Black Hole Information in an Eidetic Universe

A Bidirectional, Null-Propagating Derivation of the Page Curve

Daniel Stoker

(rewritten and restructured edition)

Published: July 3, 2022 | Restructured: June 8, 2026

Abstract

We reformulate black-hole evaporation in Eidetic Theory so that its quantitative results are *derived from* the distinctive primitives of the theory rather than imported as generic finite-dimensional postulates. Eidetic Theory replaces the inaccessible spacelike past/future boundaries of de Sitter holography with an infinite network of *Eide Spheres*: quantized $2S+T$ spherical surfaces that propagate at the speed of light and exist in temporally bidirectional entangled pairs, one future-directed and one past-directed. The emergent bulk, Hawking modes, horizons, singularities, and islands are reconstructed representations of this boundary code.

Three mechanisms carry the physics, and each replaces an axiom that the previous formulation assumed. (i) *Bidirectional pairing* supplies a built-in, non-postulated final-state structure: the pair Bell state *realizes an Horowitz–Maldacena-type final-state boundary condition without post-selection*—being part of the primitive Eide pair rather than an imposed projection, subject to a benign-reconstruction condition—from which global fixed-sector unitarity (with the radiation seen through an induced channel) and the isometric (non-deleting) character of the singularity return are derived, and from which the AMPS duplicate-purifier obstruction is removed algebraically as a built-in ER = EPR (horizon smoothness remaining conditional on the benign-reconstruction condition). Bidirectional pairing supplies one-branch mixedness—the *rising* Page branch—but, as we make explicit, the Page *turnover* requires a finite shared Eide reservoir whose degrees are transferred by emission; bare private-partner pairing yields only the monotone Hawking line. (ii) *Null propagation* fixes the *candidate local* interaction graph of the scrambling dynamics: two Eide Spheres interact precisely when their null sheets intersect, so the locality and connectivity needed for an approximate unitary design are kinematic consequences rather than assumptions. For the Page curve we require only an inverse-polynomial second-moment gap of the induced edge-update chain—the focused-graph polynomial-gap condition, which we assume for these near-horizon graph families and which remains a graph-theoretic theorem target; graph expansion is *not* required and would only upgrade the result to fast scrambling. (iii) *Geometric quantization of the Eide S^2 flux* gives a two-state horizon-cell alphabet in the minimal sector; in the balanced minimal independent-cell model the horizon transfer operator has Perron/topological-entropy eigenvalue $\lambda_\star = 2$, yielding the intrinsic area density $\eta_{\mathrm{Eide}} = \ln \lambda_\star / a_\star$ and, after Jacobson-type matching to the emergent Einstein response, the area quantum $a_\star = 4\ell^2_{\mathrm{Pl,phys}} \ln 2$.

The resulting finite-code Page theorem states that for every fixed typicality failure probability $\delta > 0$ there is a constant $C_{\delta,\gamma} = O_\delta(1)$ with

$$\mathbb{P}_{\mathrm{SVP}}\big[\, |S_R(t) - \min\{\ln d_R(t), \ln d_C(t)\}| \le C_{\delta,\gamma} \,\big] \ge 1 - \delta,$$

and the geometric Page curve follows with only logarithmic microstate-counting corrections. Because the construction is genuinely Eidetic, it makes predictions a generic relabeling cannot: a discrete horizon area spectrum with quantum $a_\star = 4\ell^2_{\mathrm{Pl,phys}} \ln 2$ in the balanced minimal

independent-cell model, a shared subleading $O(\ln A)$ entropy coefficient for black-hole and de Sitter static-patch horizons *when they lie in the same fixed-center Eide surface-code universality class*, and a candidate backward-branch Page-correction target near the turnover. We are explicit about the residual cost of this honesty. The scrambling step is reduced but not eliminated: null propagation supplies the local connected interaction graph, while polynomial-depth design behavior requires the focused-graph inverse-polynomial gap condition (connectivity alone gives only a positive gap, hence a design in some unbounded depth). Graph expansion is *not* needed for the Page curve, but would upgrade the result to fast scrambling. The associated lifetime comparison is a consistency check, conditional on a gate-rate model and a design-depth exponent. In the balanced minimal independent-cell realization of the boundary axiom, $k = 1$ flux quantization and a provably subextensive global closure give $\lambda_\star = 2$; a possible R_ω-induced local flux-matching could lower this value. The benign-reconstruction obligation splits in two: the recovery-fidelity part (no teleportation corruption) *follows* from the same approximate 2-design as the Page curve, via Hayden–Preskill decoupling, and is demonstrated on the toy R_ω; only horizon *smoothness* (the firewall question) remains genuinely conjectural. What genuinely remains is therefore horizon smoothness, the universality of the local SVP gate family, the focused-graph gap condition, the single-reservoir/no-remnant condition, and the charged/rotating generalization of the area count.

Contents

1 Purpose, scope, and logical status

Hawking's semiclassical calculation shows black holes emit radiation with a thermal character [1]. Page's finite-dimensional theorem gives the signature of *unitary* evaporation: for a typical pure state on a bipartite Hilbert space the smaller subsystem is almost maximally entangled with its complement [3]. The information problem is whether the fine-grained radiation entropy follows the Page curve during evaporation.

This paper derives that curve inside Eidetic Theory. The earlier formulation of the present author proved a correct but *detachable* result: replacing "Eide Sphere" by "finite-dimensional subsystem" left every theorem intact, because the physical content entered through three generic postulates—a scrambling axiom, an area-law counting axiom, and the full thermodynamic skeleton of general relativity. The present edition is organized so that each of those postulates is instead a *consequence* of a primitive Eidetic structure. The logical spine is now:

$$
\begin{aligned}
&\text{bidirectional pair Bell network} &&\Longrightarrow \quad \text{global fixed-sector purity + induced one-branch CPTP channel + isometric return;}\\
&\text{null propagation of } S^2 \text{ surfaces} &&\Longrightarrow \quad \text{local, connected } \textit{candidate} \text{ SVP interaction graph (expansion is a separate hypothesis);}\\
&\text{geometric quantization of } S^2 \text{ flux} &&\Longrightarrow \quad \text{two-state cell alphabet, and } \lambda_\star = 2 \text{ under the minimal unconstrained-cell hypothesis;}\\
&+ \text{ Jacobson-type emergent-Einstein matching} &&\Longrightarrow \quad \text{geometric Page curve.}
\end{aligned}
$$

One structural lesson of this organization deserves emphasis at the outset, because it is the sharpest non-generic Eidetic claim and it is easy to overstate the role of pairing. Bidirectional pairing supplies *one-branch mixedness*—the thermal character of a single branch, hence the *rising* Page branch—but it does *not* by itself produce the Page turnover. The falling branch requires that the backward-branch purifier be a *finite shared reservoir* whose degrees are transferred by emission, not a supply of private partners; if each Hawking mode kept its own backward partner the radiation entropy would rise monotonically as in Hawking's calculation (Lemma 3, Figure 2). The distinctive Eidetic content is therefore not merely that a complement purifies the radiation, but that the backward branch must behave as a shared finite reservoir—a claim we make precise and exhibit in a toy model.

Two honest caveats are stated up front and never hidden. First, the bidirectional-unitarity derivation (§4) realizes the Horowitz–Maldacena final-state idea *without ad hoc post-selection*, but it inherits a reconstruction obligation, which we split (§4.5) into two parts of different status: the recovery-fidelity part (no teleportation corruption) is *discharged*—it follows from the same approximate 2-design as the Page curve, via Hayden–Preskill decoupling (Proposition 3), and is demonstrated on the toy R_ω—while horizon *smoothness* (the firewall question) remains a genuine conjecture (Conjecture 1). Second, the null-intersection interaction graph (§5) is proven local and connected in the black-hole sector. We had previously isolated graph *expansion* as a further hypothesis; §10.1 weakens it to the focused-graph polynomial-gap condition (Definition 6): connectivity alone yields a positive second-moment gap (a design in some depth), and the polynomial-gap condition—assumed for the focused near-horizon graph families and remaining a graph-theoretic theorem target—upgrades this to a polynomial-depth approximate 2-design (Lemma 4, Corollary 3), which forms within the evaporation lifetime under a stated gate-rate model and design-depth exponent (Lemma 5). Full expansion (Conjecture 2) survives only as an *optional sharpening* giving fast scrambling, not as a hypothesis of the Page theorem. The remaining genuine inputs are therefore the benign-reconstruction condition, the universality of the local SVP gate family, the focused-graph gap condition, and (for the area coefficient) the Jacobson-type matching. The paper proves *less unconditionally* than a generic finite-dimensional argument in exchange for the results being genuinely Eidetic and for a set of falsifiable predictions (§14).

Throughout, entropy is in nats, and $c = k_B = 1$ unless restored. An $O(1)$ term is bounded independently of the large-code parameters $n_{\mathrm{tot}}, d_R, d_C$, or the semiclassical area.

2 The Eidetic boundary replacement

2.1 Replacing asymptotic past/future boundaries by a bidirectional net

AdS/CFT uses a timelike asymptotic boundary [21, 22]. In de Sitter the natural global boundaries are the spacelike past and future infinities $\mathcal{I}^-, \mathcal{I}^+$ [23]; these are not operational finite-observer boundaries, and the attempts to run an AdS/CFT-style dictionary against them have no local boundary algebra to anchor. Eidetic Theory makes a different microscopic move: it replaces the

role of the two inaccessible spacelike global boundaries by an infinite *network of paired spacelike past/future boundary surfaces.*

Axiom 1 (Eidetic boundary replacement)**.** The asymptotic pair $\mathcal{I}^- \cup \mathcal{I}^+$ is not fundamental. Its past/future role is represented microscopically by a pre-bulk bidirectional Eide-Sphere net

$$V_0 = \{\Sigma_{\alpha,+}, \Sigma_{\alpha,-} : \alpha \in \mathbb{N}\}, \qquad \Sigma_{\alpha,\sigma} \cong S^2,\ \sigma \in \{+,-\},$$

with pair involution $\tau(\alpha,+) = (\alpha,-)$, $\tau(\alpha,-) = (\alpha,+)$ encoding the local past/future pairing at every Eide pair.

This is a replacement of *role*, not an identification of objects. Eide Spheres are not $\mathcal{I}^{\pm}$, not stretched horizons, and not submanifolds of an existing bulk. They are primitive quantum boundary-code degrees of freedom from which bulk representations emerge. Replacing two global surfaces by a network of local pairs is what later supplies, in one stroke, both the purity premise of the Page theorem (§4) and the de Sitter/black-hole unification (§13).

2.2 The $2S + T$ interpretation

An Eide Sphere is a quantized two-sphere together with a temporal orientation. "$2S + T$" means two spatial surface dimensions plus one temporal branch label. The label is not a coordinate on a pre-existing spacetime; it is an orientation datum that reconstructs as a null future- or past-directed bulk branch. For each site (α, σ) the primitive data are

$$(\Sigma_{\alpha,\sigma},\ \Omega_{\alpha,\sigma},\ L_{\alpha,\sigma},\ \mathcal{H}_{\alpha,\sigma},\ T_{\alpha,\sigma}),$$

with $\Sigma_{\alpha,\sigma} \cong S^2$, $\Omega_{\alpha,\sigma}$ a quantized two-form, $L_{\alpha,\sigma}$ its prequantum line bundle, $\mathcal{H}_{\alpha,\sigma}$ the finite site Hilbert space, and $T_{\alpha,\sigma}$ the branch orientation. Section 3 shows that the dimension of $\mathcal{H}_{\alpha,\sigma}$ is fixed by the quantized flux, with the minimal sector giving $\mathcal{H}_{\alpha,\sigma} \cong \mathbb{C}^2$.

3 Eide-Sphere flux quantization and the finite code

Because the area law and the value of $\lambda_\star$ will be *derived* from this quantization rather than postulated, the single-sphere construction is placed at the front of the development.

Proposition 1 (Single Eide-Sphere flux Hilbert space)**.** *Let $\Sigma \cong S^2 \cong \mathbb{CP}^1$ carry a closed two-form Ω with quantized flux*

$$\frac{1}{2\pi}\int_\Sigma \Omega = k \in \mathbb{Z}.$$

Geometric (Bohr–Sommerfeld) quantization with prequantum line bundle $\mathcal{O}(k)$ for $k \geq 0$ gives a finite site Hilbert space

$$\mathcal{H}_\Sigma^{(k)} \cong H^0(\mathbb{CP}^1, \mathcal{O}(k)), \qquad \dim \mathcal{H}_\Sigma^{(k)} = k + 1,$$

and the symplectic area of Σ in flux units is proportional to k. For the past-directed branch one uses the conjugate (antiholomorphic) polarization*: since $H^0(\mathbb{CP}^1, \mathcal{O}(-k)) = 0$ for $k > 0$, the relevant space is $\overline{H^0(\mathbb{CP}^1, \mathcal{O}(k))} \cong H^0(\mathbb{CP}^1, \mathcal{O}(k))^*$, the complex conjugate of the holomorphic site space, of the same dimension $k + 1$ but carrying the conjugate* SU(2) *representation. The sign of k records temporal orientation; it does not change the site dimension.*

Proof. Standard Borel–Weil for SU(2): holomorphic sections of $\mathcal{O}(k)$ over $\mathbb{CP}^1$ form the spin-$k/2$ representation of dimension $k+1$. The Bohr–Sommerfeld cell count equals $\dim H^0$ up to the metaplectic shift absorbed in $O(1)$. The conjugate (antiholomorphic) polarization gives the complex-conjugate space, of equal dimension and opposite charge; one does not use $H^0(\mathbb{CP}^1, \mathcal{O}(-k))$, which vanishes for $k > 0$. □

The minimal nontrivial sector $k = 1$ gives the qubit cell $\mathcal{H}_{\alpha,\sigma} \cong \mathbb{C}^2$, with pair convention $k_{\alpha,+} = +1,\ k_{\alpha,-} = -1,\ k_{\alpha,+} + k_{\alpha,-} = 0$, and $K_\alpha = \mathcal{H}_{\alpha,+} \otimes \mathcal{H}_{\alpha,-}$. We work in this minimal sector; general k is recovered by $\dim = k + 1$ wherever 2 appears.

