How Fast Can a Horizon Learn Which Path You Took?

Registered research scaffold — no numerical result

Brecht Corbeel
Absolute Digital Publishers
absolutedigitalpublishers.com

September 20, 2026

Abstract

Semiclassical field calculations show that a stationary charged or massive body held in a spatial superposition near a Killing horizon can entangle with soft field modes and lose interference visibility. That mechanism is inherited, not proposed here. This draft registers, without solving, an operational translation from a conditional horizon-state overlap $V(T) = |\langle h_0(T)|h_1(T)\rangle|$ to three information quantities that must not be conflated: the trace-distance/Helstrom minimum-error decision information $I_{\mathrm{dec}}(T)$, the Holevo capacity ceiling $\chi_H(T) = h_2\left(\frac{1+V(T)}{2}\right)$, and an irreducible receipt distinguishability that survives the best admissible future attempt at purification. For equal-prior pure conditional states, $V(T)^2 + D(T)^2 = 1$. When an imported field calculation yields $V(T) = e^{-\Gamma T}$, the deficit thresholds $T_\chi(\delta)$ and $T_{\mathrm{dec}}(\delta)$ at which $\chi_H \geq 1 - \delta$ and $I_{\mathrm{dec}} \geq 1 - \delta$ can be inverted exactly from binary entropy; that inversion is an analytic identity reported here, not a horizon calculation, because the rate Γ, its clock convention, and its validity domain are not derived by this scaffold. The program registers, but does not execute, a comparison of rate-provider domains across Schwarzschild, Kerr, Rindler, ordinary-matter, and near-extremal charged-black-hole models, including screening and quantum-spectral-gap regimes in which the accessible literature already establishes a vanishing rate and no finite recording time exists through that channel. The construction is rejected if a boundary-resolved field calculation finds no conditional horizon-state distinguishability where the reduced model predicts a finite record, if a global capacity bound is reported as a causal collector's accessible information, if the receipt clock is not shown robust to enlarging the admissible future-control family, or if the translation changes no parameter ranking relative to visibility alone. No numerical field result, capacity map, receipt time, or external review is reported.

Keywords: black hole decoherence; which-path information; Killing horizons; Holevo information; trace distance; Helstrom discrimination; quantum Darwinism (comparison class).
Series: Revolutionary Life 75, Series D, D02.
Status: Registered scaffold. The horizon field calculation, the Kerr/extremal limit treatment, the finite-time purification module, the 60-source priority review, and external review are incomplete; Section 10 enumerates each open debt and what it blocks.

1 The question is a reading, not a witness

A black hole horizon has no eye, no pointer state, and no moment of comprehension. It is nonetheless established, at the level of semiclassical quantum field theory, that a charged or massive body held in a stationary spatial superposition near a Killing horizon can entangle with soft radiative field modes crossing or approaching that horizon, correlating each path with a distinct conditional field state [1, 2]. If those conditional states become distinguishable, the interference lost by an outside interferometer and the which-path evidence available on or behind the horizon are two descriptions of one entanglement. That mechanism is inherited. It is not this scaffold's contribution, and no part of this paper is entitled to present it as a new finding.

The operational question this project registers is narrower than "does a horizon decohere a superposition." It is: once a field calculation supplies a conditional-state overlap $V(T)$, what is a decoherence rate actually allowed to mean? A visibility curve reports how much fringe contrast an external interferometer has lost. By itself it does not say how many bits an optimal single-shot decision extracts, what an idealized Holevo-class receiver could learn in principle, whether the relevant conditional states are pure or mixed, which spacetime region physically holds the evidence, or whether an admissible future control could still return part of the lost coherence. Treating every exponential decay as a single "learning curve" silently answers four different questions with one symbol.

This draft separates them. It defines a visibility clock, a minimum-error decision clock, a Holevo-ceiling clock, and a receipt clock that only closes after the best physically admissible future purification has been exhausted. It requires every rate object to carry an explicit clock convention – Killing time, local proper time, or coordinate/laboratory time – and forbids multiplying rates and durations expressed in different conventions. And it treats the phrase "the horizon learned the path" as shorthand for a state-discrimination task performed by a specified receiver with specified access, never as a claim about consciousness, measurement collapse, or a privileged physical fact about which branch is real.

None of the definitions below has been evaluated on an actual horizon field calculation. What follows is a candidate translation layer together with the exact conditions under which it must be rejected.

2 The inherited mechanism is not this paper's contribution

Danielson, Satishchandran, and Wald show that a charged or massive body held open in a stationary spatial superposition outside a black hole sources distinct coherent states of low-frequency radiation on the two branches, and that this radiation constitutes a genuine independent degree of freedom rather than a merely constrained near field; consequently, sufficiently slow opening and closing motions can suppress ordinary radiated power to infinity while the horizon-crossing flux still grows with holding time and decoheres the superposition [1]. The result was generalized from black-hole event horizons to a broader class of bifurcate Killing horizons, including Rindler and de Sitter horizons, unifying the acceleration-horizon and black-hole cases under one derivation [2]. A causality argument with an earlier ancestor shows why this is forced rather than optional: if a horizon-crossing observer could in principle distinguish the two branches' long-range fields, an already spacelike-separated interference outcome could not consistently depend on that later, optional measurement; quantized radiation and vacuum fluctuations remove the tension [10].

Gralla and Wei supply the horizon-side calculation this scaffold imports: for nonextremal bifurcate Killing horizons in electromagnetic and Klein–Gordon analogues, the large-holding-time conditional overlap takes the exponential form $V(T) = \exp(-C\kappa T/2)$, with C a dimensionless decohering flux coefficient and κ the surface gravity in the chosen time normalization; the same paper evaluates C on the symmetry axis of a Kerr black hole and identifies an electromagnetic extremal-Kerr limit in which the relevant horizon flux is expelled and both C and κ vanish [3]. Wilson-Gerow, Dugad, and Chen reformulate the same decoherence as an open-system problem: a worldline-localized random-force description in which steady decoherence tracks Ohmic friction and finite-temperature field fluctuations, valid for horizons at nonzero temperature [4]. Danielson, Satishchandran, and Wald give a local two-point-function derivation of the same Schwarzschild effect directly in the laboratory frame, distinguishing Unruh, Hartle–Hawking, and Boulware vacuum choices and explicitly noting a subsequent sign correction to two of their published equations [6]. Danielson, Kudler-Flam, Satishchandran, and Wald solve a genuinely different problem: given the radiation that has already crossed a horizon cross-section by some cutoff, what future field control best preserves the branches' coherence, finding that the optimal continuation can be counterintuitive, including cases where reopening an already-recombined superposition further reduces final decoherence [7].

