Elegant Black Hole Physics from Informational Tunneling (q→τ)

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 17, 2025

Across nature, coherence emerges when a realized model q aligns with the generative structure τ. We apply this cross-scale motif to black holes via a single principle: Informational Alignment. Physical evolution is a constrained gradient flow that monotonically reduces the divergence $D(q \parallel \tau)$ while a convex informational curvature regularizer $R(K)$caps curvature invariants K. In the weak-curvature regime, the theory reproduces GR + QFT; near extreme curvature, $R(K)$ halts geodesic focusing, replacing singularities with smooth cores and taming Kerr's inner-horizon instabilities—without firewalls or ad hoc matter. The same principle underwrites unitary evaporation: generalized-entropy extremality (quantum extremal surfaces and islands) arises as Euler–Lagrange conditions of the information-regularized action, so information leaves in outward correlations after the Page time while the horizon remains smooth for infallers. Interior geometry remains "it from bit": entanglement and complexity generate an expanding interior volume that is later encoded in the radiation's entanglement wedge. The result is a single, covariant upgrade that keeps all tested predictions of GR, renders black holes globally normal (no geodesic incompleteness, no Cauchy-horizon pathology), and yields crisp observational signatures in ringdown, high-spin tails, and shadows.

Keywords: informational tunneling, informational alignment, curvature cap, black holes, Kerr instability, singularity resolution, generalized entropy, islands, unitarity, entanglement, complexity, Logos, cosmology, cognition.

Outline

1. Premise: Alignment as Physics
   1.1 Coherence as $q \to \tau$alignment
   1.2 Divergence $D(q \parallel \tau)$as a physical Lyapunov functional
2. Principle and Action
   2.1 Informational Alignment Principle (IAP): $\dot{D} \leq 0$under conservation
   2.2 Information-regularized action: $S = \int \square\sqrt{-g}\,[\frac{1}{16\pi G}R + \mathcal{L}_m] + \int \square\sqrt{-g}\,R(K)$
   2.3 Choosing $R(K)$: convexity, bounded invariants, GR limit
3. Field Equations and Limits
   3.1 Modified Einstein equations and Bianchi consistency
   3.2 Recovery of GR + QFT for $K \ll K_*$
4. Static, Spherical Interiors
   4.1 Core radius $r_0 \sim (GM)^{1/3} L_{\backslash *}^{2/3}$and finite invariants
   4.2 TOV with effective focusing cap; absence of singularity
5. Rotating (Kerr) Black Holes
   5.1 Mass inflation as informational over-focusing
   5.2 IAP-induced saturation and restoration of global hyperbolicity
6. Thermodynamics and Information
   6.1 Generalized entropy as stationarity condition; islands from variation
   6.2 Smooth horizons, unitary evaporation, Page curve
7. Geometry from Entanglement
   7.1 Interior volume/complexity growth under $\dot{D} \leq 0$
   7.2 Entanglement-wedge reconstruction and τ-aligned readout
8. Observational Signatures
   8.1 Ringdown shifts $\sim \mathcal{O}((r_0/r_s)^2)$
   8.2 High-spin tail modifications; sub-extremal spin ceiling
   8.3 Shadow micro-corrections and constraints on K_*
9. Beyond Black Holes
   9.1 Stars/galaxies as $q \approx \tau$attractors
   9.2 Divergent minds vs. informational singularities

1. Premise: Alignment as Physics

1.1 Coherence as q→τ alignment

- Objects:
  - q: the realized state/model of a system (fields + geometry + control parameters).
  - τ: the generative structure that “makes” the data (dynamics, constraints, symmetries, boundary conditions).
- Claim (Alignment Postulate): Physical evolution tends to reduce the *mismatch* between $q$and τ; coherence is the regime where $q \approx \tau$.
- Operational reading across domains:
  - Thermo/transport: near-equilibrium constitutive laws are the $q \approx \tau$ limit of the microdynamics.
  - GR: a spacetime is coherent when its stress–energy and curvature satisfy the same constraints (Bianchi identities, conservation)—i.e., the realized metric $q$matches the constraint-generating structure τ.
  - Computation/cognition: learning minimizes predictive error so that the model $q$recovers the data-generating process τ.
- Geometric picture: Put q on a statistical/quantum state manifold with metric given by Fisher information/quantum Fisher. Alignment is geodesic approach of $q$toward the constraint-set defined by τ.
- Phenomenology: “Radiance” (useful work, low dissipation, clean signals) appears when $q \to \tau$; turbulence, decoherence, and runaway focusing mark $q \to \tau$.

1.2 Divergence D(q‖τ) as a physical Lyapunov functional

- Definition (choice of D):
  - Classical: $D_{\mathrm{KL}}(q \parallel \tau) = \sum_i \quad q_i \ln{(q_i/\tau_i)}$or $\int q \ln{(q/\tau)}\, dx$.
  - Quantum (algebraic): $D(\rho \parallel \sigma) = \mathrm{Tr}[\rho(\ln \rho - \ln \sigma)]$, with ρ the realized state and σ the τ-consistent reference (e.g., maximum-entropy state under τ’s constraints).
- Lyapunov property: along physical evolution,

$$\frac{d}{dt} D(q \parallel \tau) \ \le\ 0,$$

with equality iff q lies on the τ-manifold (stationarity). This is guaranteed when dynamics are contractive under coarse-graining (data-processing inequality) or implement gradient flow of D under an appropriate metric (Jordan–Kinderlehrer–Otto/Wasserstein or quantum analogs).

- Structure:
  - Generator: $\dot{q} = -\mathcal{G}^{-1}(q)\,\nabla_q D(q \parallel \tau)$(steepest descent), where $\mathcal{G}$ is a positive operator induced by dissipation/transport (classical) or modular flow (quantum).
  - Stationary set: $\nabla_q D = 0 \iff q \in \mathcal{M}_\tau$ (the constraint surface defined by τ).
  - Local convexity: near $\mathcal{M}_\tau$, D is second-order positive, fixing linear response and fluctuation–dissipation relations.
- GR tie-in (semiclassical): choose D over observables that include geometry (via the modular Hamiltonian for wedges/surfaces). Then extremizing generalized entropy $S_{\mathrm{gen}} = \frac{A}{4G} + D_{\mathrm{out}}$ is equivalent to stationarity of D under geometric variations—identifying quantum extremal surfaces as Euler–Lagrange conditions of the same Lyapunov functional.
- Black-hole interior intuition: unchecked focusing corresponds to increasing D (model decouples from constraints). Adding an informational-curvature regularizer forces the flow back into $dD/dt \le 0$, preventing singularities while retaining GR when curvature is small.
- Diagnostics & measurables:
  - Near-alignment: variance–response matches Fisher predictions; ringdown follows minimal Dtrajectories.
  - Misalignment: rising D flags inner-horizon instabilities; saturation by the curvature regularizer restores $dD/dt \le 0$.

