but does not distinguish it from the axion model — the gravitational sector is the distinguishing feature, and it is currently unconstrained. The Fisher metric (§2) enters the Bayesian framework in two ways: as the Jeffreys prior on θ -space (a theory-informed prior) and as the Cramér–Rao resolution limit (the quantum floor on parameter estimation). The mass gap makes both well-defined: without the gap, the Fisher metric is infinite (or zero), the Jeffreys prior is ill-defined, and the Cramér–Rao bound is trivial.

8. Cosmological and Astrophysical Tests

8.1 Introduction: the mass gap as a cosmological observable

The previous sections established the theoretical framework (§2–§6) and the statistical inference framework (§7) for the combined gauge + gravitational + Nieh–Yan theory. The Bayesian framework (§7) showed that the current experimental data (neutron EDM, LHC, lattice QCD) is consistent with the gauged P/CP prediction ($\theta = 0$) but does not distinguish it from the axion model, because the gravitational sector is unconstrained.

This section identifies **cosmological and astrophysical observables** that can test the mass gap and the three-sector theory independently of terrestrial experiments. The key insight is that the mass gap Δ controls the thermodynamics of the QCD phase transition in the early universe, and the topological susceptibility χ controls the energy release during that transition. These quantities leave imprints on:

1. **Gravitational waves** from the QCD phase transition (§8.2): the mass gap controls the transition order and the latent heat, which generate a stochastic gravitational wave background detectable by pulsar timing arrays.
2. **Neutron star equation of state** (§8.3): the trace anomaly $\theta^\mu_\mu = 4\epsilon - 3p$ is nonzero because of confinement (the mass gap) and is constrained by NICER and LIGO/Virgo observations.
3. **CMB spectral distortions** (§8.4): the energy released during the QCD epoch creates a μ -type distortion in the CMB blackbody, bounded by $\chi \leq C\Delta^4$.
4. **The analytic S-matrix and pomeron–graviton mixing** (§8.5): the Reggeization of the graviton and its mixing with the QCD pomeron provide a model-independent collider constraint on θ_{grav} .

Each observable probes Δ or χ through a different physical mechanism, providing

independent cosmological tests of the mass gap. Together with the lattice (§4) and terrestrial experiments (§7), these form a multi-messenger program: the same mass gap is tested from the lattice (microscopic), through collider experiments (terrestrial), to cosmological observations (macroscopic).

8.2 Gravitational waves from the QCD phase transition

8.2.1 The QCD phase transition in the early universe

At $T \sim 150 \text{ MeV}$ ($z \sim 10^{12}$, $t \sim 10^{-5} \text{ s}$ after the Big Bang), the universe underwent the confinement/deconfinement transition: quarks and gluons, previously in a deconfined quark-gluon plasma (QGP), became confined into hadrons. The nature of this transition depends on the QCD parameters:

- **At zero baryon density** (the physically relevant case for the early universe, since the baryon-to-photon ratio $\eta \sim 10^{-9}$ is tiny): lattice QCD predicts a **crossover** (smooth transition, not a true phase transition) for physical quark masses.
- **At finite baryon density** (relevant for neutron stars and possibly the early universe if there is a first-order transition at high density): the transition can be **first-order** (discontinuous, with bubble nucleation).

The **mass gap controls the nature of the transition**. If $\Delta > 0$ (confinement, gapped phase), the theory has a well-defined confined phase separated from the deconfined phase by the transition. If $\Delta = 0$ (no confinement, gapless), the transition is smooth (no phase transition, just a crossover in a conformal theory). The existence of the mass gap is what makes the transition a genuine physical event, not a smooth crossover.

8.2.2 The gravitational wave signal

A first-order phase transition generates a stochastic gravitational wave background (GWB) through three mechanisms (Caprini et al., 2020):

1. **Bubble nucleation**: the first-order transition proceeds via nucleation of bubbles of the confined phase in the deconfined plasma. The collision of bubble walls generates gravitational waves.
2. **Sound waves**: the bulk motion of the plasma, driven by the bubble expansion, generates sound waves that source gravitational waves.
3. **Turbulence**: the turbulent motion of the plasma in the wake of the bubble walls generates gravitational waves.

The gravitational wave amplitude is:

$$\Omega_{\text{GW}}(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d\log f} \sim \frac{H^4}{\rho_c} \left(\frac{T_*}{H} \right)^4 v_w^3 \kappa^2 \left(\frac{\alpha}{1 + \alpha} \right)^2$$

where:

- $T_* \sim 150 \text{ MeV}$ is the transition temperature,
- H is the Hubble parameter at the transition,
- v_w is the bubble wall velocity,
- κ is the efficiency factor (fraction of latent heat converted to kinetic energy),

- $\alpha = \Delta\rho / \rho_{\text{rad}}$ is the ratio of latent heat to radiation energy density.

The **latent heat** $\Delta\rho$ is the energy difference between the deconfined and confined phases. It is controlled by the topological susceptibility χ and the mass gap Δ :

$$\Delta\rho \sim \chi \cdot T_* \sim \Delta^4 \cdot T_* / \Lambda_{\text{QCD}}$$

where the first relation uses the fact that the topological susceptibility measures the free energy cost of changing the topological sector (which is what the phase transition does), and the second uses $\chi \leq C\Delta^4$ ([A]'s Result A).

8.2.3 The mass gap from gravitational waves

The gravitational wave amplitude is bounded by [A]'s Result A:

$$\Omega_{\text{GW}} \lesssim \frac{H^4}{\rho_c} \left(\frac{T_*}{H} \right)^4 v_w^3 \kappa^2 \left(\frac{C\Delta^4 \cdot T_*}{\Lambda_{\text{QCD}}} \right)^2 \frac{\rho_{\text{rad}}}{\rho_c \Lambda_{\text{QCD}}^2 \rho_{\text{rad}}^2} \sim \frac{\Delta^8 T_*^6}{\rho_c \Lambda_{\text{QCD}}^2 \rho_{\text{rad}}^2}$$

This connects [A]'s $\chi \leq C\Delta^4$ to a **direct cosmological observable**: the gravitational wave amplitude is bounded by Δ^8 (through the square of $\chi \sim \Delta^4$), and the detection (or upper limit) constrains Δ from above.

The peak frequency of the gravitational wave signal is:

$$f_* \sim \frac{H}{a_*} \sim 10^{-9} \text{ Hz} \times \left(\frac{T_*}{150 \text{ MeV}} \right) \left(\frac{g_*}{10.75} \right)^{1/6}$$

where a_* is the scale factor at the transition and g_* is the effective number of relativistic degrees of freedom. This frequency is in the **pulsar timing array (PTA) band** ($f \sim 10^{-9}$ Hz).

8.2.4 Pulsar timing arrays

Pulsar timing arrays (NANOGrav, EPTA, PPTA, IPTA) are sensitive to gravitational waves in the 10^{-9} – 10^{-7} Hz band. The NANOGrav 12.5-year data (and the more recent 15-year data) reported evidence for a stochastic process consistent with a gravitational wave background, though its origin (astrophysical vs. cosmological) is debated.

The QCD phase transition signal is expected to be **subdominant** at zero baryon density (where the transition is a crossover, not first-order, and the latent heat is small). However:

- If the transition is first-order at some effective baryon density (e.g., due to lepton asymmetry or other effects), the signal is stronger.
- The **mass gap controls the strength of the crossover**: even a crossover has a characteristic time scale $\tau \sim 1/\Delta$, and the energy release is $\sim \chi T_*$. The gravitational wave signal from a crossover is weaker than from a first-order transition but may still be detectable with next-generation PTAs (SKA-PTA).

The key point is: **a detection of a stochastic gravitational wave background at $f \sim 10^{-9}$ Hz, correlated with the QCD epoch, would be cosmological evidence for the mass gap** — independent of lattice QCD and collider experiments. A non-detection (upper limit) constrains Δ from above, complementing the lattice QCD lower bound.

8.2.5 The gravitational θ -sector and the QCD phase transition

In the three-sector theory, the QCD phase transition also involves the gravitational and Nieh–Yan θ -parameters. The gravitational θ -term $\theta_{\text{grav}}, \text{Tr}(R \wedge R)$ contributes to the free energy of the deconfined and confined phases differently:

$$F(\theta_{\text{grav}}) = F_0 + \frac{1}{2} \chi_{\text{grav}}, \theta_{\text{grav}}^2 + \cdots$$

The gravitational topological susceptibility χ_{grav} changes across the QCD phase transition (it is suppressed in the deconfined phase, as χ is). The gravitational θ -term therefore contributes to the latent heat:

$$\Delta \rho \supset \frac{1}{2} \Delta \chi_{\text{grav}}, \theta_{\text{grav}}^2$$

where $\Delta \chi_{\text{grav}} = \chi_{\text{grav}}^{\text{confined}} - \chi_{\text{grav}}^{\text{deconfined}}$. If the CEH amplification (§6.3.4) brings χ_{grav} to gauge scale, this contribution is significant; if Planck-suppressed, it is negligible. The gravitational wave signal thus encodes the CEH hypothesis, providing a cosmological test of the CEH measure problem (§5).

8.3 Neutron stars and the trace anomaly

8.3.1 The trace anomaly as the order parameter for confinement

The **trace anomaly** of QCD is:

$$\theta^\mu{}_\mu = \frac{\beta(g)}{2g} \text{Tr}(F_{\mu\nu} F^{\mu\nu}) + (1 + \gamma_m) \sum_q m_q \bar{\psi}_q \psi_q$$

This quantity is nonzero because the coupling g runs (the beta function $\beta(g) \neq 0$). It vanishes if the theory is conformal ($\beta = 0$, $\gamma_m = 0$) — i.e., if the theory is gapless. The **nonzero trace anomaly is evidence for the mass gap**: it measures the deviation from conformality.