Two further results bound the regime in which any of the above applies. Biggs and Maldacena show that ordinary finite-temperature dissipative matter, modeled through fluctuating multipole operators, can reproduce qualitatively and in the matched electromagnetic case quantitatively comparable decoherence to a horizon, without any causal horizon at all [5]. Biggs and Trezzi study near-extremal charged Reissner–Nordström black holes with quantized near-horizon geometry and find, through quartic order in the interaction and conjectured to all orders for a stated opposite-dipole configuration, that a spin-induced spectral gap makes the late-time rate vanish below a threshold energy above extremality [8]. Batista, Landulfo, Mann, and Matsas give a nonperturbative Unruh–DeWitt-detector treatment of the same mechanism for a gapless detector coupled to a massive scalar field on uniformly accelerated superposed trajectories, supporting the semiclassical picture beyond leading order in that specific model [9].

Separately, Arrasmith, Albrecht, and Zurek analyze the reversed role assignment – a black hole itself held in a spatial superposition, decohered by its own Hawking radiation revealing its position [16]. That is a different physical problem: here the black hole is environment, not system, and the conditional states, rates, and validity domains do not transfer between the two setups.

No numerical rate, coefficient, spectral gap, or vanishing theorem from any of the above is recomputed, extended, or claimed as new physics anywhere in this scaffold. Every subsequent equation in this paper accepts an overlap or a rate as an import with a declared spacetime, channel, clock convention, and citation, and translates it into operational information quantities. The translation is the candidate contribution; the decoherence mechanism is not.

3 A record is three numbers, not one exponential

Consider an equal superposition of two spatially separated alternatives coupled to an environment that includes horizon-accessible field degrees of freedom,

$$|\Psi(0)\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes |h_{\mathrm{in}}\rangle\,, \tag{1}$$

evolving, under the idealization that the joint system-plus-field state remains pure, to

$$|\Psi(T)\rangle = \frac{|0\rangle\,|h_0(T)\rangle + |1\rangle\,|h_1(T)\rangle}{\sqrt{2}}. \tag{2}$$

Tracing out the field multiplies the off-diagonal element of the reduced path density matrix by the conditional-state overlap

$$\nu(T) = \langle h_0(T)|h_1(T)\rangle, \qquad V(T) = |\nu(T)|. \tag{3}$$

Under balanced arms and ideal readout, $V(T)$ is the surviving fringe visibility; that identification is ordinary two-path complementarity, quantified by Englert's visibility–distinguishability inequality and confirmed independently by atom-interferometer measurements of visibility and which-way distinguishability [11, 12]. Equation (2) and the identification of $|h_0(T)\rangle\,, |h_1(T)\rangle$ with a declared field-theoretic conditional state are inherited from [1–3]; nothing in this section supplies a new mechanism for changing $\nu(T)$.

Assume for this section that $|h_0(T)\rangle$ and $|h_1(T)\rangle$ are normalized, pure, and prepared with equal prior probability – an assumption that Section 5 restricts to a declared horizon subalgebra and that must be revisited whenever the field state is genuinely mixed. Define the conditional-state trace distance

$$D(T) = \tfrac{1}{2}\,\|\,|h_0(T)\rangle\langle h_0(T)| - |h_1(T)\rangle\langle h_1(T)|\,\|_1\,. \tag{4}$$

For two pure states this reduces to a closed form in the overlap,

$$D(T) = \sqrt{1 - V(T)^2}, \qquad V(T)^2 + D(T)^2 = 1. \tag{5}$$

Equation (5) is an exact algebraic identity for pure, equal-prior conditional states; it is not a horizon result, and it fails without modification once the conditional states are mixed, have unequal priors, or are restricted to a proper subalgebra of the full field state (Section 5).

The **Holevo capacity ceiling** of the equal-prior pure ensemble $\{\frac{1}{2}, |h_0(T)\rangle\} \cup \{\frac{1}{2}, |h_1(T)\rangle\}$ is the von Neumann entropy of the average conditional state,

$$\chi_H(T) = h_2\!\left(\frac{1 + V(T)}{2}\right), \qquad h_2(p) = -p\log_2 p - (1-p)\log_2(1-p). \tag{6}$$

Holevo's theorem makes $\chi_H(T)$ an upper bound on the classical information any single measurement can extract from the ensemble, not a guarantee that a one-copy receiver attains it [13]; Fuchs and Caves quantify exactly how far an ensemble's accessible information can fall short of this ceiling [15]. Equation (6) is a direct application of that inherited bound to the pure-state overlap of Eq. (3); it is the paper's proposed *operational* object, not a new information-theoretic result.

Definition 3.1 (Deficit threshold times)**.** *For a declared confidence deficit $\delta \in (0, 1)$, the **Holevo-ceiling time** is*

$$T_\chi(\delta) = \inf\{T : \chi_H(T) \geq 1 - \delta\}, \tag{7}$$

*and, with the minimum-error decision information $I_{\mathrm{dec}}(T)$ of Section 4 substituted for $\chi_H(T)$, the **minimum-error decision time** $T_{\mathrm{dec}}(\delta)$ is defined identically.*

Both quantities in Definition 3.1 are deterministic transformations of an already-assumed overlap curve $V(T)$; neither is a new field-theory result, and neither can be evaluated without first importing Γ or an equivalent rate object from a specific spacetime, channel, and clock convention (Section 6). A code path that evaluates Eq. (7) before that import has occurred is computing an analytic identity, not a horizon-record time; Appendix A makes this distinction explicit and reports only the identity.