Takeaway: Treat $D(q \parallel \tau)$as the energy-like, covariant Lyapunov that physics already (often implicitly) minimizes. Coherence is simply the fixed point $q \to \tau$; black-hole pathologies are regimes where the unregularized flow lets D blow up, and our "elegant fix" is to guarantee monotone descent with a single curvature-aware term.

2. Principle and Action

2.1 Informational Alignment Principle (IAP): $D \le 0$ under conservation

- Objects.
  q = realized model (state + geometry).
  τ = generative structure (constraints, symmetries, boundary data).
  $D(q\|\tau)$ = divergence measuring misalignment (classical KL, quantum relative entropy, or a coarse-grained surrogate tied to accessible observables).

- Postulate (single arrow of physical updating).
  Along physical evolution,

- $d/dt\, D(q\|\tau) \le 0$

with equality iff q lies on the τ–manifold (stationarity). The flow is generated as a constrained steepest descent:

$\dot{q} = -\, G^{-1}(q)\, \nabla_q D(q\|\tau)$ subject to (i) $\nabla_\mu T^{\mu\nu}=0$, (ii) Bianchi: $\nabla_\mu G^{\mu\nu}=0$,

where G is a positive (problem-dependent) Onsager/metric operator. The conservation constraints guarantee that alignment never "cheats" by violating local energy–momentum balance or diffeomorphism invariance.

- Physical reading.
  In weak curvature, the IAP reduces to familiar relaxation (hydro, kinetic, linear response).
  In strong curvature, the same monotonicity must be preserved; this is where the curvature regularizer enters (next subsection).

2.2 Information-regularized action:

$S = \int\sqrt{-g}\, [\, (1/16\pi G)\, R + L_m\,] + \int\sqrt{-g}\, \mathcal{R}(K)$

- Purpose.
  Add one covariant, observer-independent term $\mathcal{R}(K)$ that prevents runaway geodesic focusing by penalizing large curvature without altering tested low-curvature physics.

- Notation (to avoid collisions).
  R = Ricci scalar (as usual).
  $\mathcal{R}(K)$ = *regulator* (a function of an invariant K).
  K may be chosen as, e.g., the Kretschmann scalar $K \equiv R_{\alpha\beta\gamma\delta} R^{\alpha\beta\gamma\delta}$. Other choices (linear combos of R^2, $R_{\mu\nu}R^{\mu\nu}$, Weyl^2) are possible; use one invariant knob to keep the principle minimal.

- Field equations (schematic).
  Varying S with respect to $g^{\mu\nu}$ gives
- $G_{\mu\nu} = 8\pi G \, (T_{\mu\nu} + T^{(reg)}_{\mu\nu})$,

where

$$T^{(reg)}_{\mu\nu} \equiv -(2/\sqrt{-g})\, \delta/\delta g^{\mu\nu} \int \sqrt{-g}\, \mathcal{R}(K)$$

$$= -1/2\, g_{\mu\nu}\, \mathcal{R}(K) + 2\, \mathcal{R}'(K) \bullet Q_{\mu\nu}[g],$$

with $Q_{\mu\nu}[g]$ a covariant, divergence-free tensor built from curvature (for $K = R_{\alpha\beta\gamma\delta}R^{\alpha\beta\gamma\delta}$, $Q_{\mu\nu}$ contains terms like $R_{\mu\alpha\beta\gamma}\, R_{\nu}{}^{\alpha\beta\gamma}$ and covariant derivatives of curvature). Because S is diffeo-invariant,

$$\nabla^\mu T^{(reg)}_{\mu\nu} = 0 \quad \text{and} \quad \nabla^\mu (G_{\mu\nu} - 8\pi G\, T^{(tot)}_{\mu\nu}) = 0$$

hold identically; the Bianchi identity is preserved.

- IAP as a variational statement.
  The dynamics implied by S ensure that, for admissible coarse-grainings, the physical evolution implements a contraction of $D(q\|\tau)$. At the semiclassical level, extremizing generalized entropy $S_{gen} = A/4G + D_{out}$ reproduces QES/islands as Euler–Lagrange conditions of the same principle, so unitarity and a smooth horizon coexist.

2.3 Choosing $\mathcal{R}(K)$: convexity, bounded invariants, GR limit

To stay elegant and predictive, $\mathcal{R}(K)$ must satisfy:

- (i) Convexity & monotone penalty.
- $\mathcal{R}'(K) \geq 0, \quad \mathcal{R}''(K) \geq 0 \quad$ for $K \geq 0$,

so large curvature always costs more and the penalty is well-posed.

- (ii) Low-curvature transparency (GR limit).
- $\mathcal{R}(0) = 0, \quad \mathcal{R}'(0) = 0, \quad$ and $\mathcal{R}(K) \approx O(K^2/K_*)$ as $K \ll K_*$,

ensuring post-Newtonian tests, gravitational waves, lensing, etc., are unchanged to current precision.

- (iii) Effective curvature cap near K_*.
  The regulator must *saturate focusing* before geodesic incompleteness. Two minimalist families:

A. Born–Infeld–type soft cap

$\mathcal{R}$_BI(K) = (K_* / λ) [√(1 + λ K / K_*) − 1],

with 0 < λ ≤ 1. Then $\mathcal{R}$'(K) = (1/2) / √(1 + λ K/K_*), which diminishes the incremental "gain" from further curvature, self-limiting mass-inflation and central focusing.

B. Inequality via auxiliary field (hard cap, variationally clean)
Add a Lagrange multiplier Λ(x) enforcing K ≤ K_*:

S = ∫√−g [(1/16πG) R + L_m + Λ(x) (K_* − K)] ,

with KKT conditions: Λ ≥ 0, K ≤ K_*, Λ (K_* − K) = 0.
In weak curvature, Λ = 0 and GR is recovered; when evolution tries to push K upward, Λ turns on and halts it *covariantly*, guaranteeing no singularity.

- (iv) Stability/causality.
  Choose $\mathcal{R}$ so the linearized spectrum around GR has no tachyons/ghosts in the low-K regime (e.g., keep corrections beyond current detectors' reach; or use the auxiliary-field form, which avoids higher time-derivatives in the principal part of the equations).
- (v) One scale, universal.
  K_* sets the onset of regularization (Planckian by default). With K_* fixed, the interior core scale in Schwarzschild follows
- r_0 ~ (G M)^{1/3} L_*^{2/3}, with L_* ≡ K_*^{−1/4},

giving r_0 << r_s for astrophysical M (so exteriors remain Kerr/Schwarzschild to excellent approximation).

- (vi) Observable consequences (small, crisp, testable).
  Ringdown shifts ~ O((r_0/r_s)^2), suppressed inner-horizon turbulence at high spin, slightly sub-extremal spin ceiling, and at most percent-level shadow corrections unless L_* >> L_P.

Net effect.
With a single convex $\mathcal{R}$(K) added to the Einstein–Hilbert action, the IAP's .D ≤ 0 becomes globally enforceable: horizons stay smooth, Kerr's Cauchy horizon cannot trigger a pathology, and central singularities are replaced by finite-curvature cores—all while standard GR is recovered wherever K << K_*.