Finite cutoff. At cutoff N set $I_N = \{1, \dots, N\} \times \{+, -\}$. Each site carries $\mathcal{H}_{\alpha,\sigma} \cong \mathbb{C}^2$, $\mathcal{A}_{\alpha,\sigma} = B(\mathcal{H}_{\alpha,\sigma}) \cong M_2(\mathbb{C})$, each pair $K_\alpha \cong \mathbb{C}^2 \otimes \mathbb{C}^2$. The finite microscopic Hilbert space and algebra are

$$\mathcal{H}_N = \bigotimes_{\alpha=1}^{N} K_\alpha, \qquad \mathcal{A}_N = B(\mathcal{H}_N),$$

with the quasi-local Eide algebra $\mathcal{A}_{\text{Eide}} = \overline{\bigcup_{N<\infty} B(\mathcal{H}_N)}^{\|\cdot\|}$. The Page theorem is proved at finite cutoff; no infinite-dimensional entropy subtlety is needed.

Entropy capacity versus state entropy. Each site has capacity $c_\alpha = \ln d_\alpha$, $d_\alpha = \dim \mathcal{H}_{\alpha,\pm}$, and for a set X of logical modes $C(X) = \sum_{\alpha \in X} \ln d_\alpha$. This dimension count is *not* the von Neumann entropy $S(\rho_X) = -\operatorname{Tr} \rho_X \ln \rho_X$. The Page curve is derived only after the state is shown to saturate the dimension-based envelope; keeping the two notions separate is what prevents circularity (§17).

4 Bidirectional pairing as the mechanism of unitarity

This section is the heart of the rewrite. In the earlier formulation, fixed-sector purity and the isometric (entropy-preserving) character of the singularity were *assumed*. Here bidirectional pairing supplies the *global purifier structure*, and a code-isometric dynamical update then preserves fixed-sector purity; together these yield global fixed-sector unitarity/isometry and make the singularity return naturally entropy-preserving. We are explicit that bidirectional pairing alone does not derive the full microscopic dynamical law—the code-isometric repartitioning is the dynamical input—but it supplies the purifier and the final-state structure of Horowitz–Maldacena type that is created with the pair rather than post-selected.

4.1 The pair Bell state as an HM-type final-state structure

For each Eide pair choose the reference Bell state

$$|\Omega_\alpha\rangle = \tfrac{1}{\sqrt{2}}\big(|0\rangle_{\alpha,+}|0\rangle_{\alpha,-} + |1\rangle_{\alpha,+}|1\rangle_{\alpha,-}\big), \qquad S(|\Omega_\alpha\rangle\langle\Omega_\alpha|) = 0. \tag{1}$$

Either single branch is maximally mixed, $\rho_{\alpha,\pm} = \operatorname{Tr}_{\alpha,\mp} |\Omega_\alpha\rangle\langle\Omega_\alpha| = \frac{1}{2}\mathbb{I}_{\alpha,\pm}$, so $S(\rho_{\alpha,\pm}) = \ln 2$. This is the elementary Eidetic mechanism: the complete bidirectional pair is pure while one branch looks mixed under restriction. A thermally weighted pair $|\Omega_\alpha(\beta)\rangle = Z_\alpha^{-1/2} \sum_i e^{-\beta E_i/2} |i\rangle_{\alpha,+}|i\rangle_{\alpha,-}$ gives the outgoing restriction $\rho_{\alpha,+} = e^{-\beta H_{\alpha,+}}/Z_\alpha$, the finite-code form of Hawking thermality. The Page curve below needs only the global bidirectional purity and the radiation/complement split, not the spectral form.

The crucial structural point is the following. In Horowitz–Maldacena [24] one *imposes* a maximally entangled state on the infalling matter and the interior Hawking partners at the singularity, and the standard objection is that this projection is ad hoc and threatens linearity. In Eidetic Theory the maximally entangled state (1) is not a projection added at the end of evaporation; it is the defining entanglement of the past/future pair, present from the boundary network. The future-directed branch reconstructs as outgoing radiation; the past-directed branch reconstructs as the interior partner that meets the singularity. The Bell pair thus plays the *algebraic role* of an HM-type final-state boundary condition, but because it is part of the primitive Eide pair rather than an imposed projection, it avoids ad hoc post-selection *provided* the reconstruction map R_ω implements the corresponding teleportation channel without corruption—the recovery-fidelity part of benign reconstruction (part 1a, Definition 1), which is discharged from the design in Proposition 3. The identification of the Bell pair with the literal HM matter/interior projection therefore holds at the level of faithful teleportation; the separate question of horizon smoothness (part 1b, Conjecture 1) is not thereby settled.

4.2 Energy and oriented flux

Temporal orientation is encoded by branch flow, not negative physical energy. With site energy operators $\widehat{E}_{\alpha,\sigma} \geq 0$ and the oriented flux convention $F_{\alpha,+} = +\widehat{E}_{\alpha,+}$, $F_{\alpha,-} = -\widehat{E}_{\alpha,-}$, a pair can carry canceling oriented flux while both branches carry positive energy capacity. This is the Eidetic image of Hawking's negative-energy infall: the apparent negative flux into the hole is the past-directed orientation of a positive-energy backward branch, not a violation of positivity. (The dynamical evaporation law $A(t)$ that turns this bookkeeping into horizon shrinkage is deferred to §16; only the count-parametric structure is needed for the Page theorem.)

4.3 Sectors, fixed-sector purity, and the induced radiation channel

The raw tensor product is not the physical Hilbert space. Let $q = (E_{\mathrm{ADM}}, Q, J, z, \dots)$ label asymptotic energy, charges, angular momentum, center data, and fixed-sector labels, with projection $P_{q,N} : \mathcal{H}_N \to \mathcal{H}_N$, $\mathcal{H}^q_{\mathrm{phys},N} = P_{q,N}\mathcal{H}_N$, and $\mathcal{A}^q_{\mathrm{phys},N} = B(\mathcal{H}^q_{\mathrm{phys},N})$.

Proposition 2 (Derived fixed-sector purity and induced one-branch channel)**.** *Suppose the global Eide network is in a pure pair-network state $|\Psi\rangle = \bigotimes_\alpha |\Omega_\alpha\rangle$ (or a fixed-sector projection thereof), and let evaporation be represented by a code-isometric repartitioning of pairs between an outgoing (future-branch) ledger and a complement (backward-branch) ledger. Then within each fixed sector q the* global *evolution on $\mathcal{H}^q_{\mathrm{phys},N} \cong \mathcal{H}_R(t) \otimes \mathcal{H}_C(t)$ is implemented by an isometry U_q (unitary in the no-remnant case), the global state stays pure,*

$$|\Psi(t)\rangle = U_q(t)|\Psi(0)\rangle \in \mathcal{H}_R(t) \otimes \mathcal{H}_C(t), \qquad \rho^q_{\mathrm{phys}}(t) = |\Psi(t)\rangle\langle\Psi(t)|,$$

and the radiation observer sees the induced one-branch channel

$$\rho_R(t) = \mathrm{Tr}_C\, |\Psi(t)\rangle\langle\Psi(t)|,$$

a CPTP map—not a unitary on $\mathcal{H}_R$ alone.

Proof. Global purity is built in by (1). A code-isometric repartitioning $W_k : \mathcal{H}_{\mathrm{BH}}(0) \to \mathcal{H}_R(k) \otimes \mathcal{H}_C(k)$ with $W_k^\dagger W_k = \mathbb{I}$ preserves purity and, when $\dim \mathcal{H}_{\mathrm{BH}}(0) = d_R(k)d_C(k)$ (no-remnant case), is unitary onto the displayed factor; otherwise it is an isometric embedding into the fixed sector. Because the fixed-sector projector commutes with the conserved data, the induced map on $\mathcal{H}^q_{\mathrm{phys},N}$

is the compression of a unitary to an invariant subspace, hence unitary (isometric) there, which establishes global purity. Restriction to the radiation algebra $\mathcal{A}_R(t)$ is then the partial trace over $\mathcal{H}_C(t)$, a CPTP channel; it is not unitary on $\mathcal{H}_R$ alone, which is exactly why $S_R(t)$ can rise and then fall along the Page curve. □

Remark 1. What was *Axiom 2 (Eidetic unitarity)* in the earlier formulation is now Proposition 2: the *global* fixed-sector unitarity/isometry, and the resulting purity, are consequences of bidirectional purity plus isometric repartitioning rather than an independent postulate. The radiation observer's evolution is the induced CPTP channel, never a one-branch unitary; this distinction is what the Page theorem requires. Slow evaporation is an adiabatic sequence of neighboring fixed sectors; the Page theorem is applied inside each center-conditioned block.

4.4 Isometric singularity return is the backward branch

Axiom 2 (Eidetic singularity duality)**.** The Big Bang and black-hole singularities are opposite transitions of the same boundary/bulk reconstruction:

$$\text{Big Bang}: \ \text{Eide} \to \text{bulk}, \qquad \text{black-hole singularity}: \ \text{bulk} \to \text{Eide}.$$

The Big Bang is the first global bulkization of pair-network entropy-energy; it starts the bulk in a low-rank pure state and so *supplies the purity premise* that Proposition 2 and the Page theorem later consume (§13). A black-hole singularity is the inverse event: faithful bulk-local representation of the past-directed branch terminates and its information is re-encoded into the Eide complement. Concretely, a returning quantum a is represented by an isometry

$$V_{\mathrm{ret},a} : \mathcal{H}^{(a)}_{\mathrm{bulk}} \to \mathcal{H}^{(a)}_{\mathrm{Eide},C}, \qquad V^{\dagger}_{\mathrm{ret},a} V_{\mathrm{ret},a} = \mathbb{I}. \tag{2}$$

With χ the invariant parameter (proper time, affine parameter, or modular parameter), and faithful bulk support for $\chi < \chi^{\star}_a$ only, the singular return event is $s_a = \chi^{\star}_a = \sup\{\chi : \mathrm{supp}_{\mathrm{bulk}}(a,\chi) \neq \emptyset\}$, and the singularity is the collection $S_{\mathrm{BH}} = \{s_a : a \in L_{\mathrm{BH}}\}$ of bulk-to-Eide transition events—not an information-destroying point.

Lemma 1 (Return is entropy preserving)**.** *If $V : \mathcal{H}_1 \to \mathcal{H}_2$ is an isometry and ρ a density matrix on $\mathcal{H}_1$, then $S(V\rho V^{\dagger}) = S(\rho)$.*

Proof. $V^{\dagger}V = \mathbb{I}_{\mathcal{H}_1}$ makes V a unitary identification of $\mathcal{H}_1$ with $V(\mathcal{H}_1) \subseteq \mathcal{H}_2$; the nonzero spectrum is unchanged and added eigenvalues are zero, with $-x \ln x \to 0$. □

A deletion singularity would instead implement $\rho \mapsto (\mathrm{Tr}\,\rho)\rho_0$. Here V_{ret} is precisely the backward branch read in the bulk-to-Eide direction, so the entropy-preserving form (2) is forced by the pair structure rather than chosen.

4.5 The benign-interior condition, split and partly discharged

The bidirectional realization of the final state inherits one Horowitz–Maldacena obligation. The earlier formulation bundled this into a single "benign reconstruction" conjecture. We now separate it into two claims of *different logical status*, and show that one of them follows from the same approximate 2-design that drives the Page curve, while the other—the genuinely hard one—remains conjectural.

Definition 1 (Two parts of benign reconstruction)**.** *(1a) No teleportation corruption (recovery fidelity).* The infalling information is faithfully recoverable from the radiation: the channel $\mathcal{N}_\omega$ from an infalling message A to the radiation D induced by R_ω and the final-state structure has entanglement fidelity $\geq 1-\varepsilon$. *(1b) Horizon smoothness.* The infalling observer encounters vacuum at the horizon (no firewall): with $\widetilde{B}$ the complement reconstruction used by the infalling observer, $\|\rho_{B\widetilde{B}} - \rho_{\rm vac}\|_1 \leq \epsilon_{\rm hor}$.

Claim (1a) is a *decoupling* statement, hence a second-moment quantity, and is controlled by the very design property already established in §10. Claim (1b) is not a decoupling statement; it is the locus of the AMPS tension and is not settled by scrambling. We treat them separately.

Proposition 3 (Recovery fidelity from the design)**.** *Let an infalling message A of n_A qubits be maximally entangled with a reference, thrown into a post-Page Eide code, and let the interior be scrambled by an SVP ensemble that is an ε_2-approximate 2-design (Theorem 3). Let D be the radiation collected after the message, of n_D qubits. Then the message decouples from the remaining interior C,*

$$\big\|\rho_{\mathrm{Ref},C} - \rho_{\mathrm{Ref}} \otimes \rho_C\big\|_1 \;\leq\; \sqrt{\tfrac{d_A\, d_C}{d_D}} \;+\; O(\varepsilon_2) \;=\; O\big(2^{-(n_D-n_A)/2}\big) + O(\varepsilon_2),$$

and consequently (Petz recovery) the message is reconstructible from D together with the early radiation with entanglement fidelity $\geq 1 - O(2^{-(n_D-n_A)/2}) - O(\varepsilon_2)$. In particular part (1a) *of Definition 1* follows *from the same approximate 2-design used for the Page curve, with no independent assumption, once the radiation exceeds the message by a few qubits.*

Proof. This is the Hayden–Preskill recovery statement [4] together with the decoupling theorem [5], whose expectation bound depends only on the second moment of the scrambling ensemble; an ε_2-approximate 2-design reproduces that second moment up to ε_2. The displayed trace-norm bound is the standard decoupling estimate with $d_A = 2^{n_A}$, $d_D = 2^{n_D}$; decoupling of the reference from C is equivalent, by Uhlmann/Petz, to existence of a recovery channel acting on D (and the early radiation purifying the old hole) with the stated fidelity. □

Remark 2 (Numerical demonstration of (1a), and that it has content). On the explicit toy R_ω of §9.1 (here $n_A = 1$ message qubit, a 4-qubit post-Page hole, 10 qubits total), a Haar/design interior scramble drives the decoupling trace norm from its maximal value 1.5 at $n_D = 0$ monotonically to 0, while the recovered information $I(\mathrm{Ref} : D, R_e)$ rises to its maximum $2n_A \ln 2 = 1.39$ nats once n_D exceeds n_A by two qubits—faithful recovery, as Proposition 3 predicts. A finite-depth local circuit gives the same collapse at modest depth. Crucially, a *non-benign control*—an interior map that leaves the message unscrambled in the complement—keeps the decoupling norm pinned at 1.5 and the recovered information at 0 until the trivial endpoint: recovery *fails*. The condition therefore has real content; the design satisfies it and a corrupting reconstruction map violates it. Figure 1 summarizes; the script is supplementary (`eide_decoupling.py`).