Remark 3.1. At $\Gamma = 0$ (no imported rate, or a rate the field calculation sets to zero), $V(T) \equiv 1$ for every T, so $D(T) \equiv 0$ and $\chi_H(T) \equiv 0$: no record accumulates, and $T_\chi(\delta) = \infty$ for every $\delta < 1$. Any implementation that instead returns a large finite number by clipping to a plotting horizon has converted a physical no-record regime into a graphics artifact, and must be treated as a bug, not a result.

4 Distinguishability translates into decision and capacity, not into each other

Trace distance is not itself a bit count; it is the quantity that fixes the performance of the best possible binary decision. Helstrom's minimum-error discrimination theorem gives the smallest achievable average error probability for guessing which of two equal-prior states was prepared,

$$p_{\mathrm{e}}(T) = \frac{1 - D(T)}{2} \overset{\text{pure}}{=} \frac{1 - \sqrt{1 - V(T)^2}}{2}, \tag{8}$$

with the pure-state specialization following from Eq. (5) [14]. For the balanced two-outcome problem, the optimal measurement realizes a symmetric binary channel between the true path label and the decision outcome, whose mutual information is

$$I_{\mathrm{dec}}(T) = 1 - h_2[p_{\mathrm{e}}(T)]\,. \tag{9}$$

Equation (8) inherits quantum minimum-error detection theory; Eq. (9) is the ordinary mutual information of the resulting classical channel. Together they answer a concrete operational question – if the path were chosen fairly and one optimized single-shot yes/no measurement were performed on the conditional field state, how many bits would that decision convey on average – and they answer only that question.

$\chi_H(T)$ from Eq. (6) and $I_{\mathrm{dec}}(T)$ from Eq. (9) are both functions of $V(T)$ alone in the pure equal-prior case, but they are not the same function. At $V = 0.5$, $\chi_H \approx 0.811$ bits while $I_{\mathrm{dec}} \approx 0.646$ bits: the two curves differ by a fixed, calculable amount at every value of V strictly between 0 and 1. This is an algebraic fact about h_2 composed with two different arguments, not a claim about any horizon.

Proposition 4.1 (Holevo-bound ordering)**.** *For every $V \in [0,1]$, $\chi_H \geq I_{\mathrm{dec}}$, with equality only at $V \in \{0,1\}$. Consequently, for any overlap curve $V(T)$ that is non-increasing in T, the deficit threshold times obey $T_\chi(\delta) \leq T_{\mathrm{dec}}(\delta)$ for every $\delta \in (0,1)$.*

Proposition 4.1 restates Holevo's theorem – a single-shot minimum-error measurement cannot exceed the ensemble's Holevo quantity – in the deficit-time language of this paper; it is not a new inequality, and Appendix A verifies it numerically as a self-check rather than reporting it as a discovery. It is nonetheless the paper's first hard rule: no plotted or tabulated pair $(T_\chi, T_{\mathrm{dec}})$ may violate this ordering, and any code producing $T_\chi > T_{\mathrm{dec}}$ has a defect in the entropy inversion, not a new physical effect.

Global capacity is not a causal collector's information. $\chi_H(T)$ bounds what an optimal measurement on the *entire* declared conditional horizon state could extract. A physically situated observer – confined to a world tube, a finite duration, and a bounded region of the horizon or its exterior – may access only a restricted subalgebra. Define, for a modeled causal collector with accessible algebra $\mathcal{A}_{\mathrm{collect}} \subseteq \mathcal{A}_{\mathrm{horizon}}$, a restricted capacity $\chi_{\mathrm{collect}}(T)$ computed identically to Eq. (6) but with the conditional states first reduced to $\mathcal{A}_{\mathrm{collect}}$, and an access gap

$$\Delta I_{\mathrm{access}}(T) = \chi_H(T) - \chi_{\mathrm{collect}}(T) \;\geq\; 0. \tag{10}$$

A large $\Delta I_{\mathrm{access}}$ means the global state is highly distinguishable while a specified physical receiver cannot approach that bound; reporting $\chi_H(T)$ alone as "what the horizon knows" silently promotes a capacity statement to an engineering claim. No value of χ_{collect} or $\Delta I_{\mathrm{access}}$ has been computed for any spacetime in this scaffold; Eq. (10) is registered as a required output, not reported as a number.

Finally, the pure equal-prior identities above are fragile. For mixed conditional states, unequal path priors, or a declared subregion rather than the full horizon algebra, $V^2 + D^2 = 1$ does not hold, χ_H must be computed from the actual mixed-state ensemble rather than from V alone, and $D(T)$ remains an operationally meaningful decision quantity only after the true reduced states, not a fitted overlap curve, are supplied.

5 The horizon state model traces to a declared algebra

Sections 4 and 3 treat $|h_0(T)\rangle$ and $|h_1(T)\rangle$ as if they already were the conditional states of a horizon subsystem. That step is not free, and this scaffold does not perform it. What follows registers the assumptions a horizon-state model must make explicit before Eqs. (5)–(10) may be applied.

The completed radiation has more than one home. The soft field sourced by the two branches can carry components that reach the horizon and components that reach null infinity; from a Rindler-type description the entangling quanta cross the horizon with negligible boost energy while from an inertial description the same process can be represented as radiation escaping to infinity [2]. These are alternative accounts of the same correlations where the derivations agree, not two independent decoherence channels to be added. A model that reports a horizon record must trace or restrict the full field state to a declared horizon subalgebra $\mathcal{A}_{\mathrm{horizon}}$; it may not infer horizon-accessible information from Alice's total visibility loss, since a photon reaching null infinity can destroy interference without leaving any distinguishable record behind the horizon, and, in principle, a conditional horizon state can be highly distinguishable in the global algebra while remaining difficult for any localized interior collector to access.