3. Field Equations and Limits

3.1 Modified Einstein equations and Bianchi consistency

Start from the information-regularized action

$S = \int\sqrt{-g}\, [\, (1/16\pi G) R + L_m \,] + \int\sqrt{-g}\, \mathcal{R}(K)$,

where K is a scalar curvature invariant (e.g., $K = R_{\alpha\beta\gamma\delta} R^{\alpha\beta\gamma\delta}$) and $\mathcal{R}(K)$ is the convex regulator (Sec. 2.3).

Varying with respect to $g^{\mu\nu}$ gives

$G_{\mu\nu} = 8\pi G\, (T_{\mu\nu} + T^{(reg)}_{\mu\nu})$,

with ordinary matter $T_{\mu\nu} \equiv -(2/\sqrt{-g})\, \delta(\sqrt{-g}\, L_m)/\delta g^{\mu\nu}$ and

$T^{(reg)}_{\mu\nu} = -\, 1/2\, g_{\mu\nu}\, \mathcal{R}(K) + 2\, \mathcal{R}'(K)\, Q_{\mu\nu}[g] + 2\, (\nabla\nabla \text{ terms})$,

where $Q_{\mu\nu}[g]$ collects algebraic curvature contributions (e.g., $R_{\mu\alpha\beta\gamma} R_{\nu}{}^{\alpha\beta\gamma}$ for K = Riem^2), and the "($\nabla\nabla$ terms)" are covariant derivatives of $\mathcal{R}'(K)$ times curvature (their explicit form is standard for f(K) theories).

Because S is diffeomorphism-invariant, the total RHS is covariantly conserved:

$\nabla^\mu T_{\mu\nu} = 0, \qquad \nabla^\mu T^{(reg)}_{\mu\nu} = 0$,

and the Bianchi identity $\nabla^\mu G_{\mu\nu} = 0$ is automatically satisfied. Hence the Bianchi consistency (local energy-momentum conservation and diffeo invariance) is preserved exactly.

Two practical realizations:

- Soft cap (Born–Infeld–type $\mathcal{R}$): higher-derivative terms appear but are suppressed by K/K_*; choose coefficients so the linearized spectrum has no ghosts/tachyons in the low-K regime.
- Hard cap (auxiliary Λ field): enforce $K \le K_*$ with $S \supset \int\sqrt{-g}\, \Lambda(K_* - K)$. *The Euler–Lagrange/KKT conditions keep $\Lambda = 0$ when $K \ll K_*$*, turning on only near the threshold. This form avoids adding higher-time-derivative principal parts, making hyperbolicity and stability easier to control.

Either way, the regulator contributes an effective stress–energy that counters runaway focusing while remaining covariant and conserved.

3.2 Recovery of GR + QFT for $K \ll K_*$

To ensure the ordinary universe is unchanged where curvature is small, require the regulator be flat to first order:

$\mathcal{R}(0) = 0$, $\mathcal{R}'(0) = 0$, $\mathcal{R}(K) = O(K^2 / K_*)$ as $K \to 0$.

Consequences:

- Field equations: For $K \ll K_*$, $T^{(reg)}_{\mu\nu} = O(K/K_*) \times$ (curvature) is parametrically tiny. The equations reduce to
- $G_{\mu\nu} = 8\pi G\, T_{\mu\nu}$ + small corrections ($\propto K/K_*$),

i.e., standard GR to any desired post-Newtonian and gravitational-wave accuracy.

- Linearized limit: Expanding around Minkowski or a weakly curved background, the graviton propagator and gauge structure are unchanged at leading order; regulator effects enter at higher order in $(\text{length})^{-4} / K_*$ and are invisible to current tests if K_* is Planckian.
- Semiclassical QFT: In curved spacetime with $K \ll K_*$, *the usual renormalized $\langle T_{\mu\nu} \rangle$ and Hawking/Unruh predictions are recovered. Near astrophysical horizons, K scales like $O(M^{-4})$ (up to order-unity factors), so for stellar to supermassive masses* $K/K_* \ll 1$ unless K_* is chosen anomalously low. Hence:
    - No horizon drama: local EFT is valid for infallers.
    - Thermodynamics intact: Bekenstein–Hawking area law and leading Hawking flux remain, with regulator corrections suppressed by K/K_*.
- Strong-curvature interior only: The regulator becomes relevant deep inside, where GR would otherwise drive $K \to \infty$ (central or inner-horizon focusing). There it supplies the minimal, covariant backreaction needed to keep $K \le K_*$ and prevent geodesic incompleteness—without spoiling the exterior.

Summary: The modified field equations remain Bianchi-consistent and reduce exactly to GR + (ordinary) QFT wherever $K \ll K_*$. Only in the extreme-curvature interior—precisely where GR breaks—does the regulator activate, delivering a nonsingular, globally hyperbolic completion while leaving all tested regimes untouched.

4. Static, Spherical Interiors

4.1 Core radius $r_0 \sim (G M)^{1/3} L_*^{2/3}$ and finite invariants

Setup ($c = 1$ units). Outside the core the spacetime is well-approximated by Schwarzschild. The leading curvature scalar that signals blow-up in GR is the Kretschmann

$$K_{Schw}(r) = R_{abcd} R^{abcd} = 48 (G M)^2 / r^6 .$$

Impose the single invariant bound $K \le K_*$ from the regulator. The would-be GR growth saturates where $K_{Schw}(r_0) \approx K_*$:

$$48 (G M)^2 / r_0^6 \approx K_* \;\Rightarrow\; r_0 \approx (48)^{1/6} (G M)^{1/3} K_*^{-1/6}.$$

Defining $L_* \equiv K_*^{-1/4}$ (the curvature length), this is

$$r_0 \approx (48)^{1/6} (G M)^{1/3} L_*^{2/3} \simeq 1.4 (G M)^{1/3} L_*^{2/3}.$$

Thus $r_0 \ll r_s$ for astrophysical M, so exterior tests of GR are untouched.

Regular core (finite invariants). Inside $r \lesssim r_0$ the regulator contributes an effective stress-energy $T^{(reg)}_{\mu\nu}$ that halts focusing. The most economical static outcome is a de Sitter–like core:

$$ds^2 \approx -(1 - (\Lambda_{eff}/3) r^2) dt^2 + (1 - (\Lambda_{eff}/3) r^2)^{-1} dr^2 + r^2 d\Omega^2,$$

with invariants

$$R = 4 \Lambda_{eff}, \qquad K_{core} = R_{abcd} R^{abcd} = (8/3) \Lambda_{eff}^2.$$

Choosing $\Lambda_{eff} = O(1/L\_{}^2)$ *makes* $K_{core} = O(1/L\_{}^4) \le K_*$ by construction. All scalars $\{R, R_{ab}R^{ab}, K\}$ remain finite; geodesics extend through $r = 0$ with the regularity conditions

$$m(r) \sim (4\pi/3) \rho_{eff}(0) r^3, \quad g_{tt}(0) \text{ finite}, \quad g_{rr}(0) \text{ finite}, \quad P'(0) = 0.$$

Matching at $r \approx r_0$ to the exterior Schwarzschild solution via the usual Israel conditions yields a thin (or smooth) transition layer whose width is set by the stiffness of $\mathcal{R}(K)$.