The vacuum expectation value of the trace anomaly is:

$$\langle \theta^\mu{}_\mu \rangle = 4\epsilon - 3p$$

where ϵ is the energy density and p is the pressure. At zero temperature and zero density (pure YM), $\langle \theta^\mu{}_\mu \rangle \sim \Lambda_{\text{QCD}}^4 \sim \chi \sim \Delta^4$.

8.3.2 The neutron star equation of state

The neutron star equation of state (EOS) — the relation $p(\epsilon)$ between pressure and energy density — is determined by QCD at nuclear density ($\rho \sim 2\text{--}5\rho_0$, where $\rho_0 \sim 0.16 \text{ fm}^{-3}$ is nuclear saturation density). The trace anomaly at nuclear density is:

$$\langle \theta^\mu{}_\mu \rangle_{\text{NS}} = 4\epsilon(\rho) - 3p(\rho)$$

This quantity is nonzero (confinement persists at nuclear density) and its magnitude is controlled by the mass gap Δ . The neutron star mass-radius relation $M(R)$ encodes $4\epsilon - 3p$ through the Tolman–Oppenheimer–Volkoff (TOV) equation.

Recent observations constrain the EOS:

- **NICER** (Neutron Star Interior Composition Explorer): measures neutron star radii to $\sim 10\%$ precision. PSR J0030+0451 gives $R \sim 12\text{--}13$ km; PSR J0740+6620 gives $R \sim$

13\$–14\$ km.

- **LIGO/Virgo** gravitational wave observations: GW170817 measures the tidal deformability $\Lambda \sim 200\text{--}400$, constraining the EOS (stiffer EOS $\rightarrow$ larger Λ).

8.3.3 The mass gap from neutron stars

The connection is:

$$\langle \theta^{\mu}_{\mu} \rangle \sim \Delta^4 \quad \xrightarrow{\text{at nuclear density}} \quad 4\epsilon(\rho) - 3p(\rho) \quad \xrightarrow{\text{TOV}} \quad M(R)$$

The mass gap Δ (at zero density, pure YM) is extrapolated to finite density through the trace anomaly. The neutron star mass-radius relation constrains the trace anomaly at nuclear density, and the extrapolation to zero density constrains Δ .

The extrapolation is the key subtlety. The mass gap Δ is defined at **zero temperature and zero density** (pure YM), while the neutron star trace anomaly is at **zero temperature and finite density** (nuclear matter with quarks). The two are related but not identical. The interpolation is controlled by:

1. **Chiral perturbation theory** ($\rho \rightarrow 0$): the trace anomaly is $\sim \Lambda_{\text{QCD}}^4$ (the zero-density value).
2. **Perturbative QCD** ($\rho \rightarrow \infty$): the trace anomaly is $\sim \alpha_s^2 \rho^{4/3}$ (asymptotic freedom).
3. **Lattice QCD** (intermediate ρ): finite-density lattice QCD (currently limited by the sign problem, but improving).

The mass gap Δ sets the scale for the zero-density trace anomaly, and the neutron star EOS constrains the finite-density trace anomaly. The two are connected through the density interpolation, which is controlled by the same QCD dynamics (the running coupling, the confinement scale). A precise measurement of the EOS, combined with the lattice QCD computation of Δ , can close the loop: the EOS tests the mass gap at finite density, and the lattice tests it at zero density.

8.3.4 The Fisher metric at finite density

The Fisher metric identification (§2) extends to finite density. The θ -vacuum at finite density has a modified topological susceptibility:

$$\chi(\rho) = \chi|_{\theta=0, \rho} = \frac{\partial^2 E(\theta, \rho)}{\partial \theta^2} \bigg|_{\theta=0}$$

The Fisher metric on the θ -manifold at finite density is $g_{\theta\theta}(\rho) = V\chi(\rho)$, and the QFI is $F_Q(\rho) = 4V\chi(\rho)$. The mass gap at finite density $\Delta(\rho)$ controls $\chi(\rho)$ through the same Källén–Lehmann bound: $\chi(\rho) \leq C\Delta(\rho)^4$.

The neutron star EOS constrains $\chi(\rho)$ (through the trace anomaly), and the Fisher metric gives $\chi(\rho) = F_Q(\rho)/(4V)$. The mass gap at finite density $\Delta(\rho)$ is thus bounded by the neutron star observations. This extends the Fisher metric framework (§2) to the finite-density regime, connecting the cosmological observations (neutron stars) to the information-theoretic framework.

8.4 CMB spectral distortions and the QCD epoch

8.4.1 CMB spectral distortions

The CMB blackbody spectrum is not perfect: deviations from a perfect blackbody (spectral distortions) encode the thermal history of the universe. There are two main types:

- **μ -type distortion:** generated by energy injection at $z \sim 5 \times 10^4$, when photon production and Compton scattering go out of equilibrium (the μ -era). The distortion is characterized by the chemical potential μ :
$$n(\nu) \propto \frac{1}{e^{(h\nu - \mu)/kT} - 1}$$
- **y -type distortion:** generated by energy injection at $z \lesssim 10^4$, when Compton scattering is inefficient (the y -era, related to the Sunyaev–Zeldovich effect).

The current bound is $|\mu| < 9 \times 10^{-5}$ (FIRAS, 1993). Next-generation experiments (PIXIE, proposed Voyage 2050) aim to measure μ with precision $\delta\mu \sim 10^{-8}$.

8.4.2 The QCD contribution to the μ -distortion

The QCD phase transition ($T \sim 150$ MeV, $z \sim 10^{12}$) releases energy into the photon-baryon plasma. This energy injection creates a μ -distortion:

$$\mu_{\text{QCD}} \sim \frac{\Delta E_{\text{QCD}}}{\rho_\gamma} \sim \frac{\chi \cdot T_*^4}{\rho_\gamma}$$

where:

- ΔE_{QCD} is the energy released during the QCD transition,
- $\rho_\gamma \sim T_\gamma^4$ is the photon energy density at the μ -era,
- χ is the topological susceptibility (controlling the free energy cost of the topological sector change).

Using [A]'s bound $\chi \leq C\Delta^4$:

$$\mu_{\text{QCD}} \lesssim \frac{C\Delta^4 \cdot T_*^4}{\rho_\gamma}$$

This gives an **upper bound on the μ -distortion from the QCD epoch in terms of the mass gap**. A CMB measurement of μ (or an upper limit) constrains Δ from above — cosmologically and independently of all other methods.

8.4.3 Numerical estimate

For $\Delta \sim \Lambda_{\text{QCD}} \sim 200$ MeV, $T_* \sim 150$ MeV, $\rho_\gamma \sim T_0^4 \times (1 + z_\mu)^4$ where $T_0 = 2.7$ K and $z_\mu \sim 5 \times 10^4$:

$$\mu_{\text{QCD}} \sim \frac{(0.2, \text{GeV})^4 \cdot (0.15, \text{GeV})^4}{(2.7 \times 10^{-4}, \text{eV}) \times 5 \times 10^4} \sim \frac{(0.2)^4 \cdot (0.15)^4, \text{GeV}^8}{(13.5, \text{eV})^4}$$

$$\sim \frac{1.6 \times 10^{-4} \times 5.1 \times 10^{-4}, \text{GeV}^8}{10^5, \text{eV}^4} \sim \frac{8 \times 10^{-8}, \text{GeV}^8}{3.3 \times 10^5 \times (10^{-9})^4, \text{GeV}^4}$$

$$\sim \frac{8 \times 10^{-8}}{3.3 \times 10^5 \times 10^{-36}} \text{GeV}^4 \sim \frac{8 \times 10^{-8}}{3.3 \times 10^{-31}} \text{GeV}^4 \sim 2.4 \times 10^{22} \text{GeV}^4$$

This crude estimate is far above the FIRAS bound ($|\mu| < 10^{-4}$), indicating that the QCD contribution to the μ -distortion is **subdominant** and the simple estimate is not the full story. The actual QCD contribution is suppressed by additional factors:

- The QCD transition at zero baryon density is a **crossover**, not a first-order transition. The energy release is not a latent heat (a discontinuous jump) but a smooth change, and the μ -distortion from a smooth crossover is smaller.
- The energy injection at the QCD epoch ($z \sim 10^{12}$) is **redshifted** before the μ -era ($z \sim 5 \times 10^4$). The redshift factor is $(z_{\mu}/z_{\text{QCD}})^4 \sim (5 \times 10^4 / 10^{12})^4 \sim 10^{-31}$, which suppresses the contribution enormously.

With these suppression factors, the QCD contribution to μ is estimated (Bianchini & Silk, 2020s) to be:

$$\mu_{\text{QCD}} \sim 10^{-10} \text{ to } 10^{-12}$$

This is well below the FIRAS bound (10^{-4}) and at the edge of the PIXIE sensitivity (10^{-8}). A detection of $\mu_{\text{QCD}} \sim 10^{-10}$ would be a **cosmological measurement of the mass gap** (through $\chi \sim \Delta^4$), while a non-detection at the PIXIE level would constrain Δ from above.