Global and local accounts must not be summed. Danielson, Satishchandran, and Wald's local two-point-function derivation reproduces the Schwarzschild decoherence rate directly from field correlators in Alice's laboratory, with the result depending on the choice of Unruh, Hartle–Hawking, or Boulware vacuum state [6]. Wilson-Gerow, Dugad, and Chen give an independent open-system derivation in terms of Ohmic friction and finite-temperature fluctuations [4]. Where these descriptions are derivations of the same physics, a provenance-tagged rate object (*global-horizon-flux* versus *local-correlator*) may be compared but must not be added; summing them without a proof of statistically independent physical origin would manufacture a decoherence rate that no single calculation predicts.

A completed record is not automatically an irrevocable one. Danielson, Kudler-Flam, Satishchandran, and Wald show that, given the radiation already past a horizon cross-section at some cutoff t_c, the future field configuration that best preserves branch coherence need not correspond to "stop perturbing everything now"; in some regimes, reopening and manipulating an already-recombined superposition further reduces final decoherence [7]. This scaffold therefore distinguishes a completed-protocol overlap $V(T)$, evaluated at a fixed final time with a fixed protocol, from an irreducible receipt quantity that optimizes over an admissible future-control family $\mathcal{A}(t_c)$:

$$D_{\rm irr}(t_c) = \inf_{u\in\mathcal{A}(t_c)} \tfrac{1}{2}\left\|\rho^u_{0,\rm H} - \rho^u_{1,\rm H}\right\|_1 , \tag{11}$$

with an associated deficit-threshold **receipt time** $T_{\rm receipt}(\delta)$ defined from $I_{\rm irr}(t_c) = 1 - h_2\left(\frac{1-D_{\rm irr}(t_c)}{2}\right)$ exactly as in Definition 3.1. $\mathcal{A}(t_c)$ must specify which controls Alice retains after t_c – their duration, energy, acceleration, localization, and boundary conditions – because an unrestricted $\mathcal{A}$ makes every record fictitiously erasable and a trivial $\mathcal{A} = \{\text{do nothing}\}$ collapses Eq. (11) into an ordinary decay threshold. Neither $D_{\rm irr}$ nor $T_{\rm receipt}$ has been evaluated for any spacetime here; no monotonicity of $D_{\rm irr}(t_c)$ in t_c is assumed, since a later cutoff may or may not grant a qualitatively new admissible control depending on how $\mathcal{A}(t_c)$ is chosen to shrink, and a failure of the expected ordering must be reported as a definition failure rather than corrected by construction.

Purity and equal priors are modeling choices, to be checked, not assumed. Sections 3–4 used $V^2 + D^2 = 1$, which holds only for pure, equal-prior, unrestricted-algebra conditional states. Restricting to $\mathcal{A}_{\rm horizon} \subsetneq \mathcal{A}_{\rm global}$ generally produces mixed reduced states even when the global state in Eq. (2) is pure, at which point $D(T)$ must be computed directly from the reduced density operators, $\chi_H(T)$ must be computed from the actual mixed ensemble, and Eq. (5) is not available as a shortcut. A horizon-state model that reports χ_H from a fitted $V(T)$ curve without first establishing that the relevant conditional states remain pure on the declared algebra has silently reintroduced the assumption it needed to justify.

Every quantity in this section requires, as an unresolved prerequisite, a boundary-resolved field calculation that this scaffold does not perform: a decomposition of the two-branch radiation state onto $\mathcal{A}_{\rm horizon}$ versus null infinity, a provenance tag on any imported rate, and an explicit admissible-control family for the receipt clock.

6 Some horizons have a recording dead zone

The exponential $V(T) = e^{-\Gamma T}$ used to state Definition 3.1 is a special case, not the generic behavior, and its rate Γ is never derived by the information layer – it is imported, with a declared spacetime, channel, and clock convention, from a specific field calculation. This section registers the comparison the numerical protocol (Section 7) is required to carry out; it computes none of the underlying rates itself.

Nonextremal Killing horizons. For nonextremal bifurcate Killing horizons in the electromagnetic and Klein–Gordon analogues treated by Gralla and Wei, the natural-unit overlap is $V(T) = \exp(-C\kappa T/2)$, so the adapter into Eq. (3) is $\Gamma = C\kappa/2$, with C a dimensionless decohering-flux coefficient and κ the surface gravity in the associated Killing-time normalization; the same reference evaluates C on the symmetry axis of a Kerr black hole [3]. Schwarzschild is the $a = 0$ special case of the same family.

Zero surface gravity. At $\kappa = 0$, Gralla and Wei find the scalar/Klein–Gordon analogue can cross over to power-law rather than exponential decay, $V(T) = (T/T_0)^{-C}$ for a reference time T_0; a threshold clock built on this law scales as a power of the deficit rather than logarithmically, and Definition 3.1 must be rederived from the power law rather than patched onto the exponential form [3].

Extremal electromagnetic screening. In the electromagnetic extremal-Kerr limit treated by the same reference, the external field does not penetrate the horizon, so the relevant decohering flux and the surface gravity both vanish, $C \to 0$, $\kappa \to 0$ [3]. This is a literature-established zero, attributed to Gralla and Wei; this scaffold does not recompute it and forbids relabeling it as "unknown" or as a positive rate obtained here.

Near-extremal quantum spectral gap. Biggs and Trezzi study near-extremal charged Reissner–Nordström black holes with a quantized near-horizon geometry and a single electromagnetically superposed charge in an opposite-dipole configuration. They find, through quartic order in the interaction and conjectured to all orders for that geometry, that a spin-induced gap in the black-hole spectrum forces the late-time decoherence rate to vanish when the energy above extremality lies below

$$E_b = \frac{G_N}{r_0^3} = \frac{\pi}{r_0 S_0} \tag{12}$$

in their natural-unit conventions, with r_0 a near-extremal radius scale and S_0 an associated entropy parameter; above the gap the quantum-corrected rate remains below the semiclassical prediction [8]. Equation (12) is quoted, not derived, and its zero is explicitly a 2026 preprint result established through quartic order for one stated charge configuration – it is not assumed to extend to arbitrary charge distributions or to the gravitational channel.