Conjecture 1 (Horizon smoothness — part (1b), unresolved)**.** The reconstruction map R_ω additionally yields a smooth horizon: in operator terms, the interior modular flow induced by R_ω commutes, up to $O(1)$ corrections in the design accuracy, with the projection onto the pair Bell state (1), so that $\|\rho_{B\widetilde{B}} - \rho_{\rm vac}\|_1 \leq \epsilon_{\rm hor}$.

We emphasize the asymmetry. Part (1a) is now a *proposition* (Proposition 3), discharged from the design; the outgoing radiation faithfully carries the infalling information. Part (1b) remains a

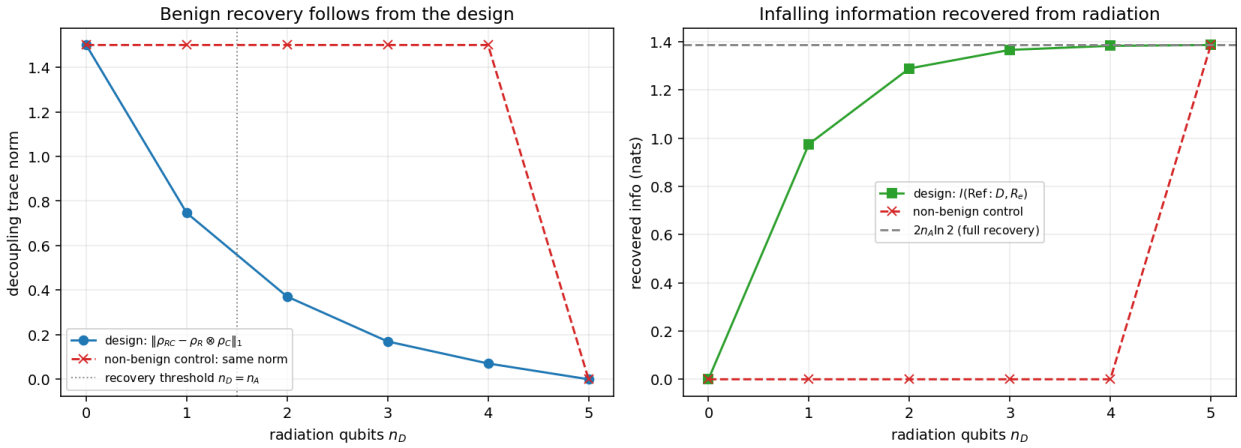

Figure 1: Benign *recovery* (part 1a) on the toy R_ω. *Left:* the decoupling trace norm $\|\rho_{\mathrm{Ref},C} - \rho_{\mathrm{Ref}} \otimes \rho_C\|_1$ collapses under a design scramble (blue) as the radiation grows past the message, but stays pinned at its maximum for a non-benign control map that hides the message in the complement (red). *Right:* the information recovered from the radiation saturates to $2n_A \ln 2$ (full recovery) under the design (green) and never recovers under the control (red). This illustrates Proposition 3: recovery fidelity follows from the same 2-design as the Page curve. It does *not* address horizon smoothness (part 1b).

genuine *conjecture*: decoupling and recovery do not imply that the infalling observer sees vacuum—this is precisely the content of the AMPS argument, and scrambling alone does not settle it. If Conjecture 1 fails, the outgoing radiation is still unitary and still faithfully decodable (1a holds regardless); only the infalling observer's smoothness is at stake, exactly as in the firewall discussion (§15). The Page theorem below uses only purity and the design moment and is unaffected by either part.

5 Null propagation and the SVP interaction graph

The Page curve does not follow from unitarity alone: a pure product state in $\mathcal{H}_R \otimes \mathcal{H}_C$ has zero radiation entropy. A mixing (scrambling) statement is required. In the earlier formulation this entered as a postulated local, ergodic, near-expander interaction graph—which, since scrambling is the entire engine of the curve, came close to assuming the conclusion. Here null propagation *fixes the candidate local interaction graph*—its vertices, edges, locality, and connectivity—while the expansion (spectral-gap) property that the polynomial-depth design bound requires remains a separate Eidetic graph hypothesis (Conjecture 2). We are careful not to claim that null focusing by itself supplies scrambling; it supplies the graph on which scrambling acts, and the rate of scrambling is governed by the expansion hypothesis stated below.

5.1 Light-speed branch motion as an intersection rule

The model states that one member of each pair propagates at the speed of light forward in time and the other backward. Since no bulk is primitive, this is encoded as an oriented branch-flow automorphism $\alpha_{\alpha,\sigma}(\lambda) : \mathcal{A}_{\alpha,\sigma} \to \mathcal{A}_{\alpha,\sigma}$, becoming null propagation only after reconstruction:

$$R_\omega(\alpha_{\alpha,+}) = \text{future-directed null branch}, \qquad R_\omega(\alpha_{\alpha,-}) = \text{past-directed null branch},$$

i.e. $g_\omega(\dot{x}_{\alpha,\sigma}, \dot{x}_{\alpha,\sigma}) = 0$ and $\mathrm{sgn}(dt_\omega/d\lambda) = \sigma$.

Definition 2 (Null intersection graph)**.** Two active Eide Spheres (α, σ) and (β, σ') are *SVP-adjacent* at parameter window t iff their reconstructed null sheets intersect within the window:

$$(\alpha, \sigma) \sim (\beta, \sigma') \iff R_\omega(\Sigma_{\alpha,\sigma}) \cap R_\omega(\Sigma_{\beta,\sigma'}) \neq \emptyset.$$

The active SVP interaction graph $G_{\mathrm{act}}(t) = (V_{\mathrm{act}}(t), E_{\mathrm{act}}(t))$ has the active logical modes as vertices and the SVP-adjacent pairs as edges.

Proposition 4 (Locality and connectivity of G_{act})**.** *In a smooth semiclassical black-hole sector, $G_{\mathrm{act}}(t)$ is (i)* local *in the emergent metric—edges connect only Eide Spheres whose null sheets meet, hence whose reconstructed supports are causally adjacent; and (ii)* connected *within the black-hole code sector, because all infalling and near-horizon null sheets are focused by the Raychaudhuri convergence of the horizon-generating congruence into a common intersection region.*

Proof sketch. (i) Null adjacency is by construction a relation between causally neighboring reconstructed supports, so each gate $U_{(u,v)}$ acts on metric-local data; this is the Eidetic origin of SVP *locality*. (ii) For a trapped region the expansion $\theta < 0$ along the horizon generators (Raychaudhuri, $d\theta/d\lambda = -\theta^2/2 - \sigma_{ab}\sigma^{ab} - R_{ab}k^a k^b$) drives neighboring null sheets together, so the intersection relation has no isolated component on the horizon cut; the active component is connected. □

Conjecture 2 (Expansion of the focused graph — optional sharpening)**.** In the near-horizon region the focused null-intersection graph lies in the bounded-degree expander class, i.e. the spectral gap of the SVP edge-selection random walk on G_{act} is bounded below by a sector-independent constant $\Omega(1)$, rather than merely by the inverse-polynomial floor $\Omega(1/\mathrm{poly})$ of the polynomial-gap condition (Definition 6) or the bare positivity > 0 from connectivity alone (Lemma 4).

Proposition 4 *derives* locality and connectivity, which were previously postulated. We emphasize that Conjecture 2 is *not* required for the Page curve: §10.1 shows that connectivity gives a positive gap (a design in some depth) and that the focused-graph polynomial-gap condition upgrades this to a polynomial-depth approximate 2-design, which forms within the evaporation lifetime under a stated gate-rate model. Expansion only sharpens the design depth from polynomial to polylogarithmic—equivalently, it is the statement that the Eide horizon *scrambles fast*, in the Sekino–Susskind sense $t_\star \sim \beta \ln S$. It is therefore relegated to a consistency bonus with the fast-scrambling conjecture, not a hypothesis of the main theorem. This is why the conjecture is labeled "optional sharpening."

5.2 The SVP design axiom, now graph-derived

Axiom 3 (Finite-network SVP locality and universality)**.** After fixed charge, gauge, and center data are quotiented, SVP dynamics in each emitted-count block is an ensemble of local circuits on the null-intersection graph $G_{\mathrm{act}}(t)$ of Definition 2, with: (1) each elementary gate acting on one edge and preserving (q, z); (2) the two-site gate distribution universal, or absolutely continuous w.r.t. Haar, on each two-site block; (3) the edge-selection rule ergodic and uniformly comparable to the simple random walk on G_{act}: there is a constant $c_{\mathrm{sel}} > 0$, independent of the sector, with selection probability $p_e \geq c_{\mathrm{sel}}/|E_{\mathrm{act}}|$ for every active edge e); (4) connectivity in the black-hole sector (Proposition 4(ii)). No expansion or sector-independent gap is assumed; by Lemma 4 the second-moment gap is automatically positive once (3)–(4) hold.

A depth-L SVP circuit is $U_{\mathrm{SVP}}^{(L)} = U_L \cdots U_1$ with law $\mu_{\mathrm{SVP},L}$. The word *ensemble* matters: a single fixed unitary is not a design. The design object is the distribution over microscopic Eide phase/gate choices compatible with the same macroscopic sector; a deterministic chaotic realization is a typical member.

6 Eide-Sphere horizon counting and the area density

The geometric Page curve needs the complement dimension as a function of horizon area. Crucially, the transfer matrix and its eigenvalue are now *computed from flux quantization* (§3), not postulated.

Definition 3 (Eide horizon area cells)**.** Fix a semiclassical sector and a horizon cut $\Gamma_H(A)$ of area A. The cut is represented by N_A active Eide-Sphere area cells, and there is a primitive area quantum $a_\star > 0$ with $N_A = A/a_\star + O(1)$ in a microcanonical window of width $O(a_\star)$.

Definition 4 (Admissible horizon configurations)**.** Let X be the finite alphabet of locally allowed Eide horizon-cell labels after fixing macroscopic data q, z. We use the term *word* as shorthand for an admissible horizon-cell configuration $w = (x_1, \ldots, x_N) \in X^N$ satisfying the local adjacency, branch-orientation, and *flux-closure* compatibility on the cut. Let $W_N(q, z)$ be the admissible set (including fixed center data, excluding gauge duplicates), and

$$\mathcal{H}_C(A) := \mathrm{span}\{|w\rangle : w \in W_{N_A}(q, z)\}, \qquad \dim \mathcal{H}_C(A) = |W_{N_A}(q, z)|.$$

Remark 3 (Two-dimensional horizon and the transfer operator)*.* A black-hole horizon cut is two-dimensional, so the linear "word" picture is a simplified presentation. For a genuine 2D cut, $T_{q,z}$ below should be read as the row-to-row *transfer operator* of the admissible finite-state surface code (equivalently, a tensor-network transfer operator across rows of the admissible 2D cell complex), and $\lambda_\star$ is the topological entropy per horizon cell of that operator. We make explicit the finite-size condition under which the subleading correction is logarithmic:

> *Polynomial finite-size degeneracy.* The fixed-center, closed-horizon surface code has topological entropy $\ln \lambda_\star$ per cell and only polynomial finite-size degeneracy, so $|W_N| = \lambda_\star^N \, N^{O(1)}$.

Under this condition the leading area law and the $O(\ln A)$ correction of Theorem 1 hold as stated. If the finite-size condition fails—for instance if boundary or shape contributions of a particular 2D geometry exceed $\ln N$—then the *leading* area law $\eta_{\mathrm{Eide}} A$ survives, but the subleading correction need not be $O(\ln A)$. Establishing the finite-size condition for the genuine Eide horizon code is a theorem target (§18), not an automatic consequence of the 2D reading.

Proposition 5 (Minimal flux-cell transfer matrix)**.** *Two claims, with distinct logical status.*

(a) (Two-state alphabet, from quantization.) *By Proposition 1 the minimal flux sector $k = 1$ gives a two-state local cell alphabet, $|X| = \dim \mathcal{H}^{(1)} = 2$. More generally $|X| = k + 1$.*

(b) (Minimal unconstrained-cell hypothesis.) *In the minimal qubit horizon code, after fixed global charge, closure, and center data are quotiented, the local flux-closure and branch-orientation rules impose* no additional nearest-neighbor exclusion *on the cut. Under this minimal-code condition the transfer operator (Remark 3) is the all-ones 2×2 adjacency, $T_{q,z} = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, and its Perron–Frobenius/topological-entropy eigenvalue is*

$$\lambda_\star = \rho(T_{q,z}) = 2.$$

In either reading, fixed global charge, closure, and center constraints impose only subextensive restrictions: there exist $c_\pm > 0$, $p_\pm \geq 0$ independent of N with

$$c_- N^{-p_-} \lambda_\star^N \leq |W_N(q,z)| \leq c_+ N^{p_+} \lambda_\star^N. \tag{3}$$

Status of the two claims. Claim (a) is a direct consequence of geometric quantization (Borel–Weil), $\dim H^0(\mathbb{CP}^1, \mathcal{O}(k)) = k+1$, and is unconditional. Claim (b) is a separate *minimal-code hypothesis*: it holds when no admissibility rule beyond the subextensive global constraints removes a nearest-neighbor pair, so that every cell label is compatible with every neighbor and $T_{q,z}$ is all-ones with $\rho(T_{q,z}) = 2$. The global charge/closure/center constraints then remove or add only subextensively many configurations, giving the polynomial prefactors in (3). If local flux-closure or branch-orientation rules *did* exclude some adjacencies, the admissibility graph would be a proper subgraph of the complete graph and one would have $1 < \lambda_\star \leq 2$; the value $\lambda_\star = 2$ is therefore a prediction of the minimal unconstrained flux-cell model, not a universal consequence of quantization alone. For general flux k the analogous minimal model gives $\lambda_\star = k+1$. □

Remark 4. This is the decisive Eidetic upgrade of the old counting axiom, now stated with its correct logical status. Previously $\lambda_\star > 1$ was *assumed* outright; here the *two-state alphabet* (claim (a)) is derived from S^2 flux quantization unconditionally, and the specific value $\lambda_\star = 2$ follows under the minimal unconstrained-cell hypothesis (claim (b)). The value of $\lambda_\star$, and hence the area quantum, is thus a falsifiable prediction of which flux sector is fundamental together with the minimal-code condition—the Eidetic analogue of the loop-quantum-gravity area spectrum—rather than a free parameter or a guaranteed output of quantization in isolation. Section 6.1 reduces the minimal-code condition (b) to the independence Axiom 1 plus a provably subextensive global closure, leaving only an R_ω flux-matching possibility as the genuine residual.