4.2 TOV with effective focusing cap; absence of singularity

Ordinary TOV (for reference).
For a static, spherical fluid,

$$dm/dr = 4\pi r^2 \rho,$$

$$dP/dr = - (\rho + P) (m + 4\pi r^3 P) / (r (r - 2 G m)).$$

In GR, reasonable equations of state allow collapse to drive $(\rho + 3P)$ large, focusing congruences and pushing $K \to \infty$.

Effective focusing cap (from the regulator).
The regulator adds a conserved $T^{(reg)}_{\mu\nu}$ that, in the static, spherical sector, is equivalent to an effective (ρ_{eff}, P_{eff}) entering the structure equations:

$$dm/dr = 4\pi r^2 (\rho + \rho_{reg}),$$

$$dP/dr = - ((\rho + \rho_{reg}) + (P + P_{reg}))$$

$$(m + 4\pi r^3 (P + P_{reg})) / (r (r - 2 G m)).$$

Crucially, near the high-curvature region the regulator enforces a cap on the focusing combination

$$(\rho + 3P)_{eff} \equiv (\rho + \rho_{reg}) + 3 (P + P_{reg}) \to \text{bounded as } K \to K_*,$$

so the Raychaudhuri term that would drive geodesic convergence cannot diverge. A minimal, covariant realization (consistent with Sec. 2–3) is:

- Soft cap: choose $\mathcal{R}(K)$ convex with $\mathcal{R}'(K)$ decreasing the incremental response at large K; this produces $P_{reg} < 0$ of de Sitter type when $K \gtrsim K_*$.
- Hard cap: enforce $K \le K_*$ via a multiplier $\Lambda(x)$; when K tries to grow, $\Lambda > 0$ sources an effective vacuum-like stress that exactly arrests further focusing.

Regular-center TOV behavior.
With the cap active for $r \lesssim r_0$, the unique regular solutions satisfy

$$m(r) = (4\pi/3) \rho_{eff}(0) r^3 + O(r^5),$$

$$P(r) = P(0) - (2\pi/3) (\rho_{eff}(0) + P_{eff}(0)) (\rho_{eff}(0) + 3 P_{eff}(0)) r^2 + O(r^4),$$

and the metric functions admit even Taylor series in r. There is no 1/r or ln r behavior, hence no curvature blow-up. The interior smoothly approximates a constant-density, positive-pressure (vacuum-like) core with

$$K(r) \to K_{core} = O(1/L_*^4) \text{ (finite).}$$

Absence of singularity (covariant statement).
Because the full, diffeo-invariant field equations are

$$G_{\mu\nu} = 8\pi G (T_{\mu\nu} + T^{(reg)}_{\mu\nu}), \quad \text{with} \quad \nabla^\mu T^{(reg)}_{\mu\nu} = 0,$$

the bound $K \le K_*$ implies:

1. No geodesic incompleteness from curvature blow-up (all polynomial invariants bounded).
2. No inner-horizon mass-inflation catastrophe in the static limit (the same cap saturates the focusing driver).
3. Exterior Schwarzschild is recovered for r ≫ r0 (since K ≪ K_* there, T^{(reg)}_{μν} → 0).

Physical picture. Gravity "self-regularizes" only where GR would otherwise fail; the TOV flow bends onto a de Sitter–like core before singular behavior can occur. The result is a static, spherical black-hole interior with a smooth center, finite invariants, and a horizon that remains locally innocuous for infallers—precisely the "normal spacetime" target.

5. Rotating (Kerr) Black Holes

5.1 Mass inflation as informational over-focusing

- Classical picture. Kerr's analytic interior contains an inner (Cauchy) horizon at $r = r_{-}$. Any tiny, late-time ingoing/outgoing flux (Price-tail debris, vacuum polarization, accretion trickle) is exponentially blue-shifted there. The effective interior mass parameter $m(v)$ grows like $m \sim e^{\kappa_{-} v}$ (with $\kappa_{-}$the inner-horizon surface gravity), driving large curvature components even when exterior perturbations are microscopic. This is the Poisson–Israel/Ori mass-inflation instability.
- Raychaudhuri/optical scalars view. The null congruences aligned with the inner horizon acquire rapidly increasing expansion/shear due to the blue-shifted stress–energy; the focusing term proportional to $\rho + p_{\parallel}$ along the null direction overcomes any dispersive relief, and curvature scalars race toward blow-up.
- Informational reading. The interior evolution leaves the τ-manifold: the realized model $q$is driven away from the generative constraints by an over-amplified response to vanishingly small external "evidence." In our language, this is informational over-focusing: $D(q \parallel \tau)$ wants to *increase* (the state decouples from valid constraints) as the inner horizon acts like a flawed "amplifier" of mismatched bits.

5.2 IAP-induced saturation and restoration of global hyperbolicity

- Single mechanism, two effects. The information-regularized action adds a conserved, geometric $T_{\mu\nu}^{(\mathrm{reg})}$sourced by the convex $\mathcal{R}(K)$. As $K \to K_{*}$, two things happen automatically:

1. Focusing cap: In the null Raychaudhuri equations, the effective driver $(\rho + p_{\parallel})_{\mathrm{eff}}$ saturates because $\mathcal{R}'(K)$ suppresses further curvature amplification. Null expansions remain finite; shear cannot diverge.
2. Blue-shift quench: The same $\mathcal{R}'(K)$ produces a vacuum-like counter-stress (locally de-Sitter–type) that redresses the exponential gain. Operationally, the would-be $e^{\kappa_{-}v}$ growth is cut off once $K \sim K_{*}$, converting the inner-horizon neighborhood from an unstable amplifier into a soft, finite-curvature transition layer.

- Geometric outcome. The ideal, null Cauchy horizon is replaced by a regular spacelike (or weakly timelike) regulator surface Σ_{reg} inside $r = r_{-}$. The interior then flows into a finite-curvature core (Sec. 4), with no geodesic incompleteness and no extendible region that would violate determinism.
    - No Cauchy pathology: Because curvature and optical scalars are bounded, there is no extendible region with ambiguous evolution; the maximal Cauchy development from generic initial data is unique—global hyperbolicity is restored.
    - Kerr exterior intact: For $r \gtrsim r_{-}$where $K \ll K_{\backslash *}$, $T_{\mu\nu}^{(\mathrm{reg})} \to 0$. Thus all tested Kerr phenomenology (ISCO, shadow, frame-dragging, ringdown at leading order) is preserved.
- Informational statement (the Lyapunov view). The IAP enforces $\dot{D} \leq 0$ everywhere, including the inner-horizon neighborhood. Mass inflation is precisely the classical route where $\dot{D} > 0$; the regulator closes this channel by making the metric-response contractive once the informational curvature hits the threshold.
- Penrose-diagram sketch (verbal). Region I (exterior Kerr) → outer horizon r_{+}→ Region II (interior) → instead of a null Cauchy horizon, a finite-curvature cap Σ_{reg} that bridges to a nonsingular core; no "other universe," no timelike singularity, no closed timelike corridors.
- Observable fingerprints (Kerr-specific).
    - High-spin remnants: A slightly sub-extremal spin ceiling (the inner-horizon gain is self-limited).
    - Late-time tails: Suppressed inner-horizon–induced turbulence alters very-late ringdown exponents for near-extremal a/M.
    - No long-lived echoes: The transition is soft and highly red-damped; any echo-like features are short and weak, unlike hard-wall models.