Ordinary-matter and warm-horizon comparison. Biggs and Maldacena show that finite-temperature ordinary dissipative matter, without any causal horizon, can produce decoherence that is qualitatively similar and, in a matched electromagnetic setting, quantitatively comparable to the horizon effect [5]. Wilson-Gerow, Dugad, and Chen's open-system Rindler description gives a rate set by Ohmic friction and finite-temperature fluctuations rather than by a black-hole-specific coefficient [4]. These are alternative rate providers, evaluated on their own physical footing, not deviations from a single "true" horizon rate.

Nonperturbative check. Batista, Landulfo, Mann, and Matsas give a nonperturbative Unruh–DeWitt treatment of the Danielson–Satishchandran–Wald mechanism for a gapless detector on uniformly accelerated superposed trajectories coupled to a massive scalar field, supporting the semiclassical exponential beyond leading order in that specific detector model [9]. It is not an electromagnetic Kerr calculation and is not evidence that any conclusion transfers across the extremal or spectral-gap boundary.

Definition 6.1 (Rate-provider status)**.** *For every point in a provider's parameter domain (mass, charge, spin, source geometry, distance, field channel, vacuum state, holding time), the numerical protocol records exactly one of four statuses:* (1) *rate derived and positive;* (2) *rate derived and zero;* (3) *approximation invalid;* (4) *rate not yet calculated. Status 4 must never be plotted as zero, and status 2 must never be replaced by a small positive floor for plotting convenience.*

Table 1 lists the rate providers this scaffold is committed to comparing once implemented. Every entry below is currently status 4 except the two literature-established zeros of Gralla–Wei's extremal Kerr limit and Biggs–Trezzi's sub-gap Reissner–Nordström limit, which carry status 2 by direct citation and are not recomputed here.

The horizon clock is therefore modular by construction. The information layer of Sections 3–4 consumes a conditional overlap or a pair of density operators with units and a validity domain attached; it does not make the six providers of Table 1 into one theory by placing their consequences on a shared axis, and a plot that does so without the status column has misrepresented theoretical ignorance and physical screening as the same color.

Provider	Clock	Status
Schwarzschild, EM/scalar	Killing	4
Kerr, EM (axis, Gralla–Wei)	Killing	4
Kerr, EM, extremal (screening)	Killing	2 (cited)
Reissner–Nordström, near-extremal, quantum gap	proper	2 (cited, sub-gap)
Rindler, open-system (warm horizon)	proper	4
Ordinary dissipative matter (no horizon)	laboratory	4

Table 1: Rate-provider registry. Status codes follow the four-way scheme above. No provider in this scaffold carries status 1; a status 1 entry may appear only after an implemented, reviewed field calculation.

7 The numerical protocol has gates, not a capacity map

The eventual numerical artifact is a versioned rate-to-record program. Its job is to expose every assumption between an imported field calculation and a sentence about learned information, not to solve quantum gravity. No item in this section has been executed against a horizon field calculation.

For each rate provider in Table 1, the configuration must record: the spacetime and horizon type; black-hole mass, angular momentum, charge, and an extremality measure; source charge or mass, branch separation, orientation, distance, opening time, holding time, and closing time; the electromagnetic, scalar, or linearized-gravity channel; the field state and its horizon/null-infinity boundary decomposition (Section 5); the time coordinate in which the rate is defined and its conversion to the laboratory clock; whether the conditional states are pure, mixed, or represented only by an asserted asymptotic overlap; the validity inequalities of the approximation and their numerical margins; and the paper version, equation identifier, unit convention, and any published correction (the local-correlator derivation, for instance, carries a documented post-publication sign correction to two equations [6]).

The program must first reproduce a provider's published overlap or decoherence curve inside its stated regime before computing D, p_{e}, χ_H, I_{dec}, or any threshold time from it. Only after that reproduction check passes may Eqs. (5)–(9) be applied. A separate, clearly labeled "exploratory" module is required for the finite-time admissible-control optimization behind $D_{\mathrm{irr}}(t_c)$; it may not be treated as validated until it reproduces benchmark cases from the optimal-purification literature [7].

Every saved output must carry a configuration hash, random seed where applicable, code commit, dependency lock, hardware/software environment record, convergence diagnostics, and confidence intervals or exact-identity flags as appropriate. Plots may be generated only from saved machine-readable output; no plotting code may interpolate, smooth, or extrapolate across a status-4 (not-yet-calculated) region of Table 1.

The registered grid varies the deficit δ, the assumed clock convention, and the rate-provider identity; exact starting values live in `simulation_config.json` and are feasibility targets, not a claim that every combination is numerically tractable. The program supplied with this scaffold, `simulate_horizon_record.py`, currently validates only the analytic prerequisites listed in Appendix B: binary-entropy identities, the pure-state relation, the Holevo-bound ordering of Proposition 4.1, round-trip threshold inversion, and an audit that no rate-provider entry in the configuration is marked as a positive rate computed by this project. Its `-simulate` option exits with a nonzero status and a message naming exactly what remains missing. Therefore this section contains no capacity map, no redundancy-style scaling law, and cites no generated figure.

8 The limits a receipt time must reproduce before it is believed

At $\Gamma = 0$ – whether because a provider genuinely has no coupling, because the extremal-screening limit applies, or because the near-extremal spectral gap of Eq. (12) is not exceeded – every accessible conditional state is orientation-independent in the relevant channel and admissible-threshold recording must be reported as never occurring through that channel, not as a small positive number. As $T \to \infty$ with $\Gamma > 0$ fixed, $\chi_H(T) \to 1$ bit: the ceiling is exactly one path bit, never more, because the underlying alternative is binary. No claim of unbounded information accumulation is compatible with the two-branch model of Eq. (2).