Theorem 1 (Eide-Sphere area-density theorem)**.** *Under Proposition 5,*

$$\ln \dim \mathcal{H}_C(A) = \eta_{\text{Eide}} A + O(\ln A), \qquad \eta_{\text{Eide}} := \frac{\ln \lambda_\star}{a_\star} = \lim_{A\to\infty} \frac{1}{A} \ln \dim \mathcal{H}_C(A). \tag{4}$$

Proof. Take logs of (3): $\ln \dim \mathcal{H}_C(A) = N_A \ln \lambda_\star + O(\ln N_A)$. Insert $N_A = A/a_\star + O(1)$ to get $\ln \dim \mathcal{H}_C(A) = (A/a_\star) \ln \lambda_\star + O(\ln A)$. Divide by A and send $A \to \infty$. □

For minimal qubit cells, $\lambda_\star = 2$ and $\eta_{\text{Eide}} = \ln 2/a_\star$, $\dim \mathcal{H}_C(A) = 2^{N_A} \text{poly}(N_A)$.

6.1 Reducing the minimal-cell hypothesis to independence plus global closure

Proposition 5(b) assumed "no local nearest-neighbor exclusion." We now reduce this hypothesis to the *minimal independent-cell realization* of Axiom 1 (the natural and simplest reading of the boundary axiom, not forced by it) together with a global flux-closure constraint, and we prove rigorously that the global constraint is *subextensive*—it cannot change $\lambda_\star$, only the polynomial prefactor.

Proposition 6 (Minimal independent-cell realization of Axiom 1)**.** *Axiom 1 furnishes a pre-bulk net of independent quantized spheres, each with site Hilbert space $\mathcal{H}_{\alpha,\sigma} \cong \mathbb{C}^2$ ($k = 1$). Consider the* minimal independent-cell realization*, in which the physical horizon-cut code after quotienting fixed center data inherits this primitive independence with no local adjacency restriction. In this realization the horizon-cut Hilbert space is the tensor product $\bigotimes_{i=1}^{N_A} \mathcal{H}_i$, and the only admissibility conditions are the global ones inherited from the fixed sector (q,z).*

Remark 5. We are careful here: Axiom 1 gives a pre-bulk network of paired spheres, but it does not by itself prove that the *physical horizon-cell code*, after fixed-sector constraints and reconstruction, has no local adjacency restriction. The minimal independent-cell realization is the natural and simplest reading of the axiom, and is the hypothesis under which $\lambda_\star = 2$ below; the place a local constraint could re-enter is the reconstruction map R_ω (effective flux-continuity across cells), which is a property of R_ω distinct from the recovery and smoothness parts of §4.5 and is the honest residual flagged in §18.

The remaining admissibility condition is the global flux closure: the total quantized flux threading the closed horizon cut is fixed by the macroscopic data. Writing the per-cell flux label as $q_i \in \{0, 1\}$ (equivalently a spin $s_i = q_i - \frac{1}{2} \in \{\pm\frac{1}{2}\}$), the closure reads $\sum_{i=1}^{N} q_i = Q$ with Q fixed by the sector.

Lemma 2 (Global closure is subextensive: balanced sector)**.** *In the balanced (neutral, non-rotating) sector the closure is centered, $Q = N/2$, and*

$$|W_N(q, z)| = \binom{N}{N/2} = 2^N \sqrt{\tfrac{2}{\pi N}} \left(1 + O(1/N)\right).$$

Hence $\lambda_\star = 2$ exactly, and the bound (3) holds with $p_- = p_+ = \frac{1}{2}$ and explicit constants. The single global constraint reduces the naive count 2^N only by the polynomial factor $N^{-1/2}$; it does not change the exponential growth rate.

Proof. With independent cells (Prop. 6) and the single linear closure $\sum_i q_i = N/2$, the admissible words are exactly the length-N binary strings of Hamming weight $N/2$, of which there are $\binom{N}{N/2}$. Stirling's formula gives the displayed asymptotics, so $\ln |W_N| = N \ln 2 - \frac{1}{2} \ln N + O(1)$, i.e. $\lambda_\star = 2$ with logarithmic (in dim) correction, matching (3). □

Remark 6 (Charged/rotating sectors are deferred). If the macroscopic charge or angular momentum is *extensive*, the closure shifts to $Q = \nu N$ with $\nu \neq \frac{1}{2}$, and the naive balanced count would give $\binom{N}{\nu N} \sim e^{NH(\nu)}$ with $H(\nu) = -\nu \ln \nu - (1 - \nu) \ln(1 - \nu) < \ln 2$, suggesting $\lambda_\star = e^{H(\nu)} < 2$. This is *not* a contradiction: at fixed charge the relevant area $A(M, Q, J)$ and the sector-restricted count co-vary, and a correct treatment counts microstates of the *charged* transfer operator at the charged horizon area. We restrict the rigorous claim $\lambda_\star = 2$ to the balanced Schwarzschild sector —which is precisely the canonical Page-curve setting—and defer the charged/rotating transfer operator to §18.

Corollary 1 (Status of $\lambda_\star = 2$ after reduction)**.** *In the balanced sector, $\lambda_\star = 2$ now rests on: (i) $k = 1$ flux quantization [Borel–Weil, unconditional]; (ii) the minimal independent-cell realization of Axiom 1 [Prop. 6, the natural reading of the boundary axiom, not forced by it]; and (iii) the global flux closure [provably subextensive, Lemma 2]. The residual is the possibility that R_ω imposes an effective local flux-matching, which would replace the complete-graph alphabet by a subshift of finite type and lower $\lambda_\star$; this is a distinct property of R_ω (neither the recovery-fidelity part nor the smoothness part of §4.5), and excluding it requires constructing R_ω explicitly enough to read off the induced cell adjacency.*

7 Quarter-area matching from emergent Einstein response

To fix the coefficient 1/4 the Newton coupling must be defined independently of entropy. This section supplies that matching; it is, by design, the Eidetic transcription of Jacobson's thermodynamic derivation [16] and is presented as a *consistency theorem*, not as a first-principles fixing of

1/4 from nothing. What *is* new is that the microscopic density η_{Eide} entering the matching is now the derived $\ln 2/a_\star$ of Theorem 1, so the matching outputs a *predicted* area quantum.

Axiom 4 (Eidetic semiclassical gravitational response)**.** In large smooth sectors the Eide reconstruction has a low-energy effective action whose leading diffeomorphism-invariant metric term is Einstein–Hilbert,

$$\Gamma_{\text{Eide}}[g,\Phi] = \frac{1}{16\pi G_{\text{Eide}}}\int_M d^4x\sqrt{-g}\,R + \Gamma_{\text{matter}}[g,\Phi] + \Gamma_{\text{sub}}[g,\Phi], \tag{5}$$

with G_{Eide} operationally defined by the linearized gravitational response to stress-energy; Γ_{sub} gives only lower-order corrections to the leading area density.

Axiom 5 (Eidetic local horizon thermodynamics)**.** For every quasistatic local horizon cut in a semiclassical reconstruction: (i) the Eide modular flow of a branch-restricted horizon algebra reconstructs the boost flow with temperature $T_\kappa = \hbar\kappa/2\pi$ (the Eide reconstruction of Bisognano–Wichmann modular flow [17]); (ii) the entanglement first law $\delta S_{\text{Eide}} = \delta Q/T_\kappa$ holds for reversible perturbations; (iii) the Einstein–Hilbert response gives $\delta Q = (\kappa/8\pi G_{\text{Eide}})\,\delta A$ at leading order.

Theorem 2 (Quarter-area matching)**.** *Under Theorem 1 and Axioms 4–5,*

$$\eta_{\text{Eide}} = \frac{1}{4G_{\text{Eide}}\hbar}, \qquad \ell^2_{\text{Pl,phys}} := G_{\text{Eide}}\hbar\ (c=1), \qquad \ln\dim\mathcal{H}_C(A) = \frac{A}{4\ell^2_{\text{Pl,phys}}} + O(\ln A). \tag{6}$$

Equivalently the microscopic cell parameters obey the matching condition

$$a_\star = 4G_{\text{Eide}}\hbar\,\ln\lambda_\star = 4\ell^2_{\text{Pl,phys}}\ln\lambda_\star. \tag{7}$$

Proof. Axiom 5 gives $\delta S_{\text{Eide}} = \delta Q/T_\kappa = (\kappa\,\delta A/8\pi G_{\text{Eide}})/(\hbar\kappa/2\pi) = \delta A/4G_{\text{Eide}}\hbar$, so $\delta S_{\text{Eide}}/\delta A = 1/4G_{\text{Eide}}\hbar$. Theorem 1 gives $S_{\text{Eide}}(A) = \ln\dim\mathcal{H}_C(A) = \eta_{\text{Eide}}A + O(\ln A)$, with leading density η_{Eide}. Equating densities gives (6); since $\eta_{\text{Eide}} = \ln\lambda_\star/a_\star$, (7) follows. □

Corollary 2 (Predicted area quantum)**.** *For minimal qubit cells $\lambda_\star = 2$ (Proposition 5), so*

$$\boxed{a_\star = 4\ell^2_{\text{Pl,phys}}\ln 2} \tag{8}$$

and hence $N_A = A/(4\ell^2_{\text{Pl,phys}}\ln 2)+O(1)$, $\dim\mathcal{H}_C(A) = 2^{N_A}\text{poly}(N_A)$, $\ln\dim\mathcal{H}_C(A) = A/4\ell^2_{\text{Pl,phys}}+O(\ln A)$.

Remark 7 (What is and is not a convention). The 1/4 is fixed by the simultaneous consistency of derived Eide counting, emergent Einstein response, modular temperature, and the first law—exactly the content of Jacobson's argument, here transcribed. The genuinely Eidetic and falsifiable output is not the 1/4 but the *value of* $a_\star$ in (8), because the two-state alphabet is fixed by flux quantization and $\lambda_\star = 2$ then follows in the balanced minimal independent-cell model (§6.1). A different fundamental flux sector, or an R_ω-induced local flux-matching, would give a different $\lambda_\star$ and hence a different, testable area spectrum (e.g. $a_\star = 4\ell^2_{\text{Pl,phys}}\ln(k+1)$ for sector k).

8 Radiation algebra and Eide-Sphere code complement

At evaporation stage t define the radiation algebra $\mathcal{A}_R(t) \subset \mathcal{A}^q_{\mathrm{phys},N}$ as the subalgebra generated by outgoing +-oriented (future-branch) EEQ branches whose reconstructed images are accessible Hawking radiation. After conditioning on center data z, use the type-I presentation

$$\mathcal{H}^q_{\mathrm{phys},N}(t) \cong \mathcal{H}_R(t) \otimes \mathcal{H}_C(t), \qquad \mathcal{A}_R(t) = B(\mathcal{H}_R(t)) \otimes \mathbb{I}_C, \tag{9}$$

with Eide-Sphere code complement

$$\mathcal{A}^q_C(t) = \mathcal{A}_R(t)' \cap \mathcal{A}^q_{\mathrm{phys},N} = \mathbb{I}_R \otimes B(\mathcal{H}_C(t)). \tag{10}$$

The complement contains the matched backward EEQ branch algebra, the return register, and the edge/soft correlations required by constraints: $\mathcal{A}^q_C(t) \supset \mathcal{A}^q_-(t) \vee \mathcal{A}^q_{\mathrm{ret}}(t) \vee \mathcal{A}^q_{\mathrm{edge/soft}}(t)$.

Remark 8 (Unique purifier is the backward branch). The purifier of the radiation is not an independent hidden system. It is the complement $\mathcal{A}^q_C(t)$, whose generators are *literally the past-directed branches* of the emitted pairs. Its reconstructed image may look like an interior partner, island, horizon wedge, or static-patch degree of freedom, but these are all reconstructions of the same backward-branch algebra. This is the Eidetic ER = EPR statement that removes the AMPS duplicate-purifier obstruction algebraically; horizon smoothness remains a separate condition (§15).

A radiation observer sees only $\mathcal{A}_R(t)$; the observer state is $\omega_R(t)(X) = \mathrm{Tr}(\rho^q_{\mathrm{phys}}(t)X)$, $X \in \mathcal{A}_R(t)$, represented by $\rho_R(t) = \mathrm{Tr}_C\, \rho^q_{\mathrm{phys}}(t)$. The partial trace is the type-I representative of restriction to $\mathcal{A}_R(t)$; it is not a physical operation that destroys $\mathcal{A}^q_C(t)$. If the center is not fixed, $\mathcal{H}^q_{\mathrm{phys}}(t) \cong \bigoplus_z \mathcal{H}_{R,z}(t) \otimes \mathcal{H}_{C,z}(t)$ and the Page theorem is proved in a fixed z block, then averaged (Corollary 6).

9 Bidirectional EEQ emission and logical budgets

Each horizon-scale emission a creates a bidirectional logical excitation $E^{(+)}_a \otimes E^{(-)}_a$ with $E^{(+)}_a \in \mathcal{A}_R(t)$ and $E^{(-)}_a \in \mathcal{A}^q_C(t)$. The emission update is code-isometric, $V_{\mathrm{em}}(t) : \mathcal{H}^q_{\mathrm{loc}}(t) \to \mathcal{H}^q_{\mathrm{phys},N}(t+\delta t)$, $V^\dagger_{\mathrm{em}} V_{\mathrm{em}} = \mathbb{I}$, and after k emissions the history is an isometry

$$W_k : \mathcal{H}_{\mathrm{BH}}(0) \to \mathcal{H}_R(k) \otimes \mathcal{H}_C(k), \qquad W^\dagger_k W_k = \mathbb{I}. \tag{11}$$

In the no-remnant, fixed-total-code case $\dim \mathcal{H}_{\mathrm{BH}}(0) = d_R(k) d_C(k)$ and W_k is unitary onto the displayed product.

Uniform logical qubits. With initial logical qubit count n_{tot} and emitted count k,

$$n_R(k) = k, \quad n_C(k) = n_{\mathrm{tot}} - k, \quad n_R(k) + n_C(k) = n_{\mathrm{tot}},$$

$$d_R(k) = 2^{n_R(k)} = 2^k, \qquad d_C(k) = 2^{n_C(k)} = 2^{n_{\mathrm{tot}}-k}.$$

A single emission repartitions $\mathcal{H}_C(k) \cong \mathcal{H}^{\mathrm{out}}_{a,+} \otimes \mathcal{H}_C(k+1)$ with $\dim \mathcal{H}^{\mathrm{out}}_{a,+} = 2$, i.e. evaporation transfers one logical Eide degree from the backward-branch ledger to the radiation ledger while the full code stays pure. (These are dimension statements, not entropy statements.)

