
Black-Hole Entropy under GG-Theory

The **GDOS** Reading of the Horizon as an **OGGEP** Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity

Katsushi Arisaka

*Department of Physics and Astronomy, Electrical and Computer Engineering,
University of California, Los Angeles, Los Angeles, CA 90095, USA*

Contact: arisaka@physics.ucla.edu

Abstract

Black holes are the strangest objects in nature. Fifty years ago Stephen Hawking discovered that they are not perfectly black: they glow faintly, and over almost unimaginable spans of time they evaporate away. This discovery created the deepest puzzle in modern physics. Whatever falls into a black hole carries detailed information about what it was; yet the faint glow that leaks back out looks featureless — the same for every black hole of the same size. If the glow truly carries nothing, then that information has been erased, and the erasure of information is something the laws of physics are not supposed to allow. For half a century no one has agreed on how the books are balanced.

This paper offers a resolution built on GG-Theory, a framework in which the world we experience is the surface of a slightly larger six-dimensional space. In that picture the edge of a black hole — its horizon — is not a wall but a doorway: matter that crosses it is not destroyed but carried into directions we cannot see, and the faint glow is simply how that crossing looks from the outside. Nothing is lost; it is only put out of reach. From this single idea, and without adjusting any free number, the paper recovers the famous formula for how much a black hole can hold, explains why the glow should eventually carry the hidden information back out, and shows why the argument that a traveler falling inward must meet a wall of fire does not go through. The same reasoning extends to the horizon of the expanding universe as a whole. The claims are set out so that future telescopes and gravitational-wave detectors can put them to the test.

Executive Summary

The puzzle. A black hole forms from matter in a definite quantum state. Hawking showed in 1974 that it then radiates, and the radiation is thermal — as random and featureless as the glow of a warm coal. Carried to its conclusion, this means a sharp, information-rich beginning evolves into a featureless end, a “pure-to-mixed” transition that ordinary quantum mechanics forbids: quantum evolution is unitary, and unitary evolution never destroys information. Don Page sharpened the demand in 1993: if the process is unitary, the entanglement entropy of the radiation must rise and then *fall* back to zero — the “Page curve” — rather than rise forever. For fifty years the central question has been what changes Hawking’s calculation so that this turnover happens.

What the standard answers achieved, and what they left out. The most successful modern attack uses the AdS/CFT correspondence and the 2019–2020 islands formula, which reproduce the Page curve in controlled, lower-dimensional toy models. These are real achievements, but they carry two limitations. They live in anti-de-Sitter universes with a supersymmetric boundary theory — not the expanding, Standard-Model universe we inhabit — and they recover the *shape* of the curve without deriving, from a microscopic count, the single most famous number in the subject: the factor of $1/4$ in the Bekenstein–Hawking entropy $S = A/4G_N$. Counting black-hole microstates from first principles had been possible only for special supersymmetric, extremal black holes far from anything in the sky.

The GG-Theory picture. GG-Theory treats the four dimensions we see as the boundary of a six-dimensional bulk: the familiar spacetime plus two non-spatial directions, a “scale” direction and a compact “phase” circle.¹ A horizon is the place where a four-dimensional observer loses causal access to part of that bulk. Describing what such an observer sees is a standard operation in quantum theory — a partial trace, here called the OGGE projection — and its output is automatically a mixed, thermal state. Hawking radiation is precisely this projection of a pure six-dimensional state; the bulk stays unitary even though the boundary description looks lossy. The black hole is the one place in nature where this projection becomes geometrically visible, which is what makes it the cleanest test of the whole framework.