The ordinary-matter sufficiency countermodel. Biggs and Maldacena's result is the strongest and most mundane rival account: a horizon is one thermal, dissipative quantum system among others, and matched ordinary conducting matter can be quantitatively comparable [5]. Wilson-Gerow, Dugad, and Chen's warm-horizon derivation makes the same point from the horizon side [4]. This scaffold accepts the demotion in advance: its goal is not to prove horizons uniquely learn anything, but to price what any conditional-environment state allows a specified reader to infer, horizon or not. The decisive test (Section 7) is a negative control – the same opening/holding/closing trajectory evaluated beside a horizonless, temperature- and low-frequency-response-matched dissipative body. If the matched environments give matched conditional states, every clock in this paper must give matched threshold times; a mismatch would indicate an error in the translation, not new physics.

The local-description sufficiency countermodel. If the local two-point-function derivation reproduces the same rate entirely from correlators available in Alice's own laboratory [6], no global horizon account is strictly needed to predict her visibility loss. That does not erase the horizon-side description; Section 5 already forbids treating the two derivations as independent evidence to be summed, and this countermodel sharpens why: if local and global accounts are the same physics viewed differently, at most one of them may contribute a rate to any given calculation.

The decoder-access countermodel. A large $\chi_H(T)$ on the full horizon algebra may require a collective measurement no physically situated collector can perform. If every modeled causal collector's $\chi_{\mathrm{collect}}(T)$ (Eq. (10)) remains near chance while $\chi_H(T)$ is large, then saying a horizon observer "learned" the path overstates a capacity bound; the access gap, not the global ceiling, is the reportable quantity in that regime.

The receipt-instability countermodel. If $T_{\mathrm{receipt}}(\delta)$ from Eq. (11) fails to converge as the admissible future-control family $\mathcal{A}(t_c)$ is enlarged toward the physically justified limit, the receipt clock is not a property of the horizon and must be reported as protocol-relative or withdrawn, following the optimal-purification finding that reopening a closed superposition can reduce final decoherence [7].

The extremal/spectral-gap countermodel. If the electromagnetic extremal-Kerr screening or the near-extremal Reissner–Nordström spectral gap applies to the modeled configuration, no finite holding time produces a record through that channel regardless of how the information layer is tuned [3, 8]. A null region is a physical prediction of the imported theory to be preserved on the parameter map, not an inconvenient boundary to interpolate across.

The role-reversal countermodel. Arrasmith, Albrecht, and Zurek's superposed-black-hole problem shows that swapping which system is superposed and which is environment changes the conditional states, the rate, and the validity domain entirely [16]. Nothing derived for a superposed charge near a black hole transfers to a superposed black hole without an independent re-derivation.

Each rival account has a decisive, preregistered comparison rather than a rhetorical dismissal: the ordinary-matter comparison requires a matched negative control; the local-description comparison requires a provenance-tagged rate registry that forbids double counting; the decoder-access comparison requires an explicit causal-collector model; the receipt comparison requires convergence under an enlarging control family; and the extremal comparison requires the four-status domain map of Table 1 to be preserved rather than smoothed. Failure on any one of these narrows or kills the proposed translation.

9 How this could be wrong

The proposed translation is rejected outright if a boundary-resolved field calculation shows the horizon-restricted conditional states remain identical (or their distinguishability vanishes) in a regime where the reduced exponential model of Section 6 assigned a positive Γ: then there is no horizon which-path record through that channel, and any finite threshold time computed from the reduced model is an artifact of an unjustified simplification, not a physical prediction. It is withdrawn as non-novel if the completed 60-source priority search finds that this Holevo/decision/receipt translation, or an equivalent operational metric, already exists in the literature under a different name, or if every preregistered comparison in Section 6 and Section 8 reduces to a monotone relabeling of $V(T)$ that changes no ranking, regime boundary, or proposed test.

More local failures are equally serious and each has a planned diagnostic. Clock mismatch: multiplying a Killing-time rate by a proper-time duration, or vice versa, can manufacture an apparent redshift enhancement or suppression with no physical content; every rate object must carry a clock tag, and the protocol refuses to multiply mismatched tags. Provenance double-counting: summing a global-horizon-flux rate and a local-correlator rate for the same physical process, when they are alternative derivations rather than independent channels, would overstate Γ and understate every threshold time. Status collapse: treating a status-4 (not-yet-calculated) parameter region as status-2 (derived zero), or treating a literature-established status-2 region as a small positive floor for plotting convenience, converts either theoretical ignorance or genuine screening into a fabricated result. Ceiling-for-decoder substitution: reporting $\chi_H(T)$ as what a specific observer learned, rather than as an upper bound that a specific measurement may fail to approach, silently deletes the Fuchs–Caves accessible-information gap [15]. Purity misuse: applying $V^2 + D^2 = 1$ after restricting to a proper horizon subalgebra, where the reduced conditional states are generically mixed, would produce a χ_H or D value with no operational meaning. Non-monotone receipt clock: if enlarging $\mathcal{A}(t_c)$ does not shrink $D_{\mathrm{irr}}(t_c)$ consistently, an implementation that assumes monotonicity by construction would silently misreport T_{receipt}.

The most serious conceptual failure is metaphorical overreach: describing χ_H, I_{dec}, or D_{irr} becoming large as “the horizon learning” or “the path becoming real” would smuggle an interpretation of quantum mechanics into what is, at every step in this paper, a state-discrimination calculation for a declared receiver with declared access. Decoherence is not collapse; a reduced path state loses interference because phase information resides in correlations with a field, and nothing in Sections 3–5 selects among interpretations of what, if anything, that entails about which branch is “real.” The paper’s restraint on this point is a feature of the definitions, not a hedge to be relaxed once numbers exist.

10 Five debts this construction has not yet paid

The threshold times defined in Definition 3.1 are at present definitions with a derivation behind them and no computation under them. Five debts stand between that and a number. They are listed in the order they must be settled, because the later ones cannot be attempted until the earlier ones close.

Debt 1 — priority. The sixty-source hard-floor review and the accompanying title, abstract and equation-level priority search remain incomplete. No sentence in Section 1 may be upgraded from proposal to result language while that is true, and the construction is withdrawn if the search shows the same operational separation already established.