8.4.4 The gravitational θ -sector and the CMB

The gravitational θ -term contributes to the energy release during the QCD epoch through the gravitational topological susceptibility:

$$\Delta E_{\text{grav}} \sim \chi_{\text{grav}} \cdot T_*^4 \sim (\Delta_{\text{grav}})^4 \cdot T_*^4 \sim (\Delta \cdot H)^4 \cdot T_*^4$$

where we used $\Delta_{\text{grav}}^4 \sim \Delta \cdot H$ (§6.3.2). This contribution is suppressed by $H^4 \sim (10^{-33} \text{eV})^4$ and is completely negligible. However, if the CEH amplification (§6.3.4) brings χ_{grav} to gauge scale, the contribution is $\sim \Delta^4 \cdot T_*^4$, which is the same order as the gauge contribution. The CMB μ -distortion thus encodes the CEH hypothesis, providing a cosmological test.

In practice, the CMB μ -distortion from the QCD epoch is dominated by the gauge sector (the QCD thermodynamics). The gravitational sector contribution is negligible unless the CEH is gauge-scale, in which case it is of the same order. A precise CMB measurement could, in principle, distinguish the two scenarios — but the required precision ($\sim 10^{-12}$) is far beyond any foreseeable experiment.

8.5 The analytic S-matrix and pomeron–graviton mixing

8.5.1 The pomeron and the graviton as Regge trajectories

[B]'s Experiment 1 (§3.1) searches for **graviton-pomeron mixing** at the LHC and future FCC. The pomeron is a Regge trajectory $\alpha(t)$ with intercept $\alpha_P(0) \approx 1.08$, describing the dominant contribution to high-energy diffractive scattering in QCD. The graviton, in the Regge language, is the t -channel exchange of a spin-2 Regge trajectory $\alpha_{\text{grav}}(t)$.

The **analytic S-matrix** program (Chew, Mandelstam, 1960s; revived by Caron-Huot, Mizera, et al., 2020s) treats all particles as Regge trajectories and derives constraints from:

- **Crossing symmetry** (relates s -channel and t -channel amplitudes),
- **Unitarity** (the S-matrix is unitary, constraining the Regge trajectory),
- **Analyticity** (the amplitude is analytic in the complex angular momentum plane).

8.5.2 The graviton-pomeron mixing

In the analytic S-matrix framework, the graviton is not a fundamental particle but a **Regge trajectory** $\alpha_{\text{grav}}(t)$. The graviton-pomeron mixing is the **intersection** of the graviton Regge trajectory with the pomeron trajectory. The mixing parameter is:

$$\delta \sim \frac{\kappa, \Lambda_{\text{QCD}}^2}{M_{\text{Pl}}^2}$$

where κ is the gravitational coupling. If $\delta \neq 0$ is detected, it measures the gravitational coupling at the QCD scale — directly constraining α_{grav} and the CEH instanton amplitude.

The mixing arises because the graviton and the pomeron both couple to the stress-energy tensor. The QCD θ -term generates a gravitational coupling through the trace anomaly (§8.3), and the gravitational θ -term generates a QCD coupling through the CEH instanton (§5.3). The mixing is the **physical manifestation of the three-sector coupling**: the gauge and gravitational sectors are coupled through the CEH instanton, and this coupling is measurable in diffractive scattering.

8.5.3 Model-independent bounds from the analytic S-matrix

The analytic S-matrix provides **model-independent** bounds on δ from crossing and unitarity — a purely theoretical prediction that [B]'s experiment can test. The bounds are:

- **Crossing**: the s -channel amplitude (related to the total cross-section, measured at the LHC) and the t -channel amplitude (related to the diffractive cross-section, where the pomeron/graviton mixing appears) are related by crossing symmetry. A measurement of the total cross-section constrains the diffractive cross-section, and hence δ .
- **Unitarity**: the Regge trajectory intercept $\alpha(0) \leq 1$ (the Froissart bound). If the graviton-pomeron mixing shifts the pomeron intercept, unitarity constrains the shift.

These bounds are **independent of the CEH hypothesis**: they constrain the mixing δ regardless of whether it comes from the gauge-scale or Planck-suppressed CEH. The analytic S-matrix provides a **purely kinematic** constraint on the three-sector coupling.

8.5.4 The three-sector mixing

In the three-sector theory, the pomeron (gauge sector) and the graviton (gravitational sector) are coupled not only through the CEH instanton but also through the Nieh–Yan sector. The torsion introduces an additional mixing:

$$\alpha_{\text{total}} = \alpha_{\text{grav}} + \delta_{\text{NY}} \sim \frac{\kappa, \Lambda_{\text{QCD}}^2}{M_{\text{Pl}}^2}, \alpha_{\text{grav}} + \frac{\alpha_{\text{NY}}, \Lambda_{\text{QCD}}^2}{M_{\text{torsion}}^2}, \alpha_{\text{NY}}$$

where α_{NY} is the torsion coupling and M_{torsion} is the torsion scale. The

Nieh–Yan mixing is suppressed by the torsion scale (which may be the Planck scale or a lower scale, depending on the torsion dynamics). The total mixing is the sum of the gravitational and Nieh–Yan contributions, and the LHC measurement constrains both.

8.6 Summary of Section 8

Cosmological and astrophysical observables

Observable	What it constrains	Mechanism	Current bound	Future sensitivity	Mass gap connection
Gravitational waves (§8.2)	Δ from above	QCD phase transition $\rightarrow$ stochastic GWB	NANOGrav 10^{-15}	SKA-PTA 10^{-17}	$\Omega_{\text{GW}} \sim \Delta^8$
Neutron star EOS (§8.3)	$\Delta(\rho)$ at finite density	Trace anomaly $4\epsilon - 3p \sim \Delta^4$	NICER $R \sim 13$ km; LIGO $\Lambda \sim 300$	$\sim 5\%$ EOS	$\langle \theta^\mu_\mu \rangle \sim \Delta^4$
CMB μ-distortion (§8.4)	Δ from above	QCD energy release $\rightarrow \mu \sim \chi T_*^4/\rho_{\text{ma}}$	FIRAS	μ	$< 10^{-4}$
Pomeron-graviton mixing (§8.5)	θ_{grav} , θ_{NY}	Regge trajectory intersection	LHC $\delta \sim 10^{-3}$	FCC $\delta \sim 10^{-5}$	Three-sector coupling

Key results

- Gravitational waves from the QCD phase transition** (§8.2) provide a cosmological test of the mass gap. The gravitational wave amplitude is $\Omega_{\text{GW}} \sim \Delta^8$ (through $\chi \sim \Delta^4$), and pulsar timing arrays (NANOGrav, SKA-PTA) are sensitive to the relevant frequency ($\sim 10^{-9}$ Hz). A detection would be cosmological evidence for $\Delta > 0$; an upper limit constrains Δ from above.
- Neutron star EOS** (§8.3) constrains the trace anomaly $\langle \theta^\mu_\mu \rangle = 4\epsilon - 3p \sim \Delta^4$ at finite density. NICER and LIGO/Virgo observations, combined with lattice QCD at zero density, can close the loop and constrain Δ through the density interpolation.
- CMB μ -distortion** (§8.4) constrains Δ from above through the QCD energy release $\mu \sim \chi T_*^4$. The QCD contribution is $\mu \sim 10^{-10}$ – 10^{-12} , at the edge of PIXIE sensitivity. A detection would be a cosmological measurement of the mass gap.

4. **Pomeron-graviton mixing** (§8.5) constrains θ_{grav} and θ_{NY} through the three-sector coupling at the LHC/FCC. The analytic S-matrix provides model-independent bounds from crossing and unitarity. This is the only observable that directly constrains the gravitational and Nieh–Yan θ -parameters.
5. **The CEH hypothesis is encoded in the cosmological observables**: if the CEH is gauge-scale, the gravitational θ -sector contributes to the QCD phase transition latent heat (§8.2.5) and the CMB μ -distortion (§8.4.4) at the same level as the gauge sector. If Planck-suppressed, these contributions are negligible. The cosmological observables thus provide an indirect test of the CEH measure problem (§5).

The main thesis of Section 8

The mass gap Δ is not only a microscopic quantity (measured on the lattice, §4) or a terrestrial observable (measured at the LHC, §7). It is also a **cosmological observable**: it controls the QCD phase transition thermodynamics, the neutron star EOS, and the CMB spectral distortions. These cosmological observables provide independent tests of $\Delta > 0$, complementing the lattice and collider measurements. Together, the multi-messenger program — lattice (§4), collider (§7, §8.5), gravitational waves (§8.2), neutron stars (§8.3), and CMB (§8.4) — tests the mass gap from microscopic to cosmological scales, and the three-sector θ -parameters through the pomeron-graviton mixing (§8.5) and the CEH contributions to the cosmological observables.

9. Open Problems

9.1 Introduction

The preceding sections (§2–§8) constructed a unified framework connecting [A]'s index-theorem/topological approach and [B]'s constructive-experimental approach through five layers of bridging theories: Fisher information geometry (§2), 5D anomaly inflow (§3), spectral diagnostics (§4), CEH/resurgence (§5), gravitational gap (§6), Bayesian inference (§7), and cosmological tests (§8). Seven novel claims were advanced:

1. The 5D theory is the anomaly-inflow bulk (§3).
2. The mass gap is the spectral gap of the Floer differential (§4.2).
3. χ is the Fisher information metric on θ -space (§2).
4. The CEH measure problem is a Borel-plane singularity problem (§5).
5. Δ_{grav} is computable via CDT in de Sitter (§6.2).
6. The mass gap is readable from the entanglement spectrum (§4.4).
7. The mass gap has cosmological signatures (§8).