The central new result: the exact $1/4$ from induced gravity. In GG-Theory gravity is not a separate force with its own constant; Newton’s constant is *induced*, generated entirely by the bulk fields. The consequence is decisive: the entanglement entropy across a horizon is not a small correction sitting on top of $A/4G_N$ — it *is* $A/4G_N$. A standard calculation (the conical-replica method) then fixes the coefficient in closed form, $1/4 = 4\pi/16\pi$, with no adjustable parameter, no supersymmetry, and no extremality, for the ordinary rotating and charged black holes that actually

¹Throughout the cohort, this “warped” six-dimensional bulk is read on the field-theory side of the holographic dictionary: the scale direction is the renormalization-group axis $r = \ln(E_{\text{GUT}}/\mu)$, not a Randall–Sundrum spatial dimension, and “warped” denotes RG running; the explicit holographic-RG construction is given in [Arisaka \(2026, A2-GG6\)](#). (This is distinct from the genuine anti-de Sitter geometries of the black-hole comparisons in this paper.)

exist. This is a first-principles derivation of the coefficient exactly where the older microscopic counts could not reach. One honest gap remains, and we flag it plainly: turning the induced Newton constant into a single pure number requires summing over all the fields of the theory, a computation we set up but do not complete.

Why the paradox dissolves. Once the horizon is a projection rather than a wall, the rest follows with little extra machinery. The Page curve’s shape is forced: for a pure total state the entropy of the radiation tracks the entropy of what remains, rising then falling. When the turnover comes is a separate question, and the paper answers it with a named mechanism — the expiry of port freshness at the Page time — that rests on one imported typicality assumption, flagged where it is used; the quantitative curve remains open work. No firewall is needed, because the “interior partner” that drove the firewall argument is not an independent four-dimensional system but part of the shared bulk; the apparent contradiction was an artifact of insisting on a four-dimensional-only description. The interior singularity becomes a coordinate artifact of that same projection, smooth in six dimensions — just as the horizon itself looked singular in Schwarzschild’s coordinates until better ones were found.

What is being counted, and the universe’s own horizon. The microstates get names at the same stroke. The entropy counts the configurations of the bulk’s compact phase circle over the horizon — the hidden-fiber holonomies, one register per Planck-sized cell, the same fiber whose windings carry electric charge elsewhere in the cohort: a horizon of area A supports $e^{A/4\ell_P^2}$ distinguishable fiber configurations. And because the argument never uses a box or a wall at infinity — the observer, not an asymptotic boundary, anchors the description — it moves without modification to the horizon of the accelerating universe itself: the de Sitter entropy, the largest entropy in nature at $\sim 10^{122}$, is the same theorem with the observer on the inside, carried at the same graded register.

The second law gets a direction and a price. The generalized second law — horizon area plus outside entropy never decreases — also stops being an analogy here. Its direction is inherited from the cohort’s arrow of time, and its production is computed exactly in the paper’s tractable class: the entropy generated when a black hole swallows something equals the statistical distance (a relative entropy) between what was thrown in and what the horizon’s thermal face would have preferred — distance to equilibrium spent, never minted.

What it predicts, and how to test it. For real astrophysical black holes the corrections to Einstein gravity are fantastically small — a part in 10^{46} for stellar-mass holes, 10^{66} for the giants imaged by the Event Horizon Telescope — so the framework predicts pure Kerr geometry to far below any foreseeable precision. Its falsifiable content is therefore carried by the *null*: any confirmed departure from Kerr in a black-hole shadow or a gravitational-wave ringdown would refute it. Crucially, the same machinery is the engine of the companion papers, so the black-hole account can also be killed from outside black-hole physics — by the dark-matter and collider predictions of the cohort. A framework that can be shot at from a neighboring field is the stronger for it.


What the paper contains, and how it is written. The argument is carried by named, graded results: five theorems — the exact $1/4$; the Page turnover as a ledger branch switch; the rotating anchor; the collapse channel; and the thermal clock face, which turns the field’s continuation “tricks” into theorems — five lemmas (when the information return begins; the second law’s direction, with its production computed exactly in closed form; the tower identity behind the species sum; the forced port; and the Euclidean shore), and a register that grades every claim as derived, conditional, or open, each with a pre-registered way to kill it. An experiments Part scores the record — the gravitational-wave ringdown and area-theorem tests, the Event Horizon Telescope shadows — and dates the decade ahead with verdicts declared before the instruments report. The technical appendices are a graduate machine room, rebuilt as a graduate course: fourteen Step-graded chapters deriving the standard material from scratch — from the Schwarzschild geometry to the islands era — with problems and complete worked solutions throughout, reference dictionaries in front, and audit registries behind. And the objections Part opens with thirty questions from the general reader, answered at length in plain English, before fifty-three technical objections, organized by referee community, are answered in the open — concessions and open commitments said in those words.