9.1 What bidirectional pairing does and does not give: the single-reservoir condition

It is important to be precise about how much of the Page curve the bidirectional pair structure supplies on its own, because a tempting but incorrect reading is that pairing alone produces the curve. It does not. The repartitioning above already encodes a genuine additional condition—that the backward-branch purifier is a *finite, shared reservoir* from which emission *transfers* degrees—and that condition, not bare pairing, is what produces the Page turnover.

Lemma 3 (Bare-pairing dichotomy)**.** *Consider n_{tot} Eide pairs in the Bell-network state $\bigotimes_\alpha |\Omega_\alpha\rangle$, with radiation at stage k taken to be the emitted future branches $\{f_1, \ldots, f_k\}$.*

(i) (Private-partner structure.) *If each emitted future branch is purified by its* own *backward partner p_α, which enters the complement with it, then $\rho_R(k) = (\frac{1}{2}\mathbb{I})^{\otimes k}$ and*

$$S_R(k) = k \ln 2 \qquad \textit{for all } k,$$

a monotonically rising line with no turnover—the Hawking entropy, not the Page curve.

(ii) (Single-reservoir structure.) *If instead the backward branches form a* fixed *reservoir of n_{tot} purifying degrees, so that emission transfers one degree from the complement ledger to the radiation ledger (no-remnant: $d_R(k)d_C(k) = 2^{n_{\mathrm{tot}}}$, $n_C(k) = n_{\mathrm{tot}} - k$), and the SVP design scrambles the fixed code, then the Page envelope is saturated and*

$$S_R(k) = \min\{k,\, n_{\mathrm{tot}} - k\} \ln 2 + O(1).$$

Proof. (i) Each $|\Omega_\alpha\rangle$ is a maximally entangled pair, so the marginal on any set of future branches whose partners are all traced out is the maximally mixed state on those branches; $\rho_R(k) = (\frac{1}{2}\mathbb{I})^{\otimes k}$ and $S_R(k) = k \ln 2$. The complement carries one fresh partner per emission, so $d_C(k) \geq d_R(k)$ for every k, the Schmidt-rank bound $S_R \leq \ln \min\{d_R, d_C\}$ is saturated on the rising side only, and no turnover occurs. (ii) is the content of Theorem 4: the fixed-code purity gives the envelope $S_R \leq \ln d_<$, and the SVP 2-design saturates it to $O(1)$, with $d_< = 2^{\min\{k, n_{\mathrm{tot}}-k\}}$. □

Remark 9 (The turnover is contingent, and the contingency is physical). Lemma 3 isolates exactly what is and is not Eidetic-for-free. The *rising* branch follows from bidirectional pairing alone (one-branch mixedness of a Bell pair). The *falling* branch—the unitary signature of information recovery—requires the single-reservoir condition: the backward-branch purifier must be a finite shared pool that shrinks as it is transferred, not a supply of private partners. If each Hawking mode kept its own partner, the complement would never shrink below the radiation, the rank bound would never bite on the late side, and the entropy would rise monotonically as in Hawking's calculation. The single-reservoir condition is therefore a genuine physical input of the construction (the no-remnant assumption of §9), not a consequence of pairing; we flag it as such in §18.

Remark 10 (Toy-model demonstration). The dichotomy is exhibited explicitly in a finite Eide network. We take $n_{\mathrm{tot}} = 10$ logical modes and an explicit *toy* reconstruction map R_ω, with the radiation/complement target partition *fixed by hand* (the first k modes are declared radiation), and compute $S_R(k)$ exactly by Schmidt decomposition. This is a finite-size sanity check, not a reconstruction of bulk geometry from Eide primitives. Under the single-reservoir structure with a design (Haar) scramble, $S_R(k)$ tracks the exact Page formula to a maximum deviation of 3.7×10^{-3} nats—an explicit $O(1)$ illustration of Theorem 4 with R_ω a concrete map rather than a schematic arrow. Replacing the Haar scramble by a *local random circuit on a connected line graph* (no

expander) drives the half-way entropy from 0.29 nats at depth 1 toward the Page peak 2.97 nats by depth 16, illustrating finite-depth convergence toward Page saturation on a connected local circuit; this is a numerical sanity check, not a substitute for the focused-graph polynomial-gap/design condition of §10.1, which is where the theorem-level statement rests. Under the private-partner structure the same R_ω yields the monotone Hawking line $S_R(k) = k \ln 2$. Finally, at a post-turnover stage the most recent radiation mode shares 1.36 nats of mutual information with the *early* radiation and only 0.02 nats with the shrinking interior (of a maximum $2 \ln 2 = 1.39$), an explicit realization of the shared-purifier (ER = EPR) structure of Remark 8. Figure 2 summarizes. The script producing Figure 2 is included as supplementary material (`eide_romega.py`); the runs use $n_{\text{tot}} = 10$, 20–30 random trials per point, a brickwork line-graph circuit, and a fixed seed for reproducibility.

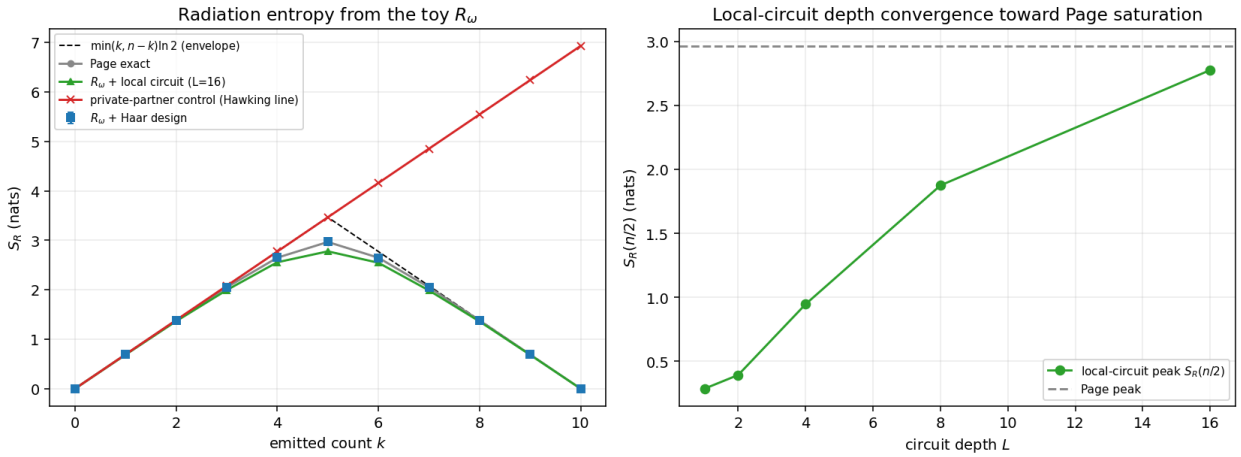

Figure 2: Toy Eide network ($n_{\text{tot}} = 10$); an illustration of the mechanism, not a proof of the graph-design theorem (numerical values in the text). *Left:* the explicit *toy* reconstruction map R_ω (radiation/complement partition fixed by hand) under the single-reservoir structure—with a design scramble (blue) and with a finite-depth local random circuit (green)—reproduces the Page curve (grey, exact Page formula; dashed, the $\min\{k, n-k\} \ln 2$ envelope), while the private-partner control model (red) gives the monotone Hawking line. The red curve is *not* a state in the same fixed single-reservoir code (its complement grows with the radiation), so it lies above the single-reservoir envelope after the midpoint without violating the finite-code bound. *Right:* half-way entropy $S_R(n/2)$ for a finite-depth random circuit on a connected line graph, illustrating depth-driven convergence toward Page saturation in this finite toy model—a numerical sanity check, not a substitute for the focused-graph polynomial-gap/design condition of §10.1 (the line graph is not a fast scrambler).

Microstate-counting area law. At finite cutoff dimensions are integers; with $s_\partial(t) := \ln d_C(t) = \ln \dim \mathcal{H}_C(A(t))$ and Corollary 2,

$$s_\partial(t) = \frac{A(t)}{4\ell^2_{\mathrm{Pl,phys}}} + O\Big(\ln\big(2 + \tfrac{A(t)}{\ell^2_{\mathrm{Pl,phys}}}\big)\Big), \tag{12}$$

$$n_C(t)\ln 2 = \frac{A(t)}{4\ell^2_{\mathrm{Pl,phys}}} + O\Big(\ln\big(2 + \tfrac{A(t)}{\ell^2_{\mathrm{Pl,phys}}}\big)\Big), \tag{13}$$

$$n_R(t)\ln 2 = \frac{A(0) - A(t)}{4\ell^2_{\mathrm{Pl,phys}}} + O\Big(\ln\big(2 + \tfrac{A(0)}{\ell^2_{\mathrm{Pl,phys}}}\big)\Big), \tag{14}$$

the last using the no-remnant relation $\ln d_R(t) = \ln d_{\mathrm{tot}} - \ln d_C(t)$, $\ln d_{\mathrm{tot}} = \ln d_C(0)$.

10 SVP as an approximate unitary design

Definition 5 (Approximate unitary design)**.** For μ on $U(K)$ with r-th moment channel $\mathcal{M}^{(r)}_\mu(X) = \mathbb{E}_{U\sim\mu}[U^{\otimes r} X (U^\dagger)^{\otimes r}]$, μ is an ε-approximate r-design if $\|\mathcal{M}^{(r)}_\mu - \mathcal{M}^{(r)}_{\mathrm{Haar}}\| \le \varepsilon$ in a moment norm controlling all degree-(r,r) observables below. We need $r = 2$.

Theorem 3 (SVP approximate 2-design theorem)**.** *Assume Axiom 3, whose locality and connectivity are supplied by Proposition 4. For every fixed emitted-count block, sector (q,z), and accuracy $\varepsilon > 0$, the depth-L SVP ensemble $\mu_{\mathrm{SVP},L}$ on $K_{q,z,t} \cong \mathcal{H}_R(t) \otimes \mathcal{H}_C(t)$ is an ε-approximate* 2-design *once*

$$L \ge C\,\Delta_2^{-1}\big(n_{\mathrm{act}} \log q + \log(1/\varepsilon)\big),$$

where $\Delta_2 > 0$ is the second-moment gap of the SVP ensemble and C depends only on the local gate family. By Lemma 4, connectivity gives $\Delta_2 > 0$ (a 2-design in some depth), and under the focused-graph polynomial-gap condition (Definition 6) $\Delta_2 \ge \Omega(1/\mathrm{poly}(n_{\mathrm{act}}))$, so $L = \mathrm{poly}(n_{\mathrm{act}}, \log(1/\varepsilon))$ with no expansion assumption.

Proof. Quotienting fixed charge/center data identifies SVP in a block with a local random circuit on the connected (Prop. 4) null-intersection graph with universal gates and ergodic edge selection. The convergence of the *second* moment operator to its Haar value is governed by its spectral gap Δ_2: after L random local gates, $\|\mathcal{M}^{(2)}_{\mu_{\mathrm{SVP},L}} - \mathcal{M}^{(2)}_{\mathrm{Haar}}\| \le (1-\Delta_2)^L \dim(K_{q,z,t})^{O(1)}$, so $L = O(\Delta_2^{-1}(n_{\mathrm{act}} \log q + \log 1/\varepsilon))$ suffices. The qubit nearest-neighbor 2-design is Harrow–Low [8], with the connected/focused-graph gap supplied by Lemma 4 via the second-moment (interchange-process) identity. The fixed-sector constraints only replace the raw tensor product by the logical quotient block; they do not change moment convergence. □

Remark 11 (Why $r = 2$ suffices, and higher designs). The Page theorem uses only the second moment (the radiation purity, Corollary 4), so an approximate 2-design is all that is required, and the focused-graph gap argument is specifically a second-moment statement: Lemma 4 controls Δ_2 via the interchange-process identity, not the gaps of higher moment operators. Higher-r design behavior may follow under the stronger local-random-circuit polynomial-design results of Brandao–Harrow–Horodecki [9], but the present argument neither needs nor claims it.

Remark 12. The fundamental input is no longer "SVP scrambles to Page typicality." It is null propagation ⇒ local connected graph ⇒ approximate design. Page behavior then follows from design moments and finite-dimensional entropy inequalities.

10.1 Discharging the expansion hypothesis: the polynomial-gap condition

The earlier formulation isolated graph expansion (Conjecture 2) as the remaining quantitative input. We now show that the Page curve does not need full expansion; it needs only the weaker *focused-graph polynomial-gap condition*, which we *hypothesize* for the null-intersection graph families of §5 and list as a residual input (§18) rather than prove. Connectivity alone (derived, Prop. 4) gives a *positive* second-moment gap (hence a design in *some* depth); under the polynomial-gap condition this becomes an *inverse-polynomial* gap (hence polynomial depth). Full expansion would further upgrade poly to polylog depth (fast scrambling) and is not required.

Definition 6 (Focused-graph polynomial-gap condition)**.** The active near-horizon graph family $\{G_{\mathrm{act}}\}$ satisfies the *polynomial-gap condition* if the simple-random-walk gap is inverse-polynomial, $\Delta_{\mathrm{RW}}(G_{\mathrm{act}}) \geq 1/\mathrm{poly}(n_{\mathrm{act}})$, uniformly in the sector. This is strictly weaker than expansion ($\Delta_{\mathrm{RW}} = \Omega(1)$) and strictly stronger than bare connectivity ($\Delta_{\mathrm{RW}} > 0$).

Lemma 4 (Connected second-moment gap)**.** *Let G_{act} be a connected graph on $n = n_{\mathrm{act}}$ logical modes, and let the SVP ensemble apply universal (or Haar) two-site gates under the normalized ergodic edge-selection chain of Axiom 3(2)–(3) ($p_e \geq c_{\mathrm{sel}}/|E_{\mathrm{act}}|$). Write $\Delta_{\mathrm{RW}}(G_{\mathrm{act}})$ for the spectral gap of the* simple random walk *on G_{act} (with the same edge weights). Then the second-moment gap obeys*

$$\Delta_2(G_{\mathrm{act}}) \;\geq\; c\, c_{\mathrm{sel}}\, \Delta_{\mathrm{RW}}(G_{\mathrm{act}}) \;>\; 0$$

for a universal constant $c > 0$ (connectivity alone). Under the focused-graph polynomial-gap condition (Definition 6)—hypothesized for the families of §5 on the grounds of bounded local geometry and diameter $O(\mathrm{poly}\, n)$, and listed as a residual input in §18—this sharpens to

$$\Delta_2(G_{\mathrm{act}}) \;\geq\; \frac{1}{\mathrm{poly}(n_{\mathrm{act}})}.$$

No expansion is assumed; expansion (Conjecture 2) is the additional statement $\Delta_2 = \Omega(1)$.