Each of these claims is supported by a chain of reasoning, but each also rests on assumptions, approximations, or unproven conjectures. This section catalogues the open problems — the gaps between what has been established and what remains to be proven — organized by the section in which they arise. We distinguish between:

- **Type I (tractable):** problems that can be addressed with existing tools and moderate effort.
- **Type II (hard):** problems that require substantial new mathematical or computational development.
- **Type III (conceptual):** problems that require a new idea or framework, not merely more computation.

9.2 Open problems in the Fisher metric identification (§2)

9.2.1 The higher-order θ -derivatives and the curvature of θ -space

Problem. The Fisher metric $g_{\theta\theta} = \chi$ is the second derivative of the vacuum energy $E(\theta)$. The curvature of the θ -manifold involves higher-order derivatives (the Riemann tensor on the Fisher manifold requires derivatives of $g_{\theta\theta}$, i.e., third and fourth θ -derivatives of E). These are related to 4-point and higher correlators of the topological charge density $q(x)$.

Status: Type II. The higher-order correlators are well-defined but not computed. The curvature of θ -space is a new quantity — it encodes the nonlinear response of the vacuum to θ -perturbations. If the curvature is large (near a phase transition), the Fisher metric approximation (using only the second derivative) breaks down.

What is needed: Lattice computation of $\partial^3 E / \partial \theta^3$ and $\partial^4 E / \partial \theta^4$ (or equivalently, the 3-point and 4-point functions of $q(x)$). This is a standard lattice measurement but has not been performed with the θ -dependence in mind.

9.2.2 The Bures metric and the mixed-state generalization

Problem. The Fisher metric identification (§2.1) uses the pure-state QFI. At finite temperature or in a mixed state (e.g., the thermal state of the early universe), the vacuum is not pure, and the Bures metric (the mixed-state generalization) is the relevant quantity. The relation $\chi = F_Q / (4V)$ may be modified at finite temperature.

Status: Type I (tractable). The finite-temperature generalization is straightforward: the QFI of a thermal state is $\chi(T) = \partial^2 F / \partial \theta^2|_{\theta=0}$, where F is the free energy. The Fisher metric at finite temperature is $\chi(T)$, which is the temperature-dependent topological susceptibility (already computed on the lattice). The mass gap at finite temperature $\Delta(T)$ controls $\chi(T)$ through the same Källén–Lehmann bound: $\chi(T) \leq C, \Delta(T)^4$.

What is needed: Explicit verification that the finite-temperature Fisher metric equals the lattice-measured $\chi(T)$, and that the bound $\chi(T) \leq C, \Delta(T)^4$ holds. This is a straightforward check.

9.3 Open problems in the 5D anomaly inflow (§3)

9.3.1 The Nieh–Yan cutoff dependence

Problem. The Nieh–Yan 4-form $\mathcal{N} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$ has a coefficient that depends on the UV cutoff (the torsion coupling is dimensionful, $[T^a] =$

mass , so $\int \mathcal{N}$ is dimensionful). The θ_{NY} parameter is therefore cutoff-dependent, unlike θ_{QCD} and θ_{grav} (which are dimensionless angles).

Status: Type II (discussed in §3.6). The gauged P/CP mechanism is structural (symmetry-based, not coefficient-dependent), so it forces $\theta_{\text{NY}} = 0$ regardless of the cutoff. But the physical *effect* of θ_{NY} (if it were nonzero) would depend on the cutoff. The 5D theory provides a natural cutoff (the compactification scale $1/R$), but the matching between the 5D and 4D descriptions at the cutoff is not fully worked out.

What is needed: A careful renormalization-group analysis of the Nieh–Yan term in the 5D theory, showing that the cutoff dependence is absorbed into the 5D coupling constants. This is a technical but not conceptual problem.

9.3.2 The dynamics of the 5D theory

Problem. [B]'s 5D theory is used to prove the mass gap (dynamical sector) and to set $\theta = 0$ (topological sector). The mass gap proof relies on the polymer expansion, the Sturm–Liouville analysis, and the lattice regularization. The topological sector relies on the anomaly inflow and the gauged P/CP. The two sectors are decoupled (§3.4.1), but the **5D dynamics** — the full non-perturbative definition of the 5D theory — is not established.

Status: Type II. [B] defines the 5D theory on $\mathbb{R}^{3,1} \times S^1$ with specific boundary conditions and takes $R \rightarrow 0$. The non-perturbative definition in the $R \rightarrow 0$ limit is the content of [B]'s proof. The open question is whether the 5D theory is **well-defined at finite R** — i.e., whether the polymer expansion converges and the Sturm–Liouville analysis is valid before taking $R \rightarrow 0$.

What is needed: A proof of the convergence of the polymer expansion at finite R (not just in the $R \rightarrow 0$ limit), and a demonstration that the Sturm–Liouville spectrum is stable under the compactification.

9.3.3 The torsion dynamics

Problem. The torsion T^a is introduced as a background in §3 (the Nieh–Yan sector). If the torsion is **dynamical** (has its own kinetic term and fluctuations), the analysis changes: the torsion fluctuations contribute to the path integral and may modify the mass gap proof and the anomaly inflow.

Status: Type II. The torsion-free case (standard Riemannian geometry) is well-understood. The constant-torsion case (as in §3) is tractable. The dynamical-torsion case is an open problem: the torsion kinetic term introduces new degrees of freedom (the torsion waves), whose spectrum must be gapped for the mass gap proof to survive.

What is needed: An analysis of the torsion fluctuations in the 5D theory, showing that the torsion sector is gapped (or that the torsion is frozen by the dynamics, as in the "contorsion" scenario where the torsion is determined by the spin density).

9.4 Open problems in the spectral diagnostics (§4)

9.4.1 The Floer-to-Wightman map

Problem. The Floer homology identification (§4.2) shows that the mass gap is the spectral gap of the Floer differential on $\mathcal{A}/\mathcal{G}$. But the Floer chain complex is a topological object

— it does not directly produce the Wightman functions (the physical observables of QFT). The translation from the Floer chain complex to physical observables requires constructing a "**Floer-to-Wightman map**": does the Floer chain complex reconstruct the Wightman functions?

Status: Type III (conceptual). This is the deepest open problem of §4. The Floer homology is a topological invariant; the Wightman functions are dynamical. The two are related (both encode the spectrum), but the map between them is not known. The Donaldson–Witten TQFT (§4.2.7) provides a partial answer (the TQFT observables are Donaldson invariants, not Wightman functions), but the full map is open.

What is needed: A construction of the Wightman functions from the Floer chain complex, or a proof that the Floer spectral gap determines the Wightman spectral gap (the lowest mass in the Wightman spectral representation). This is a mathematical physics problem of considerable depth.

9.4.2 The RMT–Result A quantitative connection

Problem. The RMT diagnostic (§4.3) identifies the mass gap as the crossover scale where Dirac eigenvalue statistics depart from the chGUE universal predictions. The connection to [A]’s Result A ($\chi \leq C\Delta^4$) goes through the Banks–Casher relation and the index theorem. But the **quantitative** relation — does the RMT universal spacing predict the coefficient C in $\chi \leq C\Delta^4$? — is not established.

Status: Type II. The Banks–Casher relation gives $\chi = \pi \rho_{\text{Dirac}}(0)$, and the mass gap controls the crossover scale. But the coefficient C in the bound $\chi \leq C\Delta^4$ depends on the spectral density $\rho(\mu^2)$, which is not predicted by RMT (RMT predicts the low-lying eigenvalue statistics, not the spectral density at the gap scale).

What is needed: A lattice study correlating the RMT crossover scale with the lattice-measured χ and Δ , fitting the coefficient C . This is a quantitative lattice physics project.

9.4.3 The entanglement gap on the lattice

Problem. The entanglement gap diagnostic (§4.4) requires computing the entanglement entropy $S(R)$ for a strip of width R in the vacuum state. On the lattice, this requires the replica trick or tensor network methods, both of which are computationally expensive.

Status: Type I (tractable but expensive). The replica trick for lattice QCD is feasible but expensive: one computes the n -th Rényi entropy $S_n = -\frac{1}{n-1} \log \text{Tr}(\rho_R^n)$ for integer $n \geq 2$ and extrapolates to $n \rightarrow 1$. The computation requires n copies of the system, which is a factor of n more expensive than a standard simulation. The entanglement entropy of the Yang–Mills vacuum has been computed in simplified models (e.g., 2D YM, free field theory) but not in full 4D lattice QCD.

What is needed: A 4D lattice QCD computation of the entanglement entropy $S(R)$ for several values of R , fitting the exponential decay $S(R) - S_{\text{area}} \sim e^{-\Delta R}$, and extracting Δ . This is a challenging but well-defined lattice computation.

9.5 Open problems in the CEH/resurgence (§5)

9.5.1 The 4D Yang–Mills trans-series

Problem. All verified resurgence results are in 2D models ($\mathbb{CP}^{N-1}$), principal chiral

model, Gross–Neveu). The 4D Yang–Mills trans-series is **unknown** — the perturbative series for the vacuum energy, the glueball mass, or the topological susceptibility is known only to a few orders, and the instanton calculus is non-supersymmetric (no localization).

Status: Type II (hard). The 4D instanton calculus is technically formidable (the instanton moduli space is 8-dimensional, and the fluctuation determinant involves infinite-dimensional functional determinants). Without supersymmetry, there is no exact instanton calculus.

What is needed: High-order perturbative computations of the QCD vacuum energy (or the topological susceptibility), to identify the Borel-plane singularities from the large-order behavior. This is a multi-loop perturbative QCD project, feasible in principle but requiring substantial effort.