A closing reflection. Seen this way, the information paradox was never really about information. It was about which description of nature we take to be the whole story. A four-dimensional observer, watching a black hole evaporate, sees a paradox only because the four-dimensional picture is a shadow — a partial trace of a larger, unitary, six-dimensional reality. The horizon is not an edge of existence but the edge of one observer’s access; what crosses it is displaced, not destroyed. That lesson does not stop at black holes. The same structure — a complete reality seen only in projection — underlies how a measurement yields a definite outcome, how heat flows in one direction, and how any open system loses track of its surroundings. The black hole is simply the most dramatic place where the universe reminds us that what we can see is not all there is.

Keywords: Black holes; event horizon; Hawking radiation; Bekenstein–Hawking entropy; information paradox; Page curve; firewalls; AMPS; ER=EPR; islands formula; AdS/CFT; Ryu–Takayanagi; replica wormholes; induced gravity; Event Horizon Telescope; LIGO ringdown; primordial black holes; GG-Theory; GDOS; OGGEP; 6D bulk.


The Axiom You Cannot See in a Laboratory, Made Visible by a Horizon

Lab: OGGEF invisible




Trace operative but undirected; observer could “in principle” resolve further

Black hole: forced



Trace forced by causal structure; no technology can recover the interior

Cosmos: forced




*modes beyond $r = c/H$
Trace forced by cosmic expansion; outside recedes faster than light*

Why this paper exists. The same partial trace, in three settings. In an ordinary laboratory (left) AX-3 operates everywhere and is operationally undirected — the trace runs over non-spatial fiber modes and no experiment reaches past it. At a black hole (center) the same trace is *forced*, by causal structure rather than by any limitation of technology: no instrument can recover the interior. In the cosmos (right) it is forced again, by expansion, for everything beyond the de Sitter horizon. In the middle and right panels the constitutional content of AX-3 stops being an axiom one accepts and becomes a geometry one can point at, and the black hole is the cleanest of the three. In full: Figure 8, p. 57.

Where the $1/4$ Comes From, and Why Nothing Was Fitted to Get It

Euclidean time τ (replica cover $\mathcal{M}_n$)



$$\int_{\mathcal{M}_n} \sqrt{g} R \supset 4\pi(1-n) A$$

$$\Rightarrow S_{\text{BH}} = \frac{A}{4G_N}, \quad \frac{1}{4} = \frac{4\pi}{16\pi}$$

The number that is usually quoted and rarely derived. On the replica cover the Euclidean-time circle is pinched at the horizon with deficit angle $2\pi(1-n)$; the curvature concentrated on the horizon contributes $4\pi(1-n)A$ to $\int \sqrt{g} R$, and fed through the induced Einstein–Hilbert action the replica formula returns $S_{\text{BH}} = A/4G_N$. The coefficient is the ratio of two normalizations: $1/4 = 4\pi/16\pi$. *Neither number knows anything about black holes.* The 4π is fixed by the geometry of spheres, the 16π by demanding Newton’s law in the weak-field limit, and both were set by consistency requirements decades before anyone computed a horizon entropy. No free parameter enters, no regularization choice survives, and supersymmetry is never used. In full: Figure 9, p. 70.

Ten Ways to Retire This Reading, and When Each One Becomes Decidable

Where this reading is exposed. Ten criteria, each with the observable, the deviation threshold at which it counts, and when it becomes decidable. The last column is the honest one, and it sorts the ten by reach. Four are live instrument rows on astrophysical data — the shadow, the ringdown, the echoes, the LISA multipoles — and the coming decade will keep squeezing them. One fires only on a well-populated primordial mass function, never on a single detection. Four are registered but presently beyond instrumental reach, and they are on the books because the register requires the exposure to be stated even where no instrument yet exists. One is internal: the species sum, the only row whose firing date the theory itself controls. Nothing here waits on a definition; every threshold is a number. In full: Table 9, p. 114.