Debt 2 — the boundary decomposition. Section 5 assumes a horizon/null-infinity split and an admissible-control family. Neither is implemented for even one solvable rate provider. Until at least one exists, no value of χ_{collect}, D_{irr} or T_{receipt} may be computed, quoted, or plotted — including as an illustrative figure, since an illustrative curve of an uncomputed quantity is indistinguishable to a reader from a result.

Debt 3 — the import path. Section 7 requires a covariant density-operator import for a rate provider and a finite-time purification module. Both are specified and neither is built. External subject-matter review is also outstanding, and is a precondition rather than a formality: the ordering of the Holevo bound in Proposition 4.1 is the kind of step where a reviewer’s objection would change the construction rather than its presentation.

Debt 4 — the failure record. Section 9 anticipates how the argument could fail. It may be rewritten only once real field-calculation residuals, convergence failures or null results exist, and every kill criterion and countermodel test listed in Section 6 must survive that rewrite. Replacing a preregistered criterion with a post hoc justification would convert a falsifiable claim into an unfalsifiable one without any visible edit to the conclusions.

Debt 5 — the ledger. Appendix B records what the scaffold program actually verifies today. Hashes for the solver, its configuration, its output and the figure-generation commands are appended once those artefacts exist.

Debts 1 and 3 are partly matters of scholarship and review. Debts 2 and 3 require code that does not exist, and debts 4 and 5 are downstream of that code. The extremal and Kerr limits remain outside the registered scope entirely.

11 A ledger has been registered, not a horizon read

This scaffold defines an operational translation layer: a conditional horizon-state overlap becomes a minimum-error decision information, a Holevo capacity ceiling, and – pending a finite-time admissible-control model – an irreducible receipt distinguishability, with every rate object required to carry a declared spacetime, channel, and clock convention. It derives the pure-state identity linking visibility and trace distance, the Holevo-bound ordering between the ceiling and decision thresholds, and the binary-entropy inversion that turns an assumed exponential overlap into two deficit-threshold times, reporting that inversion explicitly as an analytic identity rather than a physical calculation.

It reports no horizon decoherence rate, no capacity map over mass, spin, charge, or holding time, no receipt time, and no verdict on any of the six countermodels in Section 8. Those require the unresolved boundary-resolved field calculation, the covariant density-operator import for each rate provider, the finite-time purification module, the completed source review, the implemented and reviewed numerical protocol, and external subject-matter review. If the translation survives that program, it would connect an inherited horizon-decoherence mechanism to an explicit statement of which bit, for which reader, by which deadline, under which remaining control. If it does not survive – because the field calculation finds no boundary-resolved record, because every comparison in this paper turns out to be a relabeling of $V(T)$, or because the ordinary-matter negative control cannot be distinguished from the horizon case – the inherited decoherence mechanism and the established Helstrom/Holevo bounds it draws on remain entirely intact.

The scientific deliverable at this stage is a falsifiable registration of the translation, together with a functioning refusal to report a number that has not been calculated.

δ	V at χ_H thresh.	ΓT_χ	V at dec. thresh.	$\Gamma T_{\rm dec}$
0.1	0.367961	0.999777	0.226435	1.485297
0.01	0.117605	2.140426	0.058633	2.836452
0.001	0.037229	3.290676	0.016143	4.126270

Table 2: Exact binary-entropy threshold inversion for $V(T) = e^{-\Gamma T}$. Reproduced from the companion landing article's identical calculation and re-verified by `simulate_horizon_record.py -self-check` (Appendix B). Every entry is a root of a transcendental equation in the dimensionless variable ΓT; none is a physical time, since Γ is not derived by this scaffold.

A Binary-entropy threshold inversion

This appendix inverts Eqs. (6) and (9) under the assumed exponential overlap $V(T) = e^{-\Gamma T}$ to obtain the dimensionless products $\Gamma T_\chi(\delta)$ and $\Gamma T_{\rm dec}(\delta)$. Every number in this appendix is an exact numerical root of a fixed transcendental equation in one real variable; none of them uses, assumes, or estimates a physical value of Γ, a spacetime, or a channel. The appendix is therefore a mathematical identity report, not a horizon-record calculation, and Table 2 may not be cited as a physical prediction until Section 6's rate import is implemented and reviewed.

Substituting $V = e^{-x}$ with $x = \Gamma T \geq 0$ into Eq. (6) gives

$$\chi_H(x) = h_2\left(\frac{1+e^{-x}}{2}\right), \tag{13}$$

which is continuous, equals 0 at $x = 0$, and increases monotonically to 1 as $x \to \infty$, since e^{-x} decreases monotonically from 1 to 0 and $h_2\left(\frac{1+V}{2}\right)$ decreases monotonically in V on $[0,1]$ (Section 3). The threshold equation $\chi_H(x) = 1-\delta$ therefore has a unique solution $x = \Gamma T_\chi(\delta)$ for every $\delta \in (0,1)$, found here by bisection to a tolerance of 10^{-10} in x. The same argument, with χ_H replaced by $I_{\rm dec}(x) = 1 - h_2\left(\frac{1-\sqrt{1-e^{-2x}}}{2}\right)$, gives a unique $\Gamma T_{\rm dec}(\delta)$.

Table 2 confirms, as a numerical instance of Proposition 4.1, that $\Gamma T_\chi(\delta) < \Gamma T_{\rm dec}(\delta)$ at every tabulated δ: the Holevo ceiling always crosses a given deficit before the minimum-error decision does, because $I_{\rm dec} \leq \chi_H$ pointwise. The gap is not a fixed multiplicative factor – at $\delta = 0.1$ the ratio $T_{\rm dec}/T_\chi \approx 1.485$; at $\delta = 0.001$ it is ≈ 1.254 – so no single "fudge factor" converts one clock into the other across deficits, and a reduced model that reports only one threshold has discarded information the other clock would have supplied.