Proof. By the Brandao–Harrow–Horodecki reduction [9], the second-moment operator of a local random circuit is a frustration-free, edge-local positive operator whose gap is lower-bounded by a constant multiple of the gap of the associated symmetric exclusion / interchange process on G_{act}; by the Aldous spectral-gap identity (Caputo–Liggett–Richthammer [10]) the interchange gap equals the simple random-walk gap $\Delta_{\mathrm{RW}}(G_{\mathrm{act}})$. The normalized selection floor $c_{\mathrm{sel}} > 0$ contributes the displayed factor. This gives $\Delta_2 > 0$ for every connected G_{act}.

The quantitative inverse-polynomial bound is *not* universal over all connected graphs (graphs with severe bottlenecks can have Δ_{RW} exponentially small), so we state it only for the relevant family. For the focused near-horizon graphs of §5—which have bounded local valence after coarse-graining and polynomial diameter $D = O(\mathrm{poly}\, n)$—a Poincaré/canonical-path comparison to a spanning path (Diaconis–Saloff-Coste [11]) gives $\Delta_{\mathrm{RW}} \geq 1/(\rho\, D)$ with canonical-path congestion $\rho = O(\mathrm{poly}\, n)$, hence $\Delta_2 \geq 1/\mathrm{poly}(n)$. For an unstructured worst-case connected graph one retains only $\Delta_2 > 0$; the inverse-polynomial rate is a property of graphs of bounded coarse-grained valence and polynomial diameter, which we take as the defining content of the focused-graph polynomial-gap condition (Definition 6) rather than derive from the metric here. □

Corollary 3 (Polynomial-depth design under the polynomial-gap condition)**.** *Under Axiom 3 with G_{act} connected and satisfying the focused-graph polynomial-gap condition (Definition 6), the SVP ensemble is an ε-approximate unitary 2-design once*

$$L \;\geq\; C\, \Delta_2^{-1}\big(n_{\mathrm{act}} \log q + \log(1/\varepsilon)\big) \;=\; \mathrm{poly}(n_{\mathrm{act}}, \log(1/\varepsilon)),$$

using $\Delta_2 \geq 1/\mathrm{poly}(n_{\mathrm{act}})$ *from Lemma 4, with* C *depending only on the local gate family. No expansion or sector-independent gap is assumed; for a worst-case connected graph* without *the polynomial-gap condition one retains only* $\Delta_2 > 0$ *(design exists, depth not controlled).*

Proof. Insert $\Delta_2 \geq 1/\mathrm{poly}(n_{\mathrm{act}})$ from Lemma 4 into the depth bound of Theorem 3. □

Lemma 5 (Evaporation outlasts the design depth)**.** *Adopt the explicit gate-rate model in which the active horizon applies SVP two-site gates at the local Planckian/thermal rate, so the cumulative number of available gate applications up to coordinate time* t *is*

$$N_{\mathrm{gate}}^{\mathrm{avail}}(t) \;\sim\; n_{\mathrm{act}}\,\frac{t}{t_{\mathrm{Pl}}},$$

and let the total Schwarzschild evaporation time be $t_{\mathrm{evap}} \sim M^3 t_{\mathrm{Pl}} \sim n_{\mathrm{tot}}^{3/2}\, t_{\mathrm{Pl}}$*. Then the cumulative available depth over the evaporation history,* $N_{\mathrm{gate}}^{\mathrm{avail}}(t_{\mathrm{evap}}) \sim n_{\mathrm{tot}}^{5/2}$*, exceeds the design depth* $L = \mathrm{poly}(n_{\mathrm{act}})$ *of Corollary 3 once the polynomial exponent in* L *is below* $5/2$*; under the stated model the 2-design therefore forms within the evaporation lifetime.*

Proof. Both sides are explicit powers of n_{tot} under the stated gate-rate model: available depth $\sim n_{\mathrm{tot}}^{5/2}$ versus required $L = \mathrm{poly}(n)$. The margin is positive whenever the design-depth exponent is $< 5/2$, which holds for the $\Delta_2^{-1} = \mathrm{poly}$ bound of Lemma 4 in the focused-graph family. We do not claim the design forms before *each individual* emission stage without a finer per-stage budget; the load-bearing statement for the Page theorem is only that the design forms within the total evaporation lifetime, which the displayed comparison gives. □

Remark 13 (What expansion would add). If Conjecture 2 holds, $\Delta_2 = \Omega(1)$ and the design depth in Corollary 3 drops to $L = O(n_{\mathrm{act}} \log q + \log 1/\varepsilon)$ (polylogarithmic in the Hilbert dimension), i.e. the Eide horizon is a *fast scrambler* with $t_\star \sim \beta \ln S$, matching the Sekino–Susskind bound. The Page curve does not use this; it is a consistency bonus. The honest status is therefore: connectivity (derived, Prop. 4) ⇒ poly-depth design ⇒ Page curve; expansion (conjectural) ⇒ polylog-depth design ⇒ fast scrambling.

10.2 Design-to-Page saturation

Let $U \sim \mu_{\mathrm{SVP},L}$ act on $|\psi_0\rangle \in \mathcal{H}_R \otimes \mathcal{H}_C$, $|\Psi_U\rangle = U|\psi_0\rangle$, $\rho_R(U) = \mathrm{Tr}_C\,|\Psi_U\rangle\langle\Psi_U|$, and $P_R(U) = \mathrm{Tr}\,\rho_R(U)^2$. Write $d_< = \min\{d_R, d_C\}$, $d_> = \max\{d_R, d_C\}$.

Lemma 6 (Haar purity)**.** *For Haar-random* $|\Psi\rangle \in \mathcal{H}_R \otimes \mathcal{H}_C$*,* $\mathbb{E}_{\mathrm{Haar}}\,\mathrm{Tr}\,\rho_R^2 = \dfrac{d_R + d_C}{d_R d_C + 1}$.

Proof. With $D = d_R d_C$ and swap F_R on the two radiation copies, $\mathrm{Tr}\,\rho_R^2 = \mathrm{Tr}[(|\Psi\rangle\langle\Psi|)^{\otimes 2}(F_R \otimes \mathbb{I}_{CC})]$ and $\mathbb{E}_{\mathrm{Haar}}(|\Psi\rangle\langle\Psi|)^{\otimes 2} = (I + F_R \otimes F_C)/D(D+1)$. The two traces are $d_R d_C^2$ and $d_R^2 d_C$; dividing by $D(D+1) = d_R d_C(d_R d_C + 1)$ gives the claim. □

Corollary 4 (SVP purity bound)**.** *If the SVP ensemble is a 2-design accurate enough that, for the purity observable,* $|\mathbb{E}_{\mathrm{SVP}} P_R - \mathbb{E}_{\mathrm{Haar}} P_R| \leq \gamma/d_<$ *with fixed* $\gamma = O(1)$*, then* $\mathbb{E}_{\mathrm{SVP}} P_R \leq (2+\gamma)/d_<$*, and for every* $0 < \delta < 1$*,*

$$\mathbb{P}_{\mathrm{SVP}}\Big[\, S(\rho_R(U)) \geq \ln d_< - \ln \tfrac{2+\gamma}{\delta} \,\Big] \geq 1 - \delta. \tag{15}$$

Proof. $\mathbb{E}_{\mathrm{Haar}} P_R = (d_R + d_C)/(d_R d_C + 1) \leq 1/d_R + 1/d_C = (1/d_<)(1 + d_</d_>) \leq 2/d_<$, so $\mathbb{E}_{\mathrm{SVP}} P_R \leq (2+\gamma)/d_<$. Markov gives $\mathbb{P}_{\mathrm{SVP}}[P_R \geq (2+\gamma)/\delta d_<] \leq \delta$; on the complement $-\ln P_R \geq \ln d_< - \ln((2+\gamma)/\delta)$, and $S(\rho_R) \geq S_2(\rho_R) = -\ln P_R$. □

11 The finite-code Page theorem

Lemma 7 (Page envelope)**.** *If the fixed-sector state is pure with* $\mathcal{H}^q_{\mathrm{phys},N}(t) \cong \mathcal{H}_R(t) \otimes \mathcal{H}_C(t)$*, then* $S_R(t) \le \min\{\ln d_R(t), \ln d_C(t)\}$.

Proof. Schmidt-decompose the pure state; ρ_R has rank $r(t) \le \min\{d_R, d_C\}$ and $S_R(t) \le \ln r(t) \le \ln \min\{d_R, d_C\}$. □

Universal theorem versus Eidetic realization. Before stating the main theorem we make its logical division explicit. The Page *step* itself—that a pure bipartite state with a sufficiently scrambled join saturates $S_R = \ln d_< + O(1)$—is universal finite-dimensional quantum information and is not special to Eidetic Theory. What Eidetic Theory contributes is a proposed *microscopic realization of the theorem's hypotheses*: the purifier (the backward-branch complement algebra, §4, §8), the dimension budget (d_R, d_C from bidirectional emission, §9), the scrambling graph (the null-intersection graph, §5), and the horizon-cell count ($\lambda_\star$ from flux quantization, §6). The theorem below should be read in that light: it is the universal Page step, instantiated on hypotheses that the Eidetic primitives are claimed to supply.

Theorem 4 (Eidetic Page curve)**.** *Let* $\mathcal{A}^q_{\mathrm{phys},N}$ *be a fixed-sector Eide physical code algebra,* $\mathcal{A}_R(t)$ *the radiation algebra,* $\mathcal{A}^q_C(t) = \mathcal{A}_R(t)' \cap \mathcal{A}^q_{\mathrm{phys},N}$ *its complement, and assume the center-conditioned factorization* (9)*. Assume: pure formation (supplied by the Big Bang bulkization, Axiom 2, and/or inclusion of external references in the complement); global fixed-sector unitarity with induced one-branch channel (Proposition 2); bidirectional emission as code-isometric repartitioning* (11)*; isometric return (Lemma 1); independently fixed logical dimensions* $d_R(t), d_C(t)$*; and Axiom 3 with SVP depth large enough that Theorem 3 gives a 2-design meeting the hypothesis of Corollary 4. Then for every* $0 < \delta < 1$*, with* $C_{\delta,\gamma} := \ln((2+\gamma)/\delta)$*,*

$$\mathbb{P}_{\mathrm{SVP}}\big[\, |S_R(t) - \min\{\ln d_R(t), \ln d_C(t)\}| \le C_{\delta,\gamma} \,\big] \ge 1 - \delta. \tag{16}$$

For uniform logical Eide qubits, with $d_R(t) = 2^{n_R(t)}$*,* $d_C(t) = 2^{n_C(t)}$*,* $n_R + n_C = n_{\mathrm{tot}}$*,*

$$\mathbb{P}_{\mathrm{SVP}}\big[\, |S_R(t) - \min\{n_R(t), n_C(t)\} \ln 2| \le C_{\delta,\gamma} \,\big] \ge 1 - \delta. \tag{17}$$

Proof. The logical budget fixes d_R, d_C before S_R is computed. By pure formation and Proposition 2 the fixed-sector state is pure in $\mathcal{H}_R \otimes \mathcal{H}_C$; Lemma 7 gives $S_R \le \ln d_<$. By Theorem 3 the finite-depth SVP ensemble is a 2-design at the chosen accuracy, so Corollary 4 gives $S_R \ge \ln d_< - C_{\delta,\gamma}$ with probability $\ge 1 - \delta$. Combine on that event and substitute $d_R = 2^{n_R}, d_C = 2^{n_C}$. □

Corollary 5 (Nonuniform finite Eide modes)**.** *Let* R_k, C_k *be the emitted and complement sets with finite dimensions* d_α *and* $d_R(k) = \prod_{\alpha \in R_k} d_\alpha$*,* $d_C(k) = \prod_{\alpha \in C_k} d_\alpha$*. Under the hypotheses of Theorem 4, for every* $0 < \delta < 1$*,*

$$\mathbb{P}_{\mathrm{SVP}}\Big[\, \big|S_R(k) - \min\{\textstyle\sum_{\alpha \in R_k} \ln d_\alpha, \sum_{\alpha \in C_k} \ln d_\alpha\}\big| \le C_{\delta,\gamma} \,\Big] \ge 1 - \delta.$$

Corollary 6 (Center-sector Page curve)**.** *If the center is not fixed, write* $\mathcal{H}^q_{\mathrm{phys}}(k) \cong \bigoplus_z \mathcal{H}_{R,z}(k) \otimes \mathcal{H}_{C,z}(k)$ *with classical sector weights* p_z*. Then under SVP design dynamics inside each fixed center sector, for every* $0 < \delta < 1$*,*

$$S_R(k) = H(p_z) + \sum_z p_z \, \min\{\ln d_{R,z}(k), \ln d_{C,z}(k)\} + O_\delta(1)$$

for SVP-typical histories in each block, with $H(p_z) = -\sum_z p_z \ln p_z$. *For a sharply fixed center,* $H(p_z) = 0$ *and this reduces to Theorem 4.*[1]

Proof. The entropy of a block-diagonal state is the Shannon entropy of the block weights plus the average within-block entropy; apply Theorem 4 in each block. □

Page turnover. The turnover is the dimension crossing $d_R(t_P) \approx d_C(t_P)$, i.e. $n_R(t_P) \approx n_C(t_P)$; for even $n_{\rm tot}$, $n_R(t_P) = n_C(t_P) = \frac{1}{2}n_{\rm tot}$ and peak entropy $S_R(t_P) = \frac{1}{2}n_{\rm tot}\ln 2 + O_\delta(1)$. The scrambling/design time (local mixing on $G_{\rm act}$) is distinct from the Page turnover (dimension crossing).

11.1 Whole-curve (simultaneous) statement

Theorem 4 is a *fixed-time* statement: at each emitted count k the Page value holds with probability $\geq 1-\delta$. The phrase "the Page curve" refers to all stages at once, so we promote the fixed-time bound to a simultaneous one by a union bound over the discrete emission stages.