Label	Observable	Deviation threshold	Timeline
F-BH-1	EHT shadow	$> 5\%$ deviation from Kerr on M87* or Sgr A*	Ongoing, ngEHT 2030s
F-BH-2	GW ringdown	$> 1\%$ QNM deviation from Kerr in LIGO/LISA sample	Ongoing
F-BH-3	GW echoes	Post-ringdown signal at $\Delta t \sim M \log M$	Ongoing analysis
F-BH-4	PBH ratio structure	A well-populated PBH mass function whose peak <i>spacing</i> is inconsistent with the e^7 scale ladder (model-level; a single detection does not trigger it)	Ongoing
F-BH-5	Hawking spectrum	Deviation from thermal form at observable frequencies	Near-Planck BHs required; unfeasible
F-BH-6	Information loss	Demonstration of $ \psi\rangle \rightarrow \rho_{\text{mixed}}$ in gravitational setting	Theoretical consistency; no direct test
F-BH-7	Bekenstein coefficient	$S_{\text{BH}} = cA/(4G_N\hbar)$ with $c \neq 1$	Theoretical; no direct test
F-BH-8	Firewall signature	High-energy quanta at horizon accumulating in stationary BH	Accessible only with ultrafine BH spectroscopy
F-BH-9	Species sum	The sign-resolved $\sum_i \chi_i$ over the GG-A9-SM spectrum with the Kaluza–Klein tower fails to reproduce the $1/G_N$ of Step 1 (Eq. (179); L-BH-Tower; internal)	Theory-side; decidable by a dedicated computation, no instrument required
F-BH-10	EMRI multipoles	Confirmed non-Kerr multipole structure at the 10^{-4} class in a LISA extreme-mass-ratio inspiral	LISA era, mid-2030s

Overview of Contents

Part I. Black Holes, 1915–2026	19
Part II. The Inherited Framework: GG-Theory in One Part	36
Part III. The Open Problems	47
Part IV. The Horizon as Port	54
Part V. The Exact $1/4$	66
Part VI. Information: Page Curve, Firewalls, and the Interior	76
Part VII. Quantitative Predictions	98
Part VIII. Experiments — Past, Now, and Future	109
Part IX. Summary and Comparison	117
Part X. Discussions	131
Part XI. Conclusions	138
Part XII. Technical Appendices	149
Part XIII. Common Objections and Responses	212
Part XIV. References	238

List of Figures

1	A century of black-hole physics, 1915–2026.	19
2	The open-system funnel shared across the GG-Theory cohort.	21
3	The twenty-two-paper cohort map.	32
4	Roadmap of the fourteen Parts of GG-C3-BH.	34
5	6D bulk and its OGGEp projection to the 4D boundary.	39
6	Zero-parameter derivation tree of the black-hole module.	44
7	The horizon as an OGGEp port.	55
8	The horizon as the geometrically-visible OGGEp port: three settings.	57
9	The conical-deficit origin of the exact $1/4$	70
10	Information loss as geometric distance.	78
11	The Page curve as a branch switch.	82
12	Conformal (Penrose) diagram of an evaporating black hole as an OGGEp port.	86
13	The black-hole life cycle in the two-regime grammar.	95
14	The kill-test timeline, drawn.	113
15	Penrose diagram of the maximally extended (eternal) Schwarzschild solution.	156
16	The Schwarzschild effective potential (86) for two values of the specific angular.	157

List of Tables

1	Working map from the horizon’s open problems (Part III) to the most relevant inherited. . .	45
2	The real-bulk series entries held by this module.	64
3	The exact $1/4$ across the stationary black-hole family.	73