No threshold in Table 2 is a constant of nature. The choices of equal path priors, a single binary use of the channel, a bit-valued score, and a tolerated deficit δ all belong to the question an observer is asking of the conditional states, not to the underlying physics. A different loss function, a repeated decoder, or an asymmetric decision cost would produce a different pair of clocks from the same $V(T)$. Appendix B lists the self-check that re-derives this table on every run and the audit that keeps it from being mistaken for a physical result.

B Computation and test ledger

The scaffold program `simulate_horizon_record.py` implements no field-theoretic state propagation and no horizon calculation. In self-check mode it performs the following exact/numerical sanity tests:

1. endpoint, symmetry, and monotonicity checks for the binary entropy function $h_2(p)$ on $[0, 1]$;
2. the pure-state identity $V^2 + D^2 = 1$ across a grid of $V \in [0, 1]$;
3. endpoint and monotonicity checks for $\chi_H(V) = h_2\left(\frac{1+V}{2}\right)$, including $\chi_H(1) = 0$ and $\chi_H(0) = 1$;
4. Helstrom success/error consistency, $p_{\text{succ}} + p_{\text{e}} = 1$, with $p_{\text{succ}} = \frac{1}{2}$ at $V = 1$ (chance) and $p_{\text{succ}} = 1$ at $V = 0$ (perfect discrimination);
5. the Holevo-bound ordering of Proposition 4.1, $\chi_H(V) \geq I_{\text{dec}}(V)$, across the same grid;
6. the $\Gamma = 0$ limit, $V \equiv 1 \Rightarrow \chi_H \equiv 0$, confirming the no-record regime is not silently converted into a finite number;
7. round-trip inversion of $\chi_H(\Gamma T_\chi(\delta)) = 1 - \delta$ and $I_{\text{dec}}(\Gamma T_{\text{dec}}(\delta)) = 1 - \delta$ against the registered δ grid, to a stated numerical tolerance, reproducing Table 2;
8. a registered-threshold filter confirming every δ in `simulation_config.json` lies in $(0, 1)$ and yields a finite, positive, correctly ordered pair $(\Gamma T_\chi, \Gamma T_{\text{dec}})$;
9. a rate-provider honesty audit: every entry in the configuration's provider registry is checked to carry status *not_yet_calculated* or a literature-cited *derived_zero_by_literature* status with a citation key that resolves against the bibliography; no entry may carry a computed-positive-rate status, since this project has not implemented a field calculation.

The saved output labels these as scaffold validation. The program returns a nonzero exit status and a blocking message for `-simulate`, naming explicitly the horizon field calculation, the Kerr/extremal limit treatment, the finite-time purification module, the solver dependency lock, the 60-source review gate, and external review as the incomplete prerequisites. A later implementation must not silently reuse the program name while relaxing these gates: it must add state/channel/estimator tests appropriate to an actual field calculation, write a versioned output schema, and update the dependency lock and environment record before any provider may be promoted to status 1.

References

[1] D. L. Danielson, G. Satishchandran, and R. M. Wald (2022). Black holes decohere quantum superpositions. *International Journal of Modern Physics D* 31, 2241003. https://arxiv.org/abs/2205.06279. DOI https://doi.org/10.1142/S0218271822410036.

[2] D. L. Danielson, G. Satishchandran, and R. M. Wald (2023). Killing horizons decohere quantum superpositions. *Physical Review D* 108, 025007. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.108.025007.

[3] S. E. Gralla and H. Wei (2024). Decoherence from horizons: general formulation and rotating black holes. *Physical Review D* 109, 065031. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.109.065031.

[4] J. Wilson-Gerow, A. Dugad, and Y. Chen (2024). Decoherence by warm horizons. *Physical Review D* 110, 045002. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.110.045002.

[5] A. Biggs and J. Maldacena (2024). Comparing the decoherence effects due to black holes versus ordinary matter. *arXiv preprint.* https://arxiv.org/abs/2405.02227.

[6] D. L. Danielson, G. Satishchandran, and R. M. Wald (2025). Local description of decoherence of quantum superpositions by black holes and other bodies. *Physical Review D* 111, 025014. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.111.025014.

[7] D. L. Danielson, J. Kudler-Flam, G. Satishchandran, and R. M. Wald (2025). How to minimize the decoherence caused by black holes. *Physical Review D.* https://journals.aps.org/prd/abstract/10.1103/67vv-km43.

[8] A. Biggs and S. Trezzi (2026). Not all black holes decohere quantum superpositions. *arXiv preprint.* https://arxiv.org/abs/2605.23880.

[9] L. B. N. Batista, A. G. S. Landulfo, R. B. Mann, and G. E. A. Matsas (2026). Nonperturbative Danielson-Satishchandran-Wald decoherence with Unruh-DeWitt detectors. *arXiv preprint*. https://arxiv.org/abs/2605.00956.

[10] A. Belenchia, R. M. Wald, F. Giacomini, E. Castro-Ruiz, Č. Brukner, and M. Aspelmeyer (2018). Quantum superposition of massive objects and the quantization of gravity. *Physical Review D* 98, 126009. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.98.126009.

[11] B.-G. Englert (1996). Fringe visibility and which-way information: an inequality. *Physical Review Letters* 77, 2154–2157. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.77.2154.

[12] S. Dürr, T. Nonn, and G. Rempe (1998). Fringe visibility and which-way information in an atom interferometer. *Physical Review Letters* 81, 5705–5709. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.81.5705.

[13] A. S. Holevo (1973). Bounds for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission* 9, 177–183. https://www.mathnet.ru/eng/ppi903.

[14] C. W. Helstrom (1969). Quantum detection and estimation theory. *Journal of Statistical Physics* 1, 231–252. https://doi.org/10.1007/BF01007479.

[15] C. A. Fuchs and C. M. Caves (1994). Ensemble-dependent bounds for accessible information in quantum mechanics. *Physical Review Letters* 73, 3047–3050. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.73.3047.

[16] A. Arrasmith, A. Albrecht, and W. H. Zurek (2019). Decoherence of black hole superpositions by Hawking radiation. *Nature Communications* 10, 1024. https://www.nature.com/articles/s41467-019-08426-4.