Net effect: Kerr's inner-horizon catastrophe is reinterpreted as an informational over-focusing instability that our single principle saturates. The interior becomes causal and regular, while the exterior remains Kerr to experimental precision—exactly the "elegant solution" brief.

6. Thermodynamics and Information

6.1 Generalized entropy as stationarity; islands from variation

- Object of extremization. For any candidate quantum extremal surface (QES) $\mathcal{X}$ that partitions a Cauchy slice into outside region $R$and interior I,

S_gen[$\mathcal{X}$] = Area($\mathcal{X}$)/(4G) + S_out[R ∪ I] .

- IAP link. Our Informational Alignment Principle (IAP) demands monotone decrease of D(q‖τ). Varying the *same* information-regularized action that enforces the curvature bound implies the stationarity condition

δ S_gen[$\mathcal{X}$] = 0   (at fixed boundary data) ,

identifying QES surfaces as Euler–Lagrange critical points of the information-regularized dynamics.

- Island rule as a variational minimum. The entropy of radiation R is

S(R) = min_ext_{$\mathcal{X}$} { Area($\mathcal{X}$)/(4G) + S_QFT(R ∪ Island[$\mathcal{X}$]) } .

Here "Island[$\mathcal{X}$]" is the bulk region bounded by $\mathcal{X}$ whose inclusion minimizes S_gen. In our framework, the appearance of islands is not an extra axiom: it's the optimal projection of interior degrees of freedom into the exterior description that most reduces D(q‖τ).

- Consistency triangle.

    1. Geometry: δ Area enforces the usual focusing term with the regulator's finite-curvature response;
    2. Matter: δ S_out encodes the modular-Hamiltonian variation of quantum fields;
    3. Alignment: δ(S_gen) = 0 is precisely the stationarity of the Lyapunov functional that the IAP drives toward.

- Outcome. Islands emerge exactly when they are *required* to keep $\dot{D} \leq 0$ while preserving smooth local EFT for infallers and global purity for the full state.

6.2 Smooth horizons, unitary evaporation, Page curve

- Smooth horizons (no drama). For $K \ll K_*$ near the horizon of an astrophysical hole, the regulator is off: the Unruh state is valid, local curvature is mild, and infallers experience ordinary QFT in curved spacetime. No shocks or "firewalls" are needed or allowed by IAP (they would raise D).
- Unitarity from the same principle. The global state (black hole + environment) remains pure. As evaporation proceeds, the IAP steers the system to *redistribute* correlations so that D(q‖τ) decreases in the exterior description:
    - Early times (pre-Page): The optimal $\mathcal{X}$ hugs the horizon → no island; radiation is nearly thermal and entanglement entropy rises with emitted quanta.
    - Page time: Competing variations make a new extremum with an interior island lower S_gen; the entanglement wedge of the radiation jumps to include part of the interior.
    - Late times (post-Page): With the island, additional radiation reduces S(R); information flows out in fine-grained correlations even as the horizon remains smooth.
- Page curve prediction. The fine-grained entropy of the Hawking radiation rises, peaks near the Page time, then falls to zero as the black hole disappears—precisely because the minimizing QES includes an island after the transition. In our language: the exterior model q_{rad} becomes better aligned with τ (lower D) once it is allowed to encode the island's degrees of freedom.
- Compatibility with the curvature cap. The regulator only turns on deep inside (or near Kerr's inner horizon), never at the outer horizon of macroscopic holes. Thus:
    - Thermal spectrum (leading order) and BH mechanics remain intact;
    - Backreaction is included consistently, giving a gently time-dependent geometry and a well-posed, unitary evaporation history;
    - No echoes from hard walls: the core is soft/high-curvature, not reflective; any imprints on ringdown are tiny (set by r_0/r_s).
- Operational summary.
    - Local infaller view: smooth EFT across the horizon, no anomaly.
    - Distant observer view: islands/quasi-local QES implement the correct coarse-graining so that unitarity and a smooth horizon coexist.

- Global view: IAP supplies the arrow—$\dot{D} \leq 0$—that drives the system from opacity to radiance, reproducing the Page curve without extra postulates.

7. Geometry from Entanglement

7.1 Interior volume / complexity growth under $\dot{D} \leq 0$

- Setup. Let the global (pure) state evolve under the information-regularized dynamics. The Informational Alignment Principle (IAP) enforces $\dot{D} \equiv d/dt\, D(q\|\tau) \leq 0$. While exterior area shrinks during evaporation, the *interior* (maximal slice) volume and the circuit-style complexity $\mathbb{C}$ of the state typically increase for a long epoch.
- Why growth is compatible with $\dot{D} \leq 0$.
  $\dot{D} \leq 0$ is global alignment, not a ban on interior growth. In fact, reducing $D(q\|\tau)$ often requires *creating more structured entanglement*—expanding the set of correlations that make the effective interior geometry smooth and consistent with constraints. Thus:
  - Early/mid epochs: Area $\propto S_{BH}$ decreases slowly, while $\mathbb{C}$ grows (longer scrambling circuits; deeper interior slices).
  - Late epochs (post-Page): $\mathbb{C}$ still grows on intermediate windows but its readout shifts outward as the island enters the radiation's entanglement wedge.
- Quantitative sketch (copy-safe).
  Use any interior volume functional $V_{int}(t)$ (e.g., maximal-slice volume) and a coarse complexity proxy $\mathbb{C}(t)$ obeying
- $d\mathbb{C}/dt \gtrsim \kappa_{scr}$ (system-dependent), with $\kappa_{scr} \sim O(1/r_s)$ for large holes,

while

$\dot{D} = dD/dt \leq 0$.

The same microscopic scrambling that increases $\mathbb{C}(t)$ drives coarse-grained D downward by distributing correlations into τ-consistent patterns (no small-scale "overfitting" to noise).

- Regulator's role. Near would-be singular regions, the curvature cap redirects dynamics into entanglement-supported geometry instead of runaway focusing. Practically: V_{int} grows to a finite maximum set by $K \leq K_*$; curvature scalars remain bounded, and geodesics extend.