4	The four laws as mechanics, thermodynamics, and GDOS/OGGEP statements.	89
5	The three links of the second law’s direction, their owners, and their status in this paper. . .	90
6	Worked black-hole numbers (order of magnitude).	99
7	GG-Theory quantitative predictions for black-hole observables.	106
8	The dated falsifier timeline of the black-hole module.	112
9	Black-hole-specific kill criteria.	114
10	Cumulative evidence for the GG-Theory program as of July 2026, including the black-hole qualitative.	115
11	Classical puzzles of black-hole physics and their resolution in GG-Theory.	119
12	Six frameworks for black-hole physics compared across twelve desiderata.	120
13	Comparative landscape of major black-hole information approaches.	121
14	The leading frameworks against the runtime criteria.	126
15	Dictionary translating 4D black-hole concepts to their 6D bulk origin in GG-Theory.	127
16	GG-Theory axioms and postulates used in the black-hole derivations.	128
17	Single-row summary of the 1975–2026 information-paradox debate.	129
18	A compact falsification dictionary for the black-hole module.	136
19	The Bardeen–Carter–Hawking correspondence (Bardeen et al., 1973), as it stood.	162
20	Island/qES terms and their GG-Theory/GDOS readings.	192
21	The internal source ledger.	207
22	The compact inheritance table: one row per companion, consumption only — this module. .	208
23	The consolidated benchmark audit.	209
24	The named-statement deck.	210

Contents

Executive Summary	2
List of Figures	8
List of Tables	8
Preface	17
 Part I. Black Holes, 1915–2026	 19
1 The central claim: black-hole entropy and the information channel as an open-system observable	20
2 Einstein’s 1915 field equations	20
3 Schwarzschild 1916: the first exact solution	20
3.1 The Einstein-Rosen bridge	22
3.2 Physical status of the interior	22
4 The golden age, 1958–1975	22
4.1 Finkelstein 1958: the horizon is a coordinate artifact	22
4.2 Kruskal 1960: the maximal extension	22
4.3 Kerr 1963: the rotating solution	23
4.4 Penrose-Hawking singularity theorems	23
4.5 The four laws of black-hole thermodynamics	23
4.6 Bekenstein 1972–1973: entropy is real	24
4.7 Hawking 1974: black holes radiate	24
5 The information paradox (Hawking 1976)	24
5.1 The paradox stated	24
5.2 Positions in the 1976–1995 debate	25
5.3 Page curve	25
6 Holography and AdS/CFT, 1993–2000	26
6.1 ’t Hooft and Susskind 1993–1995: the holographic principle	26
6.2 Maldacena 1997: AdS/CFT	26
6.3 Ryu-Takayanagi 2006: entanglement and geometry	27
6.4 The boundary reads the bulk	27
7 The firewall paradox, 2012	27
7.1 AMPS paradox	27
7.2 ER=EPR proposal (Maldacena-Susskind 2013)	28
8 The islands formula, 2019–2020	28
8.1 Penington 2019, Almheiri et al. 2019	28
8.2 Replica wormholes	28
8.3 Status as of 2026	29
9 Observational developments, 2015–2026	29
9.1 LIGO and gravitational waves	29
9.2 Event Horizon Telescope, 2019–2026	30
9.3 Direct horizon probes are not yet sensitive enough	30
10 The cohort inheritance and the roadmap of the fourteen Parts	30
10.1 The cohort inheritance	31
10.2 Roadmap of the fourteen Parts	33
11 Part I in plain English	35
 Part II. The Inherited Framework: GG-Theory in One Part	 36
12 One geometry, no dials: the six-dimensional bulk	36
13 The boundary projection: how four-dimensional physics arises	38