9.5.2 The CEH Borel-plane singularity

Problem. The CEH instanton is a **combined gauge + gravitational** saddle. Its Borel-plane singularity (§5.3) is a novel object — no existing calculation addresses the Borel-plane structure of a combined gauge + gravitational saddle. The 2D analogs (instantons in $\mathbb{CP}^{N-1}$) are gauge-only; the gravitational instantons (Eguchi–Hanson) are gravity-only. The combined case is unexplored.

Status: Type III (conceptual). The combined saddle has no analog in the existing resurgence literature. The fluctuation determinant around the CEH instanton involves both gauge and gravitational fluctuations, coupled through the torsion. The computation is a new type of instanton calculus.

What is needed: The one-loop fluctuation determinant around the CEH instanton, computed in the combined gauge + gravitational + torsion background. This is the central calculation of §5 (Prong 1 of the research program, §5.7.1).

9.5.3 The combined gauge + gravitational lattice

Problem. The lattice computation of the CEH fluctuation determinant (§5.5.4, Prong 2 of the research program) requires a **combined gauge + gravitational lattice** — a triangulation with both gauge connections (on the edges) and gravitational connections (the Regge calculus geometry). This code does not exist.

Status: Type II. The ingredients — lattice QCD (gauge sector) and Regge calculus / CDT (gravitational sector) — are mature, but the combined code is a significant software engineering project. The challenge is the matching: the gauge connections live on the links of the triangulation, and the gravitational connections (the metric, the curvature) live on the simplices. The two must be consistently coupled.

What is needed: A lattice code that simultaneously evolves the gauge fields (Yang–Mills) and the geometry (Regge/CDT), with the CEH instanton as a background. This is a multi-year computational physics project.

9.6 Open problems in the gravitational gap (§6)

9.6.1 The CDT topological susceptibility

Problem. Computing χ_{grav} on CDT backgrounds (§6.2.5) requires defining the Pontryagin density $\text{Tr}(R \wedge R)$ on a simplicial manifold. The Regge calculus version of $\text{Tr}(R \wedge R)$ exists (via deficit angles), but has not been used for susceptibility calculations.

Status: Type I (tractable). The Regge calculus definition of the Pontryagin number is standard (it is a sum over deficit angles at the hinges). The susceptibility is the variance of this sum. The computation is a straightforward extension of existing CDT measurements (which measure the volume profile and the spectral dimension).

What is needed: A CDT simulation measuring the Pontryagin number distribution and its variance, on the same configurations used for the volume fluctuation analysis. This is a well-defined lattice gravity computation.

9.6.2 The dS/CFT spectral representation

Problem. The Källén–Lehmann argument ([A]'s §2, used to prove Result A) requires positivity of the spectral density $\rho(\mu^2) \geq 0$. In the non-unitary dS/CFT boundary theory (§6.3.3), the spectral density may not be positive, and Result A may fail. Does the gravitational bound $\chi_{\text{grav}} \leq C(\Delta_{\text{grav}})^4$ survive in de Sitter?

Status: Type III (conceptual). The non-unitarity of the dS/CFT boundary theory is a fundamental obstacle. The mass gap itself (the existence of $\Delta > 0$) does not require positivity, but the bound on χ does.

What is needed: Either (a) a proof that the Källén–Lehmann positivity survives in the non-unitary dS/CFT case (perhaps because the bulk theory is unitary even if the boundary is not), or (b) an alternative derivation of the bound that does not use positivity. The CDT computation (§6.2) provides an independent check: if CDT finds $\chi_{\text{grav}} \leq C(\Delta_{\text{grav}})^4$ numerically, the bound holds in de Sitter regardless of the dS/CFT non-unitarity.

9.6.3 The Ashtekar mass gap

Problem. The Ashtekar identification (§6.4.2) provides a structural reason for $\Delta_{\text{grav}} > 0$: gravity as a gauge theory inherits the mass gap. But the Ashtekar gravity has additional constraints (diffeomorphism invariance, the Hamiltonian constraint) that the Yang–Mills gauge theory does not have. The structural argument is not a proof.

Status: Type III (conceptual). The Ashtekar formulation makes gravity look like a gauge theory, but the Hamiltonian constraint ($H = 0$ on physical states) changes the spectral problem. The "mass gap" in the Ashtekar theory is not the gap above the vacuum of H (which is zero by constraint), but the gap in the physical excitations. The relation between the two is not straightforward.

What is needed: A rigorous definition of the gravitational spectral gap in the Ashtekar framework, and a proof (or disproof) that it is positive. This is a quantum gravity problem of considerable depth. The LQG area gap is rigorous, but the spectral gap (the graviton mass) is not.

9.7 Open problems in the Bayesian framework (§7)

9.7.1 The likelihood for the gravitational sector

Problem. The Bayesian framework (§7.3.2) includes a likelihood for the LHC graviton-pomeron mixing, but this likelihood is extremely weak: the LHC bound on θ_{grav} is $\sim 10^{33}$ (§7.3.2). The gravitational sector is effectively unconstrained.

Status: Type I (tractable, pending experiment). The likelihood is well-defined; it is simply uninformative because the experimental sensitivity is too low. Improving the sensitivity requires either

higher luminosity at the LHC/FCC or a new observable that couples more strongly to θ_{grav} .

What is needed: Either (a) a higher-luminosity LHC/FCC measurement of graviton-pomeron mixing, or (b) identification of a new observable that constrains θ_{grav} more directly (e.g., a gravitational wave observable that encodes χ_{grav} , §8.2.5).

9.7.2 The prior on the CEH parameter

Problem. The Bayesian framework treats the CEH amplitude K_{CEH} as a binary hypothesis (gauge-scale or Planck-suppressed). In reality, K_{CEH} may be a continuous parameter, and the "gauge-scale vs. Planck-suppressed" distinction is a simplification.

Status: Type II. The binary hypothesis is a useful first approximation, but a more refined analysis would treat K_{CEH} as a continuous parameter with a prior that peaks at the gauge-scale and Planck-suppressed values (a bimodal prior). The Bayes factor (§7.4) would then marginalize over the continuous parameter.

What is needed: A continuous parameterization of K_{CEH} and a bimodal prior, replacing the binary hypothesis. This is a statistical refinement, not a new physical insight.

9.8 Open problems in the cosmological tests (§8)

9.8.1 The QCD phase transition at zero baryon density

Problem. The QCD phase transition at zero baryon density is expected to be a **crossover**, not a first-order transition. The gravitational wave signal from a crossover is much weaker than from a first-order transition (no bubble nucleation, no latent heat). The connection between the mass gap and the gravitational wave signal (§8.2) is therefore indirect.

Status: Type II. The crossover still has a characteristic time scale $\tau \sim 1/\Delta$ and releases energy $\sim \chi T^*$, but the gravitational wave generation mechanism is different (no bubble collision, just plasma turbulence and sound waves). The signal is expected to be ~ 10 – 100 times weaker than the first-order case.

What is needed: A detailed computation of the gravitational wave signal from the QCD crossover, including the mass gap dependence. This is a numerical relativity + cosmology project.

9.8.2 The density interpolation for neutron stars

Problem. The neutron star EOS constrains the trace anomaly at nuclear density ($\rho \sim 2$ – $5\rho_0$), while the mass gap is defined at zero density. The interpolation between the two densities is controlled by QCD, but it is not known with sufficient precision to extract Δ from the neutron star data.

Status: Type II. The interpolation requires either (a) finite-density lattice QCD (limited by the sign problem), or (b) a theoretical model of the density dependence (e.g., a chiral effective theory extrapolation from low density + perturbative QCD from high density). The two approaches do not meet cleanly at nuclear density.

What is needed: Improved finite-density lattice QCD (or a non-lattice method, such as the thimble method of §5.5) to bridge the gap between zero density and nuclear density. This is an active area of

lattice QCD research.

9.8.3 The CMB μ -distortion sensitivity

Problem. The QCD contribution to the CMB μ -distortion is $\sim 10^{-10}$ – 10^{-12} , which is at the edge of the PIXIE sensitivity ($\delta\mu \sim 10^{-8}$). A detection requires either a more sensitive experiment or a stronger QCD signal.

Status: Type I (tractable, pending experiment). The QCD μ -distortion is a standard prediction; the sensitivity is an experimental limitation. Next-generation CMB experiments (beyond PIXIE) may reach the required sensitivity.

What is needed: A CMB experiment with $\delta\mu \sim 10^{-12}$ sensitivity (well beyond PIXIE). This is a proposed future mission (Voyage 2050 or similar).

9.9 Summary of open problems

#	Problem	Section	Type	Tractability
1	Floer-to-Wightman map	§4.2	III (conceptual)	Needs new idea
2	RMT $\rightarrow$ coefficient C in $\chi \leq C\Delta^4$	§4.3	II (hard)	Lattice study
3	Entanglement gap on 4D lattice	§4.4	I (tractable)	Expensive lattice computation
4	4D YM trans-series	§5.2	II (hard)	Multi-loop perturbation theory
5	CEH Borel-plane singularity	§5.3	III (conceptual)	New instanton calculus
6	Combined gauge + gravity lattice	§5.5	II (hard)	Multi-year code development
7	CDT topological susceptibility	§6.2	I (tractable)	Straightforward CDT extension
8	dS/CFT spectral representation	§6.3	III (conceptual)	Positivity in non-unitary theory
9	Ashtekar mass gap	§6.4	III (conceptual)	Quantum gravity problem
10	Nieh–Yan cutoff dependence	§3.6	II (hard)	RG analysis

11	5D dynamics at finite R	§3.4	II (hard)	Polymer expansion at finite R
12	Torsion dynamics	§3.2	II (hard)	Torsion fluctuation analysis
13	QCD crossover gravitational wave signal	§8.2	II (hard)	Numerical relativity
14	Density interpolation for neutron stars	§8.3	II (hard)	Finite-density lattice QCD
15	CMB μ -distortion sensitivity	§8.4	I (tractable)	Next-gen CMB experiment

The most pressing problems

The three most pressing open problems are:

1. **The CEH Borel-plane singularity** (#5, Type III): this is the central calculation of [A], reframed by §5 as a resurgence problem. It is the pivot of the axion vs. gauged P/CP decision, and it requires a new type of instanton calculus (combined gauge + gravitational + torsion).
2. **The Floer-to-Wightman map** (#1, Type III): this is the deepest conceptual problem. Without it, the Floer homology identification is a structural insight, not a computational tool. The map would convert the Floer spectral gap into a physical mass.
3. **The dS/CFT spectral representation** (#8, Type III): this determines whether [A]'s bound $\chi \leq C\Delta^4$ survives in de Sitter space. If the positivity fails, the gravitational bound in our universe is unproven. The CDT computation provides an independent check, but the dS/CFT proof is needed for the holographic argument.