Takeaway: The interior's "growth" is not a violation of alignment; it is the mechanism by which alignment is maintained—complexity rises to encode a smooth, τ-consistent interior while $\dot{D} \le 0$ ensures this growth remains contractive at the level of divergence.

7.2 Entanglement-wedge reconstruction and τ-aligned readout

- Two complementary descriptions.
  - Infalling frame: A smooth interior geometry exists; local EFT works near the horizon ($K \ll K_*$).
  - Distant frame: Bulk operators behind the horizon are encoded in boundary/exterior degrees of freedom via entanglement-wedge reconstruction.
- Islands as optimal encoders. For a radiation region R, choose the QES $\mathcal{X}$ that minimizes
- $S_{gen}[\mathcal{X}] = \mathrm{Area}(\mathcal{X})/(4G) + S_{QFT}(R \cup \mathrm{Island}[\mathcal{X}])$.

After the Page transition, $\mathrm{Island}[\mathcal{X}] \subset$ interior enters the entanglement wedge of R. Operationally, this is a τ-aligned readout: the best-aligned exterior model q_R uses interior data columns (the island) because that choice *most reduces* $D(q\|\tau)$.

- State-dependent decoding without cloning. The mapping "interior operator → exterior code" is state-dependent (code subspace). No observer can access both bulk and boundary reconstructions simultaneously, so no-cloning is preserved. Alignment view:
- $\mathrm{Readout}_R = \mathrm{argmin}_q\ D(q\|\tau)$ subject to q supported on observables in R .

The solution implicitly includes island degrees of freedom once the QES condition is met.

- Role of the regulator. The curvature regularizer never forces hard reflective walls; it stabilizes the interior so the code remains decodable (bounded K), and it prevents inner-horizon pathologies that would spoil a consistent wedge.
- Information flow picture.

1. Pre-Page: wedge(R) excludes the interior; readout is thermal-like (large D reduction unavailable).

2. Page point: new extremum with island appears; wedge(R) expands inside; D drops faster because reconstruction uses the τ-consistent interior code.

3. Late times: wedge(R) includes most of the interior relevant for decoding early infall; unitary readout with a smooth horizon for infallers is maintained.

Takeaway: Entanglement-wedge reconstruction is the mathematical form of τ-aligned readout. Islands aren't an add-on; they are the variationally optimal way to lower $D(q\|\tau)$ for the exterior description, exporting the interior's "it" into accessible "bits" without sacrificing smoothness or causality.

8. Observational Signatures

8.1 Ringdown shifts ~ $O((r_0/r_s)^2)$

- Scaling. Let $r_s = 2GM$ ($c = 1$). The curvature cap sets an interior scale
  $r_0 \approx c_0 (GM)^{1/3} L_*^{2/3}$ with $c_0 \approx (48)^{1/6} \simeq 1.4$, so
- $r_0/r_s \sim C \cdot (L_*/r_s)^{2/3}$, with $C \simeq O(1)$.

For any perturbation that samples the near-horizon geometry, the fractional shifts in quasinormal-mode (QNM) frequencies and damping times scale as

$\Delta\omega/\omega \sim \alpha \cdot (r_0/r_s)^2$, $\Delta\tau/\tau \sim \beta \cdot (r_0/r_s)^2$,

with α, $\beta = O(1)$ constants depending weakly on spin and mode (ℓ,m,n).

- Benchmarks (Planckian L_*).
  – Stellar BH (M ~ 10 M☉): (L_*/rs) ~ 10^(−40) ⇒ (r0/rs)^2 ~ 10^(−53) (undetectable).*
  *– SMBH (M ~ 10^6–10^9 M☉): (L_*/rs) ≲ 10^(−47…−50) ⇒ (r0/rs)^2 ≲ 10^(−63…−66).
- Interpretation. If L_* is Planckian, ringdown looks Kerr to exquisite precision. Any measurable $\Delta\omega/\omega$ would indicate $L_* \gg L_P$ or additional long-range effects beyond the minimal regulator.

8.2 High-spin tail modifications; sub-extremal spin ceiling

- Inner-horizon quench. In Kerr, the mass-inflation gain ~ $\exp(\kappa_- v)$ is self-limited once $K \to K_*$; the would-be Cauchy horizon becomes a soft, finite-curvature layer. This changes the very-late-time tail and the approach to extremality.
- Spin ceiling (conceptual). The regulator imposes an effective cap on the amplification that ordinarily supports near-extremal ($a/M \to 1$) inner-structure. One expects
- $a/M \le 1 - \varepsilon_*$, with $\varepsilon_* \sim \gamma \cdot (r_0/r_s)^2$,

where $\gamma = O(1)$. Thus, in this framework true extremality is *not* reached dynamically; near-extremal remnants sit a tiny distance below 1.

- Late-time tails. For high a/M, the red-shifted coupling of exterior modes to the interior is altered at order $(r0/rs)^2$, slightly modifying the power-law/oscillatory tails at very late times (beyond the main QNM ringdown). This is only testable if L_* is very large compared to L_P *or* if precision on post-ringdown tails improves by many orders of magnitude.
- Observational handle. Population studies from LIGO–Virgo–KAGRA/ET/CE: a hard pile-up *just below* $a/M = 1$ (rather than a smooth asymptote) would be consistent with a small, universal $\varepsilon\_$*; conversely, evidence of true extremality would constrain* $\varepsilon\_ \to 0$ and push K_* upward.

8.3 Shadow micro-corrections and constraints on K_*

- Photon ring & shadow edge. The regulator leaves the exterior Kerr metric intact to leading order; corrections leak in only via the tiny ratio $(r0/rs)$. The photon-sphere radius r_ph and thus the shadow diameter D_sh admit fractional shifts
- $\Delta r\_ph/r\_ph \sim O((r0/rs)^2)$, $\Delta D\_sh/D\_sh \sim O((r0/rs)^2)$.
- EHT-class bounds (order-of-magnitude). For present EHT reconstructions with percent-level uncertainties, requiring $|\Delta D\_sh/D\_sh| \lesssim 0.1$ gives a completely trivial bound ($r0/rs \lesssim 0.3$). Translating to the curvature scale:
- $r0 \approx c_0 (GM)^{1/3} L_*^{2/3}$
- $\Rightarrow L_* \approx (r0/c_0)^{3/2} (GM)^{-1/2}$,
- $\Rightarrow K_* = L_*^{-4} \approx c_0^{6} (GM)^{2} / r0^{6}$.

With $r0/rs \ll 1$ for astrophysical BHs, current shadow data allow K_* to be vastly above any laboratory scale; i.e., no tension with a Planckian K_*.

- Where shadows help. If future VLBI resolves the sub-rings and measures their spacings/phases precisely, the tiny $O((r0/rs)^2)$ lensing corrections could be integrated over many orbits to tighten bounds on L_*; absent that, ringdown is the sharper probe.
- Joint constraints. Combine (i) $\Delta\omega/\omega$ limits from QNMs, (ii) high-spin population edge ($\varepsilon\_$*), and (iii) shadow systematics. A null result across all three pushes* $L\_$ downward (toward L_P) and K_* upward; any consistent, mode- and mass-dependent deviation ~ $(r0/rs)^2$ across these channels would point to a universal curvature scale rather than object-specific exotica.