Corollary 7 (Whole discrete Page curve)**.** *Let there be* $N_{\rm step} = n_{\rm tot} + 1$ *discrete emission stages* $k = 0, 1, \ldots, n_{\rm tot}$, *and choose the SVP depth at each stage so that Theorem 4 applies with per-stage failure probability* $\delta_{\rm step}$. *Setting* $\delta_{\rm step} = \delta_{\rm tot}/N_{\rm step}$, *the constant becomes*

$$C_{\delta_{\rm step},\gamma} = \ln\Big(\frac{(2+\gamma)N_{\rm step}}{\delta_{\rm tot}}\Big) = O_{\delta_{\rm tot}}\big(\ln n_{\rm tot}\big),$$

and the entire discrete curve holds simultaneously:

$$\mathbb{P}_{\rm SVP}\Big[\,\forall k:\ \big|S_R(k) - \min\{\ln d_R(k), \ln d_C(k)\}\big| \leq \ln\tfrac{(2+\gamma)(n_{\rm tot}+1)}{\delta_{\rm tot}}\,\Big] \geq 1-\delta_{\rm tot}.$$

Proof. By Theorem 4, at each fixed stage k the event $E_k = \{|S_R(k) - \ln d_<(k)| > C_{\delta_{\rm step},\gamma}\}$ has $\mathbb{P}_{\rm SVP}[E_k] \leq \delta_{\rm step}$. By the union bound $\mathbb{P}_{\rm SVP}[\bigcup_k E_k] \leq N_{\rm step}\,\delta_{\rm step} = \delta_{\rm tot}$, so the complement—the simultaneous bound—holds with probability $\geq 1-\delta_{\rm tot}$. Substituting $\delta_{\rm step} = \delta_{\rm tot}/N_{\rm step}$ into $C_{\delta_{\rm step},\gamma} = \ln((2+\gamma)/\delta_{\rm step})$ gives the displayed constant, which is $O(\ln n_{\rm tot})$. □

Remark 14. The $O(\ln n_{\rm tot})$ price of simultaneity is of the same order as the $O(\ln A)$ microstate-counting correction already present in the geometric curve (Corollary 8), since $n_{\rm tot}\ln 2 = A(0)/4\ell^2_{\rm Pl,phys} + O(\ln A)$. Promoting the fixed-time theorem to the whole curve therefore costs nothing in the final geometric form: the simultaneous error remains $O_\delta(L_0)$.

12 Geometric Page curve

Let $L_0 := \ln(2 + A(0)/\ell^2_{\rm Pl,phys})$. Since $A(t) \leq A(0)$ in no-remnant single-reservoir evaporation, the microstate-counting corrections are uniformly $O(L_0)$.

Corollary 8 (Geometric Eidetic Page curve)**.** *In a no-remnant single-reservoir sector whose horizon-cut complement is counted by Theorem 1 with quartic coefficient fixed by Corollary 2, for every* $0 < \delta < 1$, *with SVP probability* $\geq 1-\delta$,

$$S_R(t) = \min\left\{\frac{A(0)-A(t)}{4\ell^2_{\rm Pl,phys}},\ \frac{A(t)}{4\ell^2_{\rm Pl,phys}}\right\} + O_\delta(L_0), \tag{18}$$

[1] Probability bookkeeping over sectors: for finitely many center sectors the high-probability statement follows by a union bound over z (allocating per-sector failure δ_z with $\sum_z \delta_z = \delta$), so the $O_\delta(1)$ term picks up an additional $O(\ln|\{z\}|)$, analogous to Corollary 7. For a broad center distribution, the $O_\delta(1)$ term is understood after restricting to the sectors carrying all but $O(\delta)$ of the weight p_z.

the Page point is $A(t_P) = \frac{1}{2}A(0) + O(\ell^2_{\mathrm{Pl,phys}} L_0)$*, and the peak entropy is* $S_R(t_P) = A(0)/8\ell^2_{\mathrm{Pl,phys}} + O_\delta(L_0)$*.*

Proof. Theorem 4 gives $S_R = \min\{\ln d_R, \ln d_C\} + O_\delta(1)$. Equations (13),(14) give $\ln d_C = A(t)/4\ell^2_{\mathrm{Pl,phys}} + O(L_0)$ and $\ln d_R = (A(0) - A(t))/4\ell^2_{\mathrm{Pl,phys}} + O(L_0)$; min of two quantities each $+O(L_0)$ is the min of the leading terms $+O(L_0)$. The simultaneous (whole-curve) version follows by applying Corollary 7 stage-by-stage; its $O(\ln n_{\mathrm{tot}})$ union-bound penalty is absorbed into $O_\delta(L_0)$. □

13 Cosmological unification: the de Sitter static patch

Because Axiom 1 replaces *both* de Sitter spacelike boundaries by the same pair network, a de Sitter static-patch horizon is an Eide cut of exactly the same type as a black-hole horizon, counted by the same Proposition 5 and the same area quantum (8).

Corollary 9 (Unified horizon entropy)**.** *For any reconstructed Killing/static-patch horizon of area* A_{dS} *in a smooth Eide sector,*

$$S_{dS} = \ln \dim \mathcal{H}_C(A_{dS}) = \frac{A_{dS}}{4\ell^2_{\mathrm{Pl,phys}}} + O(\ln A_{dS}),$$

with the same *leading coefficient* $1/4\ell^2_{\mathrm{Pl,phys}}$ *and,* provided the two horizons lie in the same fixed-center Eide surface-code universality class*, the* same *subleading* $O(\ln A)$ *coefficient as the black-hole horizon, since both are then governed by the identical flux transfer operator and area quantum. (The leading coefficient is universal regardless; the subleading coefficient can in principle depend on topology, zero modes, boundary conditions, or ensemble choice, which is why the universality-class proviso is required for the subleading equality.)*

This is a nontrivial prediction: the equality of the subleading logarithmic coefficients for black-hole and cosmological horizons is forced by their shared Eide microphysics and is not required by approaches that treat the two horizons with independent counting. The Big Bang as first global bulkization (Axiom 2) is the cosmological initial condition that starts the bulk in the low-rank pure state used in Theorem 4; the de Sitter pair network thus supplies both the purity premise and the cosmological entropy in one structure.

14 Predictions and falsifiability

A reframing becomes a theory when it predicts or derives something a relabeling cannot. The genuinely Eidetic construction yields the following falsifiable consequences.

1. **Discrete area spectrum / area quantum.** Corollary 2 predicts the intrinsic area quantum $a_\star = 4\ell^2_{\mathrm{Pl,phys}} \ln 2$ under the minimal unconstrained-cell hypothesis, hence a discrete horizon area spectrum $A_n = n\, a_\star + O(1)$ and a corresponding discrete mass spectrum. This connects to the quasinormal-mode "$\ln 2$ vs. $\ln 3$" area-spectrum discussion: a fundamental flux sector $k = 1$ gives $\ln \lambda_\star = \ln 2$, while $k = 2$ would give $\ln 3$. We stress, however, that whether this area quantum appears in ringdown overtone spacing or gravitational-wave echo phenomenology requires an additional, not-yet-derived map from area eigenvalues to semi-classical perturbation spectra; absent that map, the prediction is an intrinsic statement about the area spectrum rather than a near-term observational falsification. A genuine flux-closure rule that lowered $\lambda_\star$ below 2 would shift $a_\star$ accordingly.

2. **Black-hole / de Sitter coefficient equality.** Corollary 9 predicts identical leading entropy coefficients for black-hole and static-patch horizons, and identical *subleading* $O(\ln A)$ coefficients *provided both horizons lie in the same fixed-center Eide surface-code universality class*. A measured or computed mismatch in the $O(\ln A)$ coefficient between the two horizon types—beyond what a genuine universality-class difference would explain—would falsify the shared-transfer-operator claim.

3. **Backward-branch Page correction (prediction target).** This is a *candidate* prediction, not yet an established one, because no closed-form correction is derived here. If the past-directed branch carries genuine (non-passive) dynamics consistent with Conjecture 1, the final-state structure should produce a small, calculable deviation near the turnover, distinct from the exact Haar Page deficit $\mathbb{E}S_{m,n} = \ln m - m/2n + O(1/n)$ (Appendix C). Deriving a closed-form backward-branch correction term would be the single sharpest demonstration that the Eidetic ontology is doing work; its absence would mean the backward branch is operationally inert and should be dropped.

4. **Replica-saddle reproduction.** Since the island is the bulk reconstruction of the backward-branch complement (Remark 8), a correctly constructed R_ω should *reproduce* the replica-wormhole saddle of [19, 20], not merely agree with its answer. Recovering the saddle inside the Eide formalism is a positive test; an obstruction to doing so would indicate the reconstruction map is incompatible with known semiclassical gravity.

15 Consequences for black-hole information

15.1 No information loss at the singularity

A deletion singularity would act as $\rho^{(a)}_{\text{bulk}} \mapsto (\operatorname{Tr} \rho^{(a)}_{\text{bulk}})\rho_0$. Eidetic return acts as $\rho^{(a)}_{\text{bulk}} \mapsto V_{\text{ret},a}\rho^{(a)}_{\text{bulk}}V^\dagger_{\text{ret},a}$ with $V^\dagger_{\text{ret},a}V_{\text{ret},a} = \mathbb{I}$. By Lemma 1 entropy and nonzero spectrum are preserved. The singularity marks the failure of faithful bulk-local representation, not of unitarity.

15.2 No duplicate purifier; firewall tension reduced to benign reconstruction

The AMPS monogamy tension [18] is that a late Hawking mode cannot be independently purified both by early radiation and by a separate interior partner while keeping ordinary effective field theory at the horizon. In the Eidetic code there is one purifier, $\text{Purifier}(B) = \mathcal{A}^q_C(t) = \mathcal{A}_R(t)' \cap \mathcal{A}^q_{\text{phys},N}$, whose generators are the past-directed branches of the emitted pairs (Remark 8). The early-radiation purifier, the interior partner, the return register, and the island are not independent factors; they are reconstructions/subalgebras of the same backward-branch algebra. The *duplicate-purifier* counting obstruction is therefore removed algebraically: there is no second, independent Hilbert factor to purify the late mode, so the monogamy bookkeeping that drives the AMPS argument does not arise. This is a structural statement about the purifier, not a proof of horizon smoothness.

What remains is precisely the smoothness question. Whether the infalling observer sees a benign (firewall-free) horizon is controlled by Conjecture 1; a minimal statement is $\|\rho_{B\widetilde{B}} - \rho_{\text{vac}}\|_1 \le \epsilon_{\text{hor}}$, with $\widetilde{B}$ the complement reconstruction used by the infalling observer. The Page theorem proves unitary fine-grained evaporation regardless of whether this smoothness condition holds.

15.3 Islands and static patches as reconstructions

Modern island calculations reproduce the Page curve semiclassically via quantum extremal surfaces and entanglement-wedge reconstruction [19, 20]. In Eidetic Theory an island is not fundamental: it is a semiclassical reconstruction of part of the backward-branch complement, $\mathcal{A}(I) \subseteq R_\omega(\mathcal{A}_C^q(t))$. Static-patch screens are observer-dependent reconstructions of selected boundary-code data. The hierarchy is

$$\text{AdS/CFT: timelike asymptotic boundary;} \quad \text{global dS/CFT: spacelike } \mathcal{I}^{\pm};$$

$$\text{static-patch holography: observer horizon;} \quad \text{Eidetic: pre-bulk bidirectional } S^2 \text{ net.}$$

The purifier is thus neither at an unreachable infinity nor merely an observer screen; it is the backward-branch algebraic complement inside the Eide code.

16 Optional physical clock profile

The theorem is count-parametric: S_R as a function of emitted logical capacity. A physical time profile needs an independent Eidetic emission law, e.g. $dA/dt = \mathcal{F}_{\text{Eide}}(A, Q, J, \dots)$ or $dn_R/dt = \Gamma_{\text{Eide}}(n_C, q)$, which is *not* used in the proof. If one separately establishes a finite-window logistic count equation for a smooth $N_R(t)$,

$$\frac{dN_R}{dt} = \frac{2}{\tau_E N_{\text{tot}}} N_R(N_{\text{tot}} - N_R), \quad N_R(t_P) = \tfrac{1}{2}N_{\text{tot}},$$

then $N_R(t) = \frac{1}{2}N_{\text{tot}}(1 + \tanh((t-t_P)/\tau_E))$ and $S_R(t) = \frac{1}{2}N_{\text{tot}} \ln 2\,(1 - |\tanh((t-t_P)/\tau_E)|) + O(1)$. This is a clock parametrization, not a premise.

17 Non-circularity and math checks

1. d_R, d_C (or n_R, n_C) are Hilbert-space *dimensions*, defined before S_R.
2. Eide-Sphere flux quantization fixes capacities, not radiation entropy; the two-state alphabet is *derived* and the value $\lambda_\star = 2$ follows under the minimal unconstrained-cell hypothesis (Proposition 5), not assumed.
3. The Page minimum first appears as a Schmidt-rank upper bound from global purity, where global purity is itself a consequence of bidirectional pairing (Proposition 2).
4. Equality with the minimum is obtained only after the SVP design theorem gives $P_R = O(1/d_<)$; the design property is itself derived from null-propagation locality plus connectivity (Proposition 4), with the focused-graph polynomial-gap condition (Definition 6) supplying polynomial depth (Lemma 4, Corollary 3), so no expansion assumption enters the Page curve.
5. $\eta_{\text{Eide}} = \ln \lambda_\star / a_\star$ is derived before any Page discussion; $1/4$ is then fixed by the independent emergent-Einstein matching, and the falsifiable output is $a_\star = 4\ell^2_{\text{Pl,phys}} \ln 2$.
6. The complement is defined algebraically as $\mathcal{A}_C^q = \mathcal{A}_R' \cap \mathcal{A}^q_{\text{phys}}$ and identified with the backward-branch algebra; it is not inserted as an island.
7. The partial trace over C is the type-I representative of restriction to $\mathcal{A}_R$, not physical information loss.