13.1	Which particles? The observer-indexed vacuum, and horizon thermality as openness	40
14	Zero free parameters: BI-6 and the integer ledger	41
15	The runtime: how a black hole’s branches are selected	42
16	What this buys the black hole	43
17	Part II in plain English	45
Part III. The Open Problems		47
18	What the information paradox really asks	47
18.1	Possible resolutions	47
18.2	Mathur’s theorem	48
19	The Bekenstein–Hawking entropy counting problem	48
19.1	How many microstates?	48
19.2	The species problem	49
20	The singularity: what’s inside?	50
20.1	Singularity replacements	50
21	Firewalls revisited	50
21.1	The AMPS argument in detail	50
21.2	The GG-Theory resolution	51
22	The Page curve derivation	51
23	The state of play	51
24	Part III in plain English	53
Part IV. The Horizon as Port		54
25	The horizon as an OGGEF port	54
25.1	The horizon is a special locus	55
25.2	The geometrically visible port: black holes as the cleanest test of OGGEF	56
26	Hawking radiation as the OGGEF output	57
26.1	The standard derivation	57
26.2	GDOS re-derivation	57
26.3	The Hawking temperature	58
26.4	The explicit Bogoliubov derivation: why the trace is thermal	58
27	The rotating anchor and the collapse channel: two theorems at the model level	59
28	The thermal clock face as a theorem	62
29	The real-bulk register at the horizon	63
30	Part IV in plain English	65
Part V. The Exact $1/4$		66
31	The Bekenstein–Hawking Entropy and the Exact $1/4$ from Induced Gravity	66
31.1	Entropy as entanglement	66
31.2	Strategy: gravity is induced, so the horizon entropy <i>is</i> the entanglement	67
31.3	Step 1 — the induced Newton constant from the canonical scale profile	67
31.4	Step 2 — one heat kernel, two quantities	68
31.5	Step 3 — the conical replica and the exact $1/4$	69
31.6	Step 4 — the microstates are the hidden fiber	71
31.7	Generality: neither supersymmetry nor extremality is used	71
31.8	Kerr, Reissner–Nordström, and the universality of the $1/4$	72
31.9	The subleading logarithmic correction	73
31.10	What is rigorous here, and what remains	74
32	Part V in plain English	74

Part VI. Information: Page Curve, Firewalls, and the Interior	76
33 The GG-Theory free-energy ledger and generalized entropy	76
33.1 Why the area term belongs to GEP	77
33.2 Why the bulk-entanglement term belongs to CEP	77
33.3 Operational entropy as a branch minimum	77
33.4 The factorization problem across the horizon	77
34 The explicit OGGEp channel for a Schwarzschild black hole	78
34.1 Setup	78
34.2 The OGGEp channel	78
34.3 Kraus representation	79
34.4 The entropy of the output state	79
34.5 Compositionality: post-evaporation state	80
35 The Page curve in GDOS	80
35.1 Setup	80
35.2 GDOS derivation	80
35.3 Connection to the islands formula	81
36 Hawking–Page and the Page curve as branch-switch events	81
36.1 Hawking–Page as a gravitational branch selection	81
36.2 The Page curve as a GFF branch minimum	82
36.3 The mechanism: port freshness and its scheduled expiry	83
36.4 Hayden–Preskill, and why no firewalls or remnants are forced	85
37 No firewalls in GDOS	85
37.1 The monogamy puzzle dissolves	85
37.2 Why no firewall is forced on the infalling observer	85
38 The interior of the black hole in GDOS	86
38.1 What’s inside?	86
38.2 Information doesn’t “fall” into the singularity	87
39 Black-hole complementarity revisited	87
39.1 Susskind-Thorlacius-Uglum 1993	87
39.2 GDOS complementarity	87
40 The four laws of black-hole mechanics as GDOS statements	88
41 The direction of the second law: the arrow consumed, the price computed	89
41.1 Three links, three owners	89
41.2 The lemma	90
41.3 The declared ceiling	92
42 The black hole as a way-station attractor: where the descent pauses, and where it ends	93
42.1 Hawking radiation is the exit channel	93
42.2 Today’s black holes are <i>conditional</i> ground states	94
42.3 Sorting the transitions: capture, milestone, and the one true GRIP-Out	95
43 The GDOS reading, in summary	95
44 Part VI in plain English	97
Part VII. Quantitative Predictions	98
45 Overview of predictions	98
46 Four black holes, end to end	98
47 EHT observations and horizon-scale geometry	99
47.1 What EHT measures	99
47.2 GG-Theory prediction for the shadow	100
47.3 Future prospects: ngEHT	100
48 LIGO/LISA ringdown spectroscopy	100
48.1 Quasi-normal modes	100
48.2 KK corrections to QNMs	101