Summary. In the minimal, elegant completion, all observable deviations are second-order in (r0/rs). For Planckian L_* they are undetectably small today; nonetheless, they define crisp, universal templates—QNM shifts, a sub-extremal spin ceiling, and photon-ring micro-corrections—that future precision can target to confirm or bound the curvature cap K_* that normalizes black-hole interiors.

9. Beyond Black Holes

9.1 Stars/galaxies as q≈τ attractors

- Alignment lens.
  In ordinary astrophysics, long-lived structures emerge when microscopic transport, radiation, and gravity self-consistently satisfy the macroscopic constraints (equations of state, hydrostatic/virial balance, angular-momentum transport). In our language, the realized model q (state of matter/fields) relaxes toward the generative constraints τ (conservation laws + boundary conditions), yielding q≈τ.
- Stars (local attractors).
  - Formation: turbulent, magnetized clouds cool and collapse until heating = cooling at the core—ignition sets a fixed point where nuclear power balances radiative/convective flux.
  - Main sequence as a manifold: varying mass/composition traces a low-dimensional τ-manifold; stellar evolution is slow drift along it as fuel and opacity laws change—$\dot{D} \le 0$ at the coarse scale (net predictive mismatch falls).
  - End states: white dwarfs, neutron stars, and stable remnants are boundary attractors where further alignment is blocked by degeneracy/strong-interaction constraints; failed alignment (over-compression) hands off to the black-hole branch handled earlier.
- Galaxies (meso-scale attractors).
  - Virialization: violent relaxation and feedback (stellar winds, SNe, AGN) drive energy and angular momentum into distributions consistent with τ (potential, baryon cycle) → coherent rotation laws, scaling relations (e.g., Tully–Fisher, fundamental plane).
  - Feedback as alignment control: star-formation and AGN feedback operate like proportional–integral controllers, preventing runaway collapse or dispersion—keeping q near the τ-implied equilibrium cones of gas supply and entropy.

- Dark sector as "informational curvature."
  What we treat phenomenologically as dark matter/energy behaves, in this frame, like missing columns of τ: latent constraints that close momentum and Friedmann balances. As cosmological inference sharpens, τ is effectively completed, reducing $D(q\|\tau)$ across large scales (better fits with fewer arbitrary priors).

Takeaway: Ordinary luminous structure is what alignment looks like in the wild: systems radiate coherence as they settle onto τ-manifolds; dissipation is the price of descending $D(q\|\tau)$.

9.2 Divergent minds vs. informational singularities

- Cognitive q and τ.
  - q (mind): the agent's working model—beliefs, priors, predictive machinery.
  - τ (world/Logos): the generative structure behind observations—regularities, norms, other minds, and reality-constraints discovered through interaction.
- Coherence and radiance.
  Learning, dialogue, and error-correcting feedback lower $D(q\|\tau)$: signals sharpen, action becomes skillful, and the person *radiates* coherence (clear explanation, reliable prediction, moral steadiness).
- Divergence modes (toward informational singularity).
  - Over-focusing (confirmation runaway): selective amplification of agreeable evidence (the cognitive analog of inner-horizon blue-shift) raises D even as subjective confidence spikes.
  - Model collapse (isolation): shrinking the hypothesis class until it cannot accommodate counterevidence; updates stall, like geodesics ending prematurely.
  - Adversarial noise (memetic turbulence): high-gain inputs saturate attention/affect circuits, decoupling q from τ (misinformation cascades, cultic closure).
- IAP for minds (practical program).
  - Contraction maps: design update rules and environments that satisfy data-processing inequalities—every epistemic step should non-increasingly transform divergence.

  - Curvature regularizer (cognitive). Introduce penalties for over-curvature in belief networks (e.g., extreme certainty without error bars, brittle causal links). The analog of $\mathcal{R}(K)$ dampens runaway confidence and reopens exploration.
  - Islands = integrative contexts. Create safe wedges where dissonant facts are integrated with supportive scaffolds (therapeutic alliance, peer review, mentoring). The agent's "entanglement wedge" enlarges to include previously excluded interior content, lowering D without trauma.
- Recovery trajectories.
  - Soft cores, not hard walls: avoid humiliation and punitive contradiction (firewalls) that shatter local smoothness; use gentle curvature caps (bounded certainties, calibrated doubt).
  - Complexity before simplicity: allow temporary complexity growth (more nuanced models, multiple perspectives) so that a simpler, truer compression is achievable later—mirroring interior-volume/complexity growth that precedes outward radiance in black holes.
- Ethical corollary.
  Institutions that maximize short-term engagement via over-focusing signals engineer informational singularities. τ-aligned pedagogy and media, by contrast, enforce $\dot{D} \le 0$ at the population scale: transparency, corrigibility, uncertainty quantification, and shared standards of evidence.

Takeaway: A "divergent mind" is the cognitive counterpart of a black-hole singularity: an update rule that lets D grow without bound. The cure is the same elegance we used in gravity—a monotone alignment principle with a soft curvature cap—so that learning remains smooth, globally coherent, and ultimately radiant.

10. Conclusion

10.1 One covariant principle, continuous q→τ, normal spacetime everywhere

Conclusion in one line: a single, covariant alignment rule—$\dot{D} \equiv d/dt\, D(q\|\tau) \le 0$, enforced by a convex curvature regulator $\mathcal{R}(K)$—turns black holes into ordinary regions of spacetime while preserving everything already confirmed about GR and QFT.

What this achieves, coherently and at once

- Continuity of q→τ: Physical evolution is a constrained gradient flow that monotonically reduces mismatch; coherence is not imposed locally but emerges globally as D falls.
- Normal interiors: The curvature bound $K \le K^*$ redirects would-be singular focusing into a smooth, finite-invariant core; geodesics extend; inner-horizon instabilities saturate rather than explode—global hyperbolicity is restored.
- Smooth horizons + unitarity: Near macroscopic horizons $K \ll K^*$, so local EFT remains intact (no drama). Fine-grained unitarity follows from the same variational principle: generalized-entropy stationarity produces quantum extremal surfaces and islands, yielding the Page curve without extra postulates.
- Single action, no patchwork: Einstein–Hilbert + matter + $\mathcal{R}$(K) is diffeomorphism-invariant and Bianchi-consistent. GR is recovered in the low-curvature limit; only the extreme interior is modified, exactly where GR fails.
- Geometry from entanglement: Interior volume and circuit complexity can grow while D decreases; later, entanglement-wedge reconstruction exports interior information to the radiation—alignment deepens even as the area shrinks.
- Falsifiability: All deviations scale ∼(r0/rs)^2 with r0 ≈ const · (G M)^(1/3) L_*^(2/3). *Ringdown micro-shifts, a sub-extremal spin ceiling, and photon-ring refinements provide clean targets; null results raise K* (push L_* down), positive results fix a universal curvature scale.