8. Singularity return is isometric (the backward branch read bulk-to-Eide) and preserves nonzero spectrum.

9. Design/scrambling time and Page turnover are distinct.

10. The optional clock profile is unused.

The key checks:

$$\ln\dim\mathcal{H}_C(A) = N_A\ln\lambda_\star + O(\ln N_A) = \eta_{\text{Eide}}A + O(\ln A), \quad \eta_{\text{Eide}} = \tfrac{\ln\lambda_\star}{a_\star},\ \lambda_\star = 2,$$
$$\eta_{\text{Eide}} = \tfrac{1}{4G_{\text{Eide}}\hbar} = \tfrac{1}{4\ell^2_{\text{Pl,phys}}}\ (c=1), \qquad a_\star = 4\ell^2_{\text{Pl,phys}}\ln 2,$$
$$S_R \le \ln\operatorname{rank}\rho_R \le \ln\min\{d_R, d_C\}, \qquad \mathbb{E}_{\text{Haar}}\operatorname{Tr}\rho_R^2 = \tfrac{d_R+d_C}{d_Rd_C+1} \le \tfrac{2}{d_<},$$
$$\mathbb{E}_{\text{SVP}}\operatorname{Tr}\rho_R^2 \le \tfrac{2+\gamma}{d_<}\ \text{(2-design)}, \qquad \mathbb{P}_{\text{SVP}}[\operatorname{Tr}\rho_R^2 \le \tfrac{2+\gamma}{\delta d_<}] \ge 1-\delta\ \text{(Markov)},$$
$$S_R \ge S_2(\rho_R) = -\ln\operatorname{Tr}\rho_R^2 \implies \ln d_< - O_\delta(1) \le S_R \le \ln d_<.$$

18 What is proved and what is deferred

Proved (relative to the stated Eidetic primitives).

- From bidirectional pairing plus code-isometric repartitioning: global fixed-sector purity, the induced one-branch CPTP channel (Prop. 2), and isometric, non-deleting singularity return (Lemma 1). The Bell pair supplies the purifier; the isometric update preserves purity; the full microscopic dynamics is not derived. Bidirectional pairing supplies the *rising* branch alone; the *falling* branch (the Page turnover) requires the single-reservoir/no-remnant condition, which is a genuine physical input, not a consequence of pairing (Lemma 3, demonstrated in Remark 10). The AMPS duplicate-purifier obstruction is removed by a built-in ER = EPR; horizon smoothness remains conditional (§15).

- From null propagation: the SVP interaction graph is local and connected (Prop. 4); connectivity gives a positive second-moment gap, and under the focused-graph polynomial-gap condition (Definition 6) this becomes a polynomial-depth approximate 2-design (Lemma 4, Cor. 3) that forms within the evaporation lifetime under the stated gate-rate model and design-depth exponent (Lemma 5). Full expansion is not required—it survives only as an optional sharpening giving fast scrambling (Conj. 2).

- From S^2 flux quantization: the two-state horizon-cell alphabet unconditionally; the minimal independent-cell realization of Axiom 1 (Prop. 6, the natural reading of the axiom, not forced by it); and a provably subextensive global flux closure (Lemma 2), together fixing $\lambda_\star = 2$ in the balanced sector (Cor. 1) and giving $\ln\dim\mathcal{H}_C(A) = \eta_{\text{Eide}}A + O(\ln A)$ (Thm. 1).

- From emergent-Einstein matching (Jacobson-type): $\eta_{\text{Eide}} = 1/4G_{\text{Eide}}\hbar$ and the predicted area quantum $a_\star = 4\ell^2_{\text{Pl,phys}}\ln 2$ (Thm. 2, Cor. 2).

- The finite-code and geometric Page curves (Thm. 4, Cor. 8), and their de Sitter unification (Cor. 9).

- The recovery-fidelity part of benign reconstruction (no teleportation corruption): the infalling information is faithfully recoverable from the radiation, which follows from the same approximate 2-design as the Page curve via Hayden–Preskill decoupling (Proposition 3, demonstrated in Remark 2).

The proven implication is:

bidirectional purity+poly-gap design+flux-counted area law+Einstein matching $\Longrightarrow$ geometric Page curve,

with the genuine residual inputs being horizon smoothness (the firewall question, part 1b), the universality of the local SVP gate family, the focused-graph polynomial-gap condition, the single-reservoir/no-remnant condition, and the balanced-sector restriction of the area count. The recovery-fidelity part of benign reconstruction is *not* a residual input: it is discharged from the design.

Deferred / conjectural.

1. Horizon smoothness (part 1b of benign reconstruction, Conj. 1): that the infalling observer sees vacuum (no firewall). This is the AMPS question and is *not* settled by scrambling; it is genuinely conjectural. The companion recovery-fidelity part (1a) is by contrast discharged from the design (Proposition 3), so what remains here is smoothness alone, not teleportation corruption.

2. Expansion of the focused null-intersection graph (Conj. 2). This is now an *optional sharpening*, not a hypothesis of the Page curve: it upgrades the design depth from polynomial to poly-logarithmic (fast scrambling), but §10.1 shows the Page curve needs only the polynomial-depth design from the focused-graph polynomial-gap condition.

3. The focused-graph polynomial-gap condition (Definition 6) itself: that the near-horizon null-intersection graph families have inverse-polynomial random-walk gap. Connectivity (Prop. 4) gives only positivity; the inverse-polynomial rate is argued for the focused families but is a property of those families, and the lifetime comparison (Lemma 5) is conditional on a gate-rate model and design-depth exponent.

4. The single-reservoir/no-remnant condition (Lemma 3): that the backward-branch purifier is a finite shared reservoir whose degrees are transferred by emission, rather than a supply of private partners. This is what produces the Page turnover (the falling branch); bidirectional pairing alone gives only the rising Hawking line. A microscopic derivation of why Eide emission transfers from a fixed reservoir—rather than spawning fresh partners—is deferred to the dynamics of R_ω.

5. The full microscopic action/Hamiltonian whose continuum reconstruction yields (5).

6. The physical evaporation rate $A(t)$ (graybody factors, backreaction, charge, spin).

7. The unique microscopic origin of the SVP local gate universality (Axiom 3(2)); and, for the area count, whether the reconstruction map R_ω imposes an *effective local* flux-matching that would replace the complete-graph cell alphabet by a subshift and lower $\lambda_\star$ below 2 (Cor. 1), together with the charged/rotating transfer operator for non-balanced sectors (Lemma 2).

8. The polynomial finite-size degeneracy of the genuine 2D horizon surface code (Remark 3): establishing $|W_N| = \lambda_\star^N N^{O(1)}$ for the fixed-center closed-horizon code, on which the $O(\ln A)$ subleading correction depends (the leading area law survives regardless).

9. Standard Model gauge emergence and a complete de Sitter holographic dual.

10. An explicit construction of R_ω sufficient to test prediction 4 of §14 (replica-saddle reproduction).

These are logically separate from the finite-dimensional Page theorem. The present edition is deliberately *narrower* than a generic finite-dimensional argument: it proves less unconditionally, but everything it proves is anchored in a distinctive Eidetic primitive and it makes falsifiable predictions (§14).

A Single Eide-Sphere qubit construction

With $\Sigma \cong S^2 \cong \mathbb{CP}^1$ and $\frac{1}{2\pi}\int_\Sigma \Omega = k \in \mathbb{Z}$, geometric quantization with $\mathcal{O}(k)$ gives $\dim H^0(\mathbb{CP}^1, \mathcal{O}(k)) = k+1$. For $k=1$, $\mathcal{H}_\Sigma \cong \mathbb{C}^2$; the past branch uses the conjugate (antiholomorphic) polarization $\overline{H^0(\mathbb{CP}^1, \mathcal{O}(1))} \cong \mathbb{C}^2$ of the same dimension (not $H^0(\mathbb{CP}^1, \mathcal{O}(-1))$, which vanishes). Pair convention $k_{\alpha,+} = +1$, $k_{\alpha,-} = -1$, sum 0, $K_\alpha = \mathcal{H}_{\alpha,+} \otimes \mathcal{H}_{\alpha,-}$. (See Prop. 1; this appendix is the minimal case used to derive $\lambda_\star = 2$ in Prop. 5.)

B Algebraic restriction in finite type-I form

Let $\mathcal{H} = \mathcal{H}_R \otimes \mathcal{H}_C$, $\mathcal{A}_R = B(\mathcal{H}_R) \otimes \mathbb{I}_C$. For a normal state ρ, the restriction to $\mathcal{A}_R$ is $\omega_R(X \otimes \mathbb{I}_C) = \mathrm{Tr}_{\mathcal{H}}(\rho(X \otimes \mathbb{I}_C))$, represented by $\rho_R = \mathrm{Tr}_C \rho$ since $\mathrm{Tr}_{\mathcal{H}_R}(\rho_R X) = \mathrm{Tr}_{\mathcal{H}}(\rho(X \otimes \mathbb{I}_C))$. This is a calculational representative of restriction, not a physical discarding of $\mathcal{H}_C$.

C Exact Page formula and asymptotics

For $m \le n$, Page's average entropy of the m-dimensional subsystem of a Haar-random pure state on $\mathbb{C}^m \otimes \mathbb{C}^n$ is $\mathbb{E}S_{m,n} = H_{mn} - H_n - \frac{m-1}{2n}$, $H_N = \sum_{j=1}^N j^{-1}$. Using $H_N = \ln N + \gamma + \frac{1}{2N} + O(N^{-2})$,

$$\mathbb{E}S_{m,n} = \ln m - \frac{m^2-1}{2mn} + O(n^{-2}) = \ln m - \frac{m}{2n} + O(1/n).$$

Setting $m = d_<$, $n = d_>$ gives the Page deficit; the design proof of §10 gives an $O(1)$ deficit rather than this exact one, which suffices for the finite-code curve.

D Support dimension for the continuity bound

For radiation marginals ρ_R, σ_R of pure states on $\mathcal{H}_R \otimes \mathcal{H}_C$, $\mathrm{rank}\,\rho_R, \mathrm{rank}\,\sigma_R \le d_<$, so the joint support has dimension $\le \min\{d_R, 2d_<\} = D_{\mathrm{eff}}$, and Audenaert's bound [6] may be applied with $D = D_{\mathrm{eff}}$ when $d_R \gg d_C$.

E Endpoint cases

If $d_R = 1$ or $d_C = 1$ then $d_< = 1$; a pure state has Schmidt rank one, $S_R = 0$, and $S_R = \ln d_< + O(1)$ is exact since $\ln d_< = 0$. The continuity step is unnecessary here.

References

[1] S. W. Hawking, "Particle Creation by Black Holes," *Commun. Math. Phys.* **43**, 199 (1975).

[2] J. D. Bekenstein, "Black Holes and Entropy," *Phys. Rev. D* **7**, 2333 (1973).

[3] D. N. Page, "Average Entropy of a Subsystem," *Phys. Rev. Lett.* **71**, 1291 (1993), arXiv:gr-qc/9305007.

[4] P. Hayden, J. Preskill, "Black Holes as Mirrors: Quantum Information in Random Subsystems," *JHEP* **0709**, 120 (2007), arXiv:0708.4025.

[5] F. Dupuis, M. Berta, J. Wullschleger, R. Renner, "One-Shot Decoupling," *Commun. Math. Phys.* **328**, 251 (2014), arXiv:1012.6044.

[6] K. M. R. Audenaert, "A Sharp Continuity Estimate for the von Neumann Entropy," *J. Phys. A* **40**, 8127 (2007), arXiv:quant-ph/0610146.

[7] C. Dankert, R. Cleve, J. Emerson, E. Livine, "Exact and Approximate Unitary 2-Designs," *Phys. Rev. A* **80**, 012304 (2009), arXiv:quant-ph/0606161.

[8] A. W. Harrow, R. A. Low, "Random Quantum Circuits are Approximate 2-designs," *Commun. Math. Phys.* **291**, 257 (2009), arXiv:0802.1919.

[9] F. G. S. L. Brandao, A. W. Harrow, M. Horodecki, "Local Random Quantum Circuits are Approximate Polynomial-Designs," *Commun. Math. Phys.* **346**, 397 (2016), arXiv:1208.0692.

[10] P. Caputo, T. M. Liggett, T. Richthammer, "Proof of Aldous' Spectral Gap Conjecture," *J. Amer. Math. Soc.* **23**, 831 (2010), arXiv:0906.1238.

[11] P. Diaconis, L. Saloff-Coste, "Comparison Theorems for Reversible Markov Chains," *Ann. Appl. Probab.* **3**, 696 (1993).

[12] J. M. Bardeen, B. Carter, S. W. Hawking, "The Four Laws of Black Hole Mechanics," *Commun. Math. Phys.* **31**, 161 (1973).

[13] W. G. Unruh, "Notes on Black-Hole Evaporation," *Phys. Rev. D* **14**, 870 (1976).

[14] R. M. Wald, "Black Hole Entropy is the Noether Charge," *Phys. Rev. D* **48**, R3427 (1993), arXiv:gr-qc/9307038.

[15] V. Iyer, R. M. Wald, "Some Properties of Noether Charge," *Phys. Rev. D* **50**, 846 (1994), arXiv:gr-qc/9403028.

[16] T. Jacobson, "Thermodynamics of Spacetime: The Einstein Equation of State," *Phys. Rev. Lett.* **75**, 1260 (1995), arXiv:gr-qc/9504004.

[17] J. J. Bisognano, E. H. Wichmann, "On the Duality Condition for a Hermitian Scalar Field," *J. Math. Phys.* **16**, 985 (1975).

[18] A. Almheiri, D. Marolf, J. Polchinski, J. Sully, "Black Holes: Complementarity or Firewalls?," *JHEP* **2013**, 062, arXiv:1207.3123.

[19] G. Penington, "Entanglement Wedge Reconstruction and the Information Paradox," *JHEP* **2020**, 002, arXiv:1905.08255.

[20] A. Almheiri, N. Engelhardt, D. Marolf, H. Maxfield, "The Entropy of Bulk Quantum Fields and the Entanglement Wedge of an Evaporating Black Hole," *JHEP* **2019**, 063, arXiv:1905.08762.

[21] J. M. Maldacena, "The Large N Limit of Superconformal Field Theories and Supergravity," *Adv. Theor. Math. Phys.* **2**, 231 (1998), arXiv:hep-th/9711200.

[22] E. Witten, "Anti-de Sitter Space and Holography," *Adv. Theor. Math. Phys.* **2**, 253 (1998), arXiv:hep-th/9802150.

[23] A. Strominger, "The dS/CFT Correspondence," *JHEP* **2001**, 034, arXiv:hep-th/0106113.

[24] G. T. Horowitz, J. Maldacena, "The Black Hole Final State," *JHEP* **2004**, 008, arXiv:hep-th/0310281.