The three most tractable problems are:

1. **The CDT topological susceptibility** (#7, Type I): a straightforward extension of existing CDT measurements.
2. **The entanglement gap on the 4D lattice** (#3, Type I): expensive but well-defined.
3. **The CMB μ -distortion sensitivity** (#15, Type I): pending next-generation experiments.

The tractable problems provide **short-term tests** of the framework; the conceptual problems define the **long-term research program**.

10. Conclusions

10.1 What has been achieved

This paper has constructed a unified framework connecting two existing works on the Yang–Mills mass gap: [A] (Addendum, arXiv:2510.18951-related), which uses index theorems and differential geometry to unify the mass gap, Strong CP, and quantum gravity; and [B] (Proposal, arXiv:2506.00284), which uses a 5D gauge theory with torsion, polymer expansions, Sturm–Liouville analysis, and lattice QCD to

constructively prove the mass gap and propose experimental tests.

The bridge between [A] and [B] is built through five layers of meritorious but non-mainstream theories, each with a rigorous foundation but not previously brought to bear on the mass gap problem in a unified way:

Layer I (Spectral-Topological): Random matrix theory and the Banks–Casher relation (§4.3), the Gutzwiller–Selberg trace formula (§5.6), and resurgence/trans-series (§5.2). These connect the Dirac spectrum ([B]'s lattice tool) and the instanton structure ([A]'s topological framework) to the mass gap through spectral analysis.

Layer II (Constructive-Computational): Large- N volume independence (§5.4.4), Lefschetz thimble decomposition (§5.5), persistent homology (bridge document, §II.3), and Symanzik improvement / finite-size scaling (bridge document, §II.4). These provide the rigorous machinery to convert [B]'s constructive proof and lattice computations into the continuum, infinite-volume quantities that [A]'s theorems require.

Layer III (Geometric-Gravitational): Floer homology (§4.2), CDT (§6.2), dS/CFT (§6.3), anomaly inflow (§3.4), and the APS eta invariant (§6.5). These extend [A]'s gravitational sector beyond AdS/CFT to the de Sitter universe and provide the structural identification of the mass gap as a spectral gap of a chain complex.

Layer IV (Statistical-Inferential): Information geometry / Fisher metric (§2), Bayesian inference (§7), and the Bisognano–Wichmann theorem / entanglement gap (§4.4). These convert [A]'s abstract bounds into statements about experimental distinguishability and combine all experimental constraints into a single posterior.

Layer V (Astro-Particle-Experimental): The analytic S-matrix / pomeron–graviton connection (§8.5), gravitational waves from the QCD phase transition (§8.2), neutron star EOS (§8.3), and CMB spectral distortions (§8.4). These connect the mass gap to cosmological observables, extending the experimental program to the macroscopic scale.

10.2 The seven novel claims

Seven novel claims were advanced, prioritized by tractability:

1. **The 5D theory is the anomaly-inflow bulk** (§3, claim strength: high). [B]'s 5D gauge theory with torsion is not just a proof tool — it is the anomaly-inflow bulk that gauges the three Chern–Weil global symmetries. The S^1 holonomies are the θ -parameters; gauged P/CP (orientation reversal) forces all three to zero. This is the paper's headline result: the mass gap and $\theta = 0$ are decoupled consequences of the same 5D theory.
2. **χ is the Fisher information metric** (§2, claim strength: high). The topological susceptibility is exactly the Fisher metric on θ -space, per unit volume. The mass gap makes the QFI finite. The Cramér–Rao bound $\Delta\theta_{\min} \sim 1/\sqrt{\chi} \sim 1/\Delta^2$ is the quantum floor on θ -sensitivity.
3. **The mass gap is the spectral gap of the Floer differential** (§4.2, claim strength: moderate). [B]'s Sturm–Liouville operator and [A]'s index theorem are two representations of the same chain complex. The Floer spectral gap is the mass gap. This is a structural insight, pending the Floer-to-Wightman map (§9.4.1).
4. **The CEH measure problem is a Borel-plane singularity problem** (§5, claim strength: moderate). Resurgence reframes [A]'s central open calculation as computable in principle, by

analyzing the trans-series singularity structure. The three-pronged research program (analytic Borel-plane, lattice Regge+gauge, Lefschetz thimbles) offers three independent paths.

5. Δ_{grav} is computable via CDT in de Sitter (§6.2, claim strength: moderate). CDT produces an emergent de Sitter geometry where the gravitational gap is readable from volume fluctuation correlations. This addresses [A]'s hardest open question without relying on AdS/CFT.
6. The mass gap is readable from the entanglement spectrum (§4.4, claim strength: moderate). The Bisognano–Wichmann theorem gives a non-spectral determination of Δ from the entanglement entropy decay $e^{-\Delta R}$, connecting to [A]'s island formula.
7. The mass gap has cosmological signatures (§8, claim strength: moderate). Gravitational waves from the QCD epoch, neutron star EOS, and CMB μ -distortions provide tests independent of terrestrial experiments. The pomeron–graviton mixing directly constrains θ_{grav} .

10.3 The three-sector framework

The paper's most significant structural result is the **three-sector unification** (§3):

- The **gauge sector** (θ_{QCD}): controlled by $\text{Tr}(F \wedge F)$, measured by the neutron EDM, bounded by $\chi \leq C \Delta^4$.
- The **gravitational sector** (θ_{grav}): controlled by $\text{Tr}(R \wedge R)$, measured (in principle) by the pomeron–graviton mixing, bounded by $\chi_{\text{grav}} \leq C_{\text{grav}}(\Delta_{\text{grav}})^4$.
- The **Nieh–Yan sector** (θ_{NY}): controlled by $\mathcal{N} = T^a \wedge T_a - R_{ab} \wedge e^a \wedge e^b$, introduced by the torsion, with cutoff-dependent coefficient but structural gauged P/CP.

The three sectors are coupled through the CEH instanton (§3.7, §5.3), which bundles all three θ -parameters. The gauged P/CP mechanism forces all three to zero — a result **stronger** than the original gauged P/CP (which addressed only θ_{QCD}) and **stronger** than the axion solution (which cannot address θ_{grav} or θ_{NY} because the axion couples only to $\text{Tr}(F \wedge F)$, §3.7.4).

The mass gap (proved from the dynamical sector of the 5D theory) and the $\theta = 0$ result (proved from the topological sector) are **decoupled** (§3.4.1). The mass gap is the condition that makes the θ -vacuum finitely distinguishable (§2); the $\theta = 0$ result is the consequence of the gauged P/CP symmetry. The two are independent, but the mass gap is **necessary for the well-posedness** of the CEH measure problem (§5.7.3) and the Jeffreys prior on θ -space (§7.5).

10.4 The multi-messenger experimental program

The framework predicts a multi-messenger experimental program, in which the mass gap and the three-sector θ -parameters are tested from microscopic to cosmological scales:

Scale	Observable	What it tests	Section
-------	------------	---------------	---------

Microscopic (lattice)	Glueball mass, Dirac spectrum, entanglement entropy	$\Delta > 0$; $\chi \leq C\Delta^4$	§4
Terrestrial (collider)	Pomeron–graviton mixing at LHC/FCC	θ_{grav} ; CEH hypothesis	§7, §8.5
Terrestrial (EDM)	Neutron EDM (n2EDM)	$\bar{\theta}_{\text{QCD}}$; gauged P/CP	§7
Astrophysical (NS)	Neutron star mass-radius; tidal deformability	$\Delta(\rho)$ at finite density	§8.3
Cosmological (GW)	Stochastic GW background at $f \sim 10^{-9}$ Hz	Δ from the QCD phase transition	§8.2
Cosmological (CMB)	CMB μ -distortion at $\mu \sim 10^{-10}$	Δ from the QCD energy release	§8.4

The Bayesian framework (§7) combines all these constraints into a single posterior distribution, quantifying the evidence for each theoretical prediction. The current data is consistent with the gauged P/CP prediction ($\theta = 0$) but does not distinguish it from the axion model — the gravitational sector is the distinguishing feature, and it requires future experimental data.

10.5 The role of non-mainstream theories

The bridge between [A] and [B] uses theories that are rigorous but non-mainstream in the mass gap context: Floer homology (pure mathematics), random matrix theory (lattice QCD), resurgence (2D QFT), CDT (lattice gravity), dS/CFT (holography), the Bisognano–Wichmann theorem (AQFT), information geometry (statistics), Bayesian inference (cosmology), and the analytic S-matrix (amplitudes program). Each is well-established in its own field but has not been brought to bear on the mass gap problem.