Broader arc: Stars and galaxies exemplify q≈τ attractors; divergent minds mirror unregularized focusing. One law—informational alignment with a soft curvature cap—governs the journey from opacity to radiance across domains. Under this principle, black holes cease to be exceptions; they become proofs that the atlas of physics can be made globally smooth.

Epilogue: From Many Puzzles to One Principle

What began as an unruly catalog—singularities, information loss, inner-horizon chaos, firewall drama—collapses to a single rule: alignment. When physical evolution is required to monotonically reduce the divergence $D(q \parallel \tau)$, and curvature is softly bounded by a convex regulator $\mathcal{R}(K)$, black holes become ordinary regions of a globally coherent spacetime. ’t Hooft’s holographic intuition then serves as the natural coarse-graining of this flow: the area term is the geometric piece of a Lyapunov functional whose stationarity selects quantum extremal surfaces and islands, reconciling smooth horizons with unitary evaporation.

Practically, the picture is simple. Singularities are replaced by finite-curvature cores; mass inflation at Kerr's inner horizon saturates; global hyperbolicity is restored; firewalls are unnecessary because near-horizon curvature is well below the regulator's threshold. What once stood as independent principles reappear as corollaries of one variational statement. Where classical GR would over-focus, the informational regulator bends the dynamics back toward consistency.

What remains is empirical and computable. First, fix the curvature scale at which regularization turns on, K_*. If K_* is Planckian, deviations hide beneath current sensitivity; if higher, they surface as universal, second-order corrections $\sim (r_0/r_s)^2$ in ringdown spectra, a sub-extremal spin ceiling, and photon-ring micro-structure. Second, sharpen the operational definition of $D(q \parallel \tau)$ on concrete algebras and code subspaces so alignment is a measurable quantity, not a slogan.

Seen this way, black holes are not exceptions but examiners. They test whether the same alignment flow governing stars and minds also governs horizons and cores. Geometry from entanglement ceases to be metaphor: interior volume and circuit complexity can grow while D falls—added structure is the cost of a more faithful model—and islands later export that structure as correlations in the radiation.

The wager is minimal yet bold: keep covariance, conservation, and every tested prediction of GR and QFT where curvature is small; add one ingredient—a convex penalty on large curvature—and insist that evolution never increases the misfit between the realized world and its generative constraints. Whether future data bound K_{star}deep in the ultraviolet or measure it directly, the moral is the same: many stories, one grammar. With $q \to \tau$ as the arrow and K_* as the only new scale, black-hole physics is not discarded—it is explained.

References

1. Youvan, D. C. (October 2025). *Heavenly Minds: Applying Informational Tunneling to Stars, Galaxies, and Divergent Intelligences*. ResearchGate. https://www.researchgate.net/publication/396524863_Heavenly_Minds_Applying_Informational_Tunneling_to_Stars_Galaxies_and_Divergent_Intelligences
2. Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D, 7*, 2333–2346.
3. Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics, 43*, 199–220.
4. Bardeen, J. M., Carter, B., & Hawking, S. W. (1973). The four laws of black hole mechanics. *Communications in Mathematical Physics, 31*, 161–170.
5. Penrose, R. (1965). Gravitational collapse and space-time singularities. *Physical Review Letters, 14*, 57–59.
6. Hawking, S. W., & Penrose, R. (1970). The singularities of gravitational collapse and cosmology. *Proceedings of the Royal Society A, 314*, 529–548.
7. Hawking, S. W., & Ellis, G. F. R. (1973). *The Large Scale Structure of Space-Time*. Cambridge University Press.
8. Ellis, G. F. R., Maartens, R., & MacCallum, M. A. H. (2012). *Relativistic Cosmology*. Cambridge University Press.
9. Finkelstein, D. (1958). Past-future asymmetry of the gravitational field of a point particle. *Physical Review, 110*, 965–967.
10. Kruskal, M. D. (1960). Maximal extension of Schwarzschild metric. *Physical Review, 119*, 1743–1745.
11. Szekeres, G. (1960). On the singularities of a Riemannian manifold. *Publicationes Mathematicae Debrecen, 7*, 285–301.
12. Oppenheimer, J. R., & Snyder, H. (1939). On continued gravitational contraction. *Physical Review, 56*, 455–459.
13. Poisson, E., & Israel, W. (1990). Internal structure of black holes. *Physical Review D, 41*, 1796–1809.
14. Ori, A. (1991). Inner structure of a charged black hole: An exact mass-inflation solution. *Physical Review Letters, 67*, 789–792.

15. Bardeen, J. M. (1968). Non-singular general relativistic gravitational collapse. In *Proceedings of GR5* (Tbilisi).

16. Frolov, V. P., Markov, M. A., & Mukhanov, V. F. (1990). Black holes as possible sources of closed and semiclosed worlds. *Physical Review D, 41*, 383–394.

17. Susskind, L., Thorlacius, L., & Uglum, J. (1993). The stretched horizon and black hole complementarity. *Physical Review D, 48*, 3743–3761.

18. Page, D. N. (1993). Information in black hole radiation. *Physical Review Letters, 71*, 3743–3746.

19. Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy. *Physical Review Letters, 96*, 181602.

20. Hubeny, V. E., Rangamani, M., & Takayanagi, T. (2007). A covariant holographic entanglement entropy proposal. *Journal of High Energy Physics, 2007*(07), 062.

21. Engelhardt, N., & Wall, A. C. (2015). Quantum extremal surfaces: Holographic entanglement entropy beyond the classical regime. *Journal of High Energy Physics, 2015*(01), 073.

22. Penington, G. (2020). Entanglement wedge reconstruction and the information paradox. *Journal of High Energy Physics, 2020*(09), 002.

23. Almheiri, A., Engelhardt, N., Marolf, D., & Maxfield, H. (2019). The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole. *Journal of High Energy Physics, 2019*(12), 063.

24. Hayden, P., & Preskill, J. (2007). Black holes as mirrors: Quantum information in random subsystems. *Journal of High Energy Physics, 2007*(09), 120.

25. Sekino, Y., & Susskind, L. (2008). Fast scramblers. *Journal of High Energy Physics, 2008*(10), 065.

26. Brown, A. R., Roberts, D. A., Susskind, L., Swingle, B., & Zhao, Y. (2016). Holographic complexity equals action. *Physical Review Letters, 116*, 191301.

27. Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In W. Zurek (Ed.), *Complexity, Entropy, and the Physics of Information*. Addison-Wesley.

28. Bousso, R. (1999). A covariant entropy conjecture. *Journal of High Energy Physics, 1999*(07), 004.

29. Wall, A. C. (2010). A proof of the generalized second law for rapidly changing fields and arbitrary horizon slices. *Physical Review D, 81*, 104018.

30. Frolov, V. P., & Zelnikov, A. (2011). *Introduction to Black Hole Physics*. Oxford University Press.