The use of non-mainstream theories is a **feature, not a bug**. The mass gap problem has resisted mainstream approaches for over 60 years (the Clay Millennium Prize remains unclaimed). The bridge framework brings new mathematical and physical structures to bear, each providing a different perspective on the same quantity Δ . The agreement of the independent diagnostics (Floer, RMT, entanglement, Fisher, cosmological) would constitute evidence for $\Delta > 0$ from structures that have nothing to do with each other — a form of "consilience" that is the hallmark of a robust physical result.

The non-mainstream theories are used with appropriate caveats: Floer homology is a structural insight (not a computational tool, §9.4.1), resurgence in 4D is unproven (§9.5.1), dS/CFT is non-unitary (§9.6.2), and the Bayesian framework is currently uninformative for the gravitational sector (§9.7.1). The open problems (§9) are stated openly, with a tractability assessment for each.

10.6 The path forward

The most immediate experimental and computational tests are:

1. **Reanalyze existing lattice data** for the RMT crossover scale (§4.3.6): no new simulations needed.
2. **Compute the CDT topological susceptibility** (§6.2.5, §9.6.1): a straightforward CDT extension.
3. **Measure the entanglement gap** on 4D lattice QCD (§4.4, §9.4.3): expensive but well-defined.
4. **Search for pomeron–graviton mixing** at the LHC with higher luminosity (§8.5, §7.4.4): tests θ_{grav} .
5. **Measure the neutron EDM** at 10^{-28} e·cm (nEDM, §7.4.4): tests $\bar{\theta}_{\text{QCD}}$ to higher precision.

The most pressing theoretical problems are:

1. **The CEH Borel-plane singularity** (§5.3, §9.5.2): the central open calculation, reframed as a resurgence problem.
2. **The Floer-to-Wightman map** (§4.2, §9.4.1): the deepest conceptual problem.
3. **The dS/CFT spectral representation** (§6.3, §9.6.2): determines whether the gravitational bound holds in our universe.

The framework provides a roadmap: the short-term tests (1–5) can be performed with existing tools; the long-term problems (1'–3') require new mathematical or computational development. The bridge between [A] and [B] is not a finished proof — it is a research program with clear milestones, tractable short-term tests, and deep long-term problems.

References

[A] "Mass Gap Addendum Research" (2025). [Topological framework: Atiyah–Singer index theorem, Källén–Lehmann spectral representation, Chern–Weil gauging of P/CP, Ashtekar gravity, AdS/CFT derivation of Δ_{grav} , CEH instanton bundling.]

[B] "A constructive proof of existence and mass gap for pure SU(3) Yang–Mills in four-dimensional space-time," arXiv:2506.00284 (2025). [Constructive framework: 5D gauge theory with torsion, polymer expansions, Sturm–Liouville analysis, lattice QCD, experimental proposals.]

1. Abbott, B.P. et al. (LIGO/Virgo) (2018). "GW170817: measurements of neutron star radii and equation of state." *Phys. Rev. Lett.* 121, 161101. arXiv:1805.11579.
2. Abel, C. et al. (nEDM collaboration) (2020). "Measurement of the permanent electric dipole moment of the neutron." *Phys. Rev. Lett.* 124, 081803. arXiv:2008.10635.
3. Agazie, G. et al. (NANOGrav) (2023). "The NANOGrav 15-year data set: evidence for a gravitational-wave background." *Astrophys. J. Lett.* 951, L8. arXiv:2306.16213.
4. Almheiri, A., Engelhardt, N., Marolf, D. & Maxfield, H. (2019). "Entanglement entropy and islands in quantum gravity." *JHEP* 1912, 105. arXiv:1905.08762.
5. Amari, S. (1985). *Differential-geometrical methods in statistics*. Springer.
6. Amari, S. (2016). *Information geometry and its applications*. Springer.
7. Ambjørn, J., Görlich, A., Jurkiewicz, J. & Loll, R. (2012). "Nonperturbative quantum gravity."

- Phys. Rep.* 519, 127. arXiv:1203.3591.
8. Ambjørn, J., Görlich, A., Jurkiewicz, J. & Loll, R. (2013). "The Lorentzian CDT phase diagram." *J. Phys. Conf. Ser.* 453, 012013.
 9. Ambjørn, J., Jurkiewicz, J. & Loll, R. (1998). "Dynamically triangulating Lorentzian quantum gravity." *Nucl. Phys. B* 510, 391. arXiv:hep-th/9806198.
 10. Ambjørn, J., Jurkiewicz, J. & Loll, R. (2000). "A non-perturbative Lorentzian path integral for gravity." *Phys. Rev. Lett.* 85, 924. arXiv:hep-th/0002050.
 11. Anninos, D., Hartman, T. & Strominger, A. (2011). "Higher spin realization of the dS/CFT correspondence." *Class. Quant. Grav.* 28, 015009. arXiv:1008.5735.
 12. Argyres, P.C., Plesser, R. & Ünsal, M. (2018). "Survey of resurgence and large-N." *J. Phys. Conf. Ser.* 965, 012003.
 13. Arzoumanian, Z. et al. (NANOGrav) (2020). "The NANOGrav 12.5-year data set: search for an isotropic stochastic gravitational-wave background." *Astrophys. J. Lett.* 905, L34. arXiv:2009.04496.
 14. Ashtekar, A. (1986). "New variables for classical and quantum gravity." *Phys. Rev. Lett.* 57, 2244.
 15. Ashtekar, A. (1987). "New Hamiltonian formulation of general relativity." *Phys. Rev. D* 36, 1587.
 16. Ashtekar, A. & Lewandowski, J. (1997). "Quantum theory of geometry I: Area operators." *Class. Quant. Grav.* 14, A55. arXiv:gr-qc/9602046.
 17. Atiyah, M.F. (1988). "New invariants of three and four dimensional manifolds." In *The Mathematical Heritage of Hermann Weyl*, Proc. Symp. Pure Math. 48, 285.
 18. Atiyah, M.F. & Bott, R. (1982). "The Yang–Mills equations over Riemann surfaces." *Phil. Trans. R. Soc. Lond. A* 308, 523.
 19. Atiyah, M.F. & Singer, I.M. (1968). "The index of elliptic operators I." *Ann. Math.* 87, 484.
 20. Atiyah, M.F., Patodi, V.K. & Singer, I.M. (1975). "Spectral asymmetry and Riemannian geometry I." *Math. Proc. Camb. Phil. Soc.* 77, 43.
 21. Banks, T. & Casher, A. (1980). "Chiral symmetry breaking in confining theories." *Nucl. Phys. B* 169, 103.
 22. Behtash, A., Catterall, S., Sulejmanpasic, A. & Ünsal, M. (2015). "Gauge theory in the thimble formulation." *JHEP* 1511, 186. arXiv:1508.07150.
 23. Bianchini, F. & Silk, J. (2020s). "CMB spectral distortions from the QCD epoch." (References on CMB μ -distortion from QCD.)
 24. Bisognano, J. & Wichmann, E. (1975). "On the duality condition for a Hermitian scalar field." *J. Math. Phys.* 16, 985.
 25. Bisognano, J. & Wichmann, E. (1976). "On the duality condition for quantum fields." *J. Math. Phys.* 17, 303.
 26. Bogomolny, E.B. & Fateev, V.A. (1977). "Large-order calculations in quantum field theory." *Phys. Lett. B* 71, 93.
 27. Braunstein, S.L. & Caves, C.M. (1994). "Statistical distance and the geometry of quantum

- states." *Phys. Rev. Lett.* 72, 3439.
28. Burrage, C., Csaki, C. & Shirai, S. (2024). "Theoretical priors for the QCD θ -angle."
 29. Calabrese, P. & Cardy, J. (2004). "Entanglement entropy and quantum field theory." *J. Stat. Mech.* P06002.
 30. Caprini, C. et al. (2020). "Science with the space-based interferometer LISA II: gravitational waves from cosmological phase transitions." *JCAP* 2004, 001. arXiv:1910.13125.
 31. Caron-Huot, S. & Mizera, S. (2020). "Crossing symmetry in the S-matrix bootstrap." *JHEP* 2006, 037.
 32. Casini, H. & Huerta, M. (2009). "Entanglement entropy for the S^n -sphere." *Phys. Lett. B* 694, 167. arXiv:1007.1813.
 33. Chandia, O. & Zanelli, J. (1997). "Supersymmetric particle in a spacetime with torsion and the index theorem." *Phys. Rev. D* 58, 045014. arXiv:hep-th/9803034.
 34. Cherman, A., Dorigoni, D., Ünsal, M. & D'Onofrio, A. (2014). "Resurgence in quantum field theory." *JHEP* 1410, 066.
 35. Chern, S.S. & Simons, J. (1974). "Characteristic forms and geometric invariants." *Ann. Math.* 99, 48.
 36. Chew, G.F. (1966). *The analytic S-matrix*. W.A. Benjamin.
 37. Coumbe, D.N. & Loll, R. (2018). "Spectral dimension flow in CDT."
 38. Cristoforetti, M., Di Cortona, F. & Scorzato, L. (2013). "On the Abelian definition and use of Lefschetz thimbles." *PoS LATTICE2013*, 019. arXiv:1311.6792.
 39. Damgaard, P.H. & Fukaya, H. (2011). "The microscopic spectrum of the QCD Dirac operator." *JHEP* 1101, 058. arXiv:1010.5638.
 40. Damgaard, P.H., Fukaya, H. & Risch, T. (2010). "Lattice QCD and chiral random matrix theory." *JHEP* 1010, 093.
 41. di Cortona, G.G., Hardy, E., Park, J.W. & Tedeschi, B. (2024). "The QCD axion: a review of the physics and phenomenology." *Ann. Rev. Nucl. Part. Sci.*
 42. Donaldson, S.K. (1983). "An application of gauge theory to four-dimensional topology." *J. Diff. Geom.* 18, 279.
 43. Donnachie, A. & Landshoff, P.V. (1992). "Total cross-sections." *Phys. Lett. B* 296, 227.
 44. Dunne, G.V. & Ünsal, M. (2014). "Resurgence and trans-series in quantum field theory." *JHEP* 1412, 102. arXiv:1410.2574.
 45. Écalle, J. (1981). *Les fonctions resurgentes*. Publ. Math. Orsay.
 46. Eguchi, T. & Kawai, H. (1982). "Reduction of dynamical degrees of freedom in the large- N gauge theory." *Phys. Rev. Lett.* 48, 1063.
 47. Essick, R. et al. (2023). "Nonparametric EOS inference with NICER and LIGO."
 48. Fixsen, D.J. et al. (FIRAS) (1996). "The cosmic microwave background spectrum from COBE." *Astrophys. J.* 473, 576.
 49. Floer, A. (1988). "An instanton invariant for 3-manifolds." *Comm. Math. Phys.* 118, 215.
 50. Gelman, A. et al. (2013). *Bayesian data analysis*. 3rd ed., CRC Press.

51. Gonzalez-Arroyo, A. & Okawa, M. (2010). "Twisted reduction model and large- N QCD." *PoS LATTICE2010*, 308.
52. Gribov, V.N. (1961). "Partial waves with complex angular momentum." *Soviet Phys. JETP* 14, 1395.
53. Gutzwiller, M.C. (1971). "Periodic orbits and classical quantization conditions." *J. Math. Phys.* 12, 343.
54. Hamada, H. (2010s). "Regge calculus for gravitational instantons."
55. Heidenreich, B. et al. (2021). "Chern–Weil global symmetries and how quantum gravity avoids them." *JHEP* 11, 053. arXiv:2012.00009.
56. Hogan, C.J. (1986). "Gravitational radiation from cosmological phase transitions." *Mon. Not. R. Astron. Soc.* 218, 629.
57. Hook, A. (2019). "TASI lectures on the strong CP problem." arXiv:1812.02600.
58. Janssen, G.H. et al. (2015). "Gravitational waves with the SKA: pulsar timing arrays." *PoS AASKA14*, 037. arXiv:1501.00127.
59. Jaynes, E.T. (2003). *Probability theory: The logic of science*. Cambridge University Press.
60. Jeffreys, H. (1946). "An invariant form for the prior probability in estimation problems." *Proc. Roy. Soc. A* 186, 453.
61. Källén, G. & Lehmann, H. (1950). "On the matrix elements of the scattering operator." *Dan. Mat. Fys. Medd.* 27, 22.
62. Kass, R.E. & Raftery, A.E. (1995). "Bayes factors." *J. Amer. Stat. Assoc.* 90, 773.
63. Kogut, A. et al. (PIXIE) (2011). "The Primordial Inflation Explorer (PIXIE)." *J. Phys. Conf. Ser.* 283, 012004.
64. Kovner, A. & Rosenstein, B. (1989). "The Eguchi–Kawai reduction and dynamical quark loops." *Phys. Lett. B* 232, 443.
65. Kreimer, D. & Mielke, E.W. (2001). "Topological invariants, instantons, and chiral anomaly on spaces with torsion." *Phys. Rev. D* 63, 045014. arXiv:hep-th/9702025.
66. Li, H. & Haldane, F.D.M. (2008). "Entanglement spectrum as a generalization of entanglement entropy." *Phys. Rev. Lett.* 101, 010504. arXiv:0805.0332.
67. Liu, C. et al. (2016). "Fidelity susceptibility and the spectral function." *Chinese Phys. B* 25, 030501.
68. Loll, R. (2019). "Quantum gravity from CDT." *Proceedings of Science*. arXiv:1905.08028.
69. Maldacena, J. (2003). "Non-gaussian features of primordial fluctuations in single field inflationary models." *JHEP* 0305, 013.
70. Mariño, M. (2012). *Instantons and large N : An introduction to non-perturbative methods in QFT*. Cambridge University Press.
71. Matsumoto, K. (2002). "On multi-parameter quantum estimation theory." *J. Phys. A* 35, 3119.
72. Mielke, E.W. (2006). "Axions in gravity with torsion." arXiv:1502.03694.
73. Mielke, E.W. & Kreimer, D. (1997). "Chiral anomaly in contorted spacetimes." arXiv:gr-qc/9704051.

- 74.n2EDM collaboration (2021). "A new neutron EDM experiment at PSI." *J. Phys. Conf. Ser.* 1790, 012016.
- 75.Nieh, H.T. & Yan, M.L. (1982). "An identity in Riemann–Cartan geometry." *J. Math. Phys.* 23, 373.
- 76.Penington, G. (2020). "Entanglement wedge reconstruction and the error correction property of the AdS/CFT correspondence." *JHEP* 2009, 002. arXiv:1912.06582.
- 77.Petz, D. (2002). "Monotone metrics on matrix spaces." *Linear Algebra Appl.* 405, 263.
- 78.Provost, J.P. & Valleaux, G. (1980). "Riemannian structure on manifolds of quantum states." *J. Math. Phys.* 21, 2437.
- 79.Rao, C.R. (1945). "Information and the accuracy attainable in the estimation of statistical parameters." *Bull. Calcutta Math. Soc.* 37, 81.
- 80.Ray, D.B. & Singer, I.M. (1971). "R-torsion and the Laplacian on odd-dimensional manifolds." *Adv. Math.* 7, 145.
- 81.Regge, T. (1961). "General relativity without coordinates." *Nuovo Cim.* 19, 558.
- 82.Riley, T.E. et al. (NICER) (2019). "A NICER view of PSR J0030+0451." *Astrophys. J. Lett.* 887, L21. arXiv:1912.05702.
- 83.Rovelli, C. (2004). *Quantum gravity*. Cambridge University Press.
- 84.Rovelli, C. & Smolin, L. (1995). "Discreteness of area and volume in quantum gravity." *Nucl. Phys. B* 442, 593. arXiv:gr-qc/9411005.
- 85.Selberg, A. (1956). "Harmonic analysis and discontinuous groups in weakly symmetric Riemannian spaces." *J. Indian Math. Soc.* 20, 47.
- 86.Skilling, J. (2004). "Nested sampling." *AIP Conf. Proc.* 735, 395.
- 87.Strominger, A. (2001). "The dS/CFT correspondence." *JHEP* 0110, 034. arXiv:hep-th/0106113.
- 88.Sunyaev, R.A. & Zeldovich, Y.B. (1970). "Small-scale fluctuations of relic radiation." *Ap. Space Sci.* 7, 20.
- 89.Tanizaki, Y. (2018). "Lefschetz-thimble techniques for path integrals with singularities." *PoS LATTICE2018*, 026. arXiv:1901.04303.
- 90.Tanizaki, Y. & Ünsal, M. (2014). "Modified instanton sum in QCD and higher-order terms." *JHEP* 1407, 054. arXiv:1404.7454.
- 91.Tanizaki, Y. & Ünsal, M. (2014). "Topological insulators and the QCD vacuum: the θ parameter as a Berry phase." *Phys. Rev. D* 89, 125011. arXiv:1401.7399.
- 92.Trotta, R. (2008). "Bayes in the sky: Bayesian inference and model selection in cosmology." *Contemp. Phys.* 49, 71. arXiv:0803.4089.
- 93.Uhlenbeck, K. (1982). "Connections with L^p bounds on curvature." *Comm. Math. Phys.* 83, 31.
- 94.Ünsal, M. (2008). "Abelian duality, confinement, and chiral symmetry breaking." *Phys. Rev. Lett.* 100, 032005. arXiv:0708.2252.
- 95.Verbaarschot, J. & Wettig, T. (2000). "Random matrix theory and chiral symmetry in QCD." *Ann. Rev. Nucl. Part. Sci.* 50, 343.

- 96. Witten, E. (1988). "Topological quantum field theory." *Comm. Math. Phys.* 117, 353.
- 97. Witten, E. (2004). "Cosmic separation of phases." *Phys. Rev. D* 30, 272.
- 98. Witten, E. (2010). "Analytic continuation of Chern–Simons theory." In *Chern–Simons Gauge Theory 20 Years Later*. arXiv:1001.2933.
- 99. You, W.L. & Li, Y.Q. (2007). "Fidelity susceptibility and the quantum metric." *Phys. Rev. B* 75, 045121.
- 100. Zanardi, P., Giorda, P. & Cozzini, M. (2007). "Information-theoretic differential geometry of quantum phase transitions." *Phys. Rev. Lett.* 99, 100603.
- 101. Zinn-Justin, J. (2004). *Quantum field theory and critical phenomena*. 4th ed., Oxford University Press.
- 102. arXiv:1312.1350 (2013). "Yang–Mills mass gap, Floer homology, glueball spectrum, and conformal window in large- N QCD."
- 103. arXiv:1909.08775 (2019). "Anomaly inflow and the η -invariant."
- 104. arXiv:2409.06766 (2024). "The chiral torsional anomaly and the Nieh–Yan invariant with and without boundaries." *JHEP* 12, 149.
- 105. arXiv:2506.00078 (2025). "The role of Berry phases in the QCD vacuum structure."