48.3	LIGO/LISA prediction	101
49	Primordial black holes and the early universe	101
49.1	Primordial black-hole formation	101
49.2	GG-Theory PBH mass spectrum	102
49.3	Observational tests	102
50	Hawking spectrum corrections	103
50.1	Standard Hawking spectrum	103
50.2	GG-Theory corrections	103
50.3	Observable prospects	103
51	Entanglement-entropy tests via analog black holes	104
51.1	Analog gravity	104
51.2	Implications for GDOS	104
52	Cosmological black holes and de Sitter horizons	104
52.1	The de Sitter horizon	104
52.2	The de Sitter entropy from induced gravity	105
52.3	Open de Sitter puzzles, stated honestly	105
52.4	Connection to the Hubble-tension companion (C2)	105
53	The numerical table of predictions	106
54	Part VII in plain English	108
Part VIII. Experiments — Past, Now, and Future		109
55	The record so far: what has been measured, and what it did to the module	109
56	Now: the instruments in mid-2026	110
57	The decade ahead: milestones, with the verdicts pre-registered	110
58	The falsifier timeline, dated	111
59	Refresh rules: the module’s data discipline	113
60	Part VIII in plain English	116
Part IX. Summary and Comparison		117
61	Comparison with the leading frameworks: the deeper analysis	122
61.1	Classical general relativity	122
61.2	The fuzzball program	122
61.3	AdS/CFT correspondence	123
61.4	Firewall proposal	123
61.5	The islands formula in GDOS	123
61.6	ER=EPR in GDOS	124
61.7	What the runtime taxonomy adds that no leading framework has	124
62	Part IX in plain English	130
Part X. Discussions		131
63	What is new	131
64	Related work and precursors	133
65	The role of GG-Theory in the broader program	133
66	What the module explains, what remains open, and future directions	134
67	Falsifiability and kill criteria	136
68	Part X in plain English	137
Part XI. Conclusions		138
69	The 6D picture, one more time	138

70	The paradox was never about information	139
71	What is a horizon?	140
72	On determinism, probability, and the view from nowhere	141
73	Black holes as Rosetta stones	143
74	A closing reflection	143
75	Part XI in plain English	146
	Author’s Note	146
	Acknowledgments	148
 Part XII. Technical Appendices		149
A	Summary of abbreviations	151
B	Glossary of key symbols	152
C	Cross-cohort naming	152
D	Numerical parameters for reference	153
E	The Schwarzschild geometry from Einstein’s equations	153
E.1	Step 1 — the vacuum solution, worked	153
E.2	Step 2 — the horizon is not a place where physics breaks: Kretschmann, Finkelstein, Kruskal	155
E.3	Step 3 — orbits, the photon sphere, the ISCO, and the shadow	156
E.4	Step 4 — what the classical solution assumes, and cannot answer	157
E.5	Step 5 — the GG-Theory reading	158
F	Classical black-hole mechanics: the four laws	160
F.1	Step 1 — surface gravity and the zeroth law	160
F.2	Step 2 — the first law, verified	161
F.3	Step 3 — the Raychaudhuri equation and the area theorem	161
F.4	Step 4 — the third law, and the analogy table	162
F.5	Step 5 — what the analogy could not explain	162
F.6	Step 6 — the GG-Theory reading	162
G	Bekenstein’s entropy and the generalized second law	164
G.1	Step 1 — Wheeler’s teacup: the crisis	165
G.2	Step 2 — Bekenstein’s 1973 proposal	165
G.3	Step 3 — the Bekenstein bound, derived	165
G.4	Step 4 — the generalized second law, and the species problem	166
G.5	Step 5 — what the argument assumed, and the GG-Theory reading	166
H	Explicit derivation of Hawking temperature in GDOS	168
H.1	Step 1 — the stage: a scalar field on a fixed curved background	168
H.2	Step 2 — the Bogoliubov machinery	169
H.3	Step 3 — the warm-up with an exact answer: the accelerated observer	170
H.4	Step 4 — from Rindler to the black hole: two routes to κ	171
H.5	Step 5 — restoring units, and the numbers	172
H.6	Step 6 — what GDOS adds: the custody reading	173
I	Euclidean derivation of temperature and entropy	174
I.1	Step 1 — the cigar and the period	175
I.2	Step 2 — the on-shell action, term by term	176
I.3	Step 3 — the thermodynamics of the saddle	176
I.4	Step 4 — the GG-Theory reading	177
J	Hawking–Page free-energy notebook	179
J.1	Step 1 — temperature from regularity, and the two branches	179
J.2	Step 2 — the free-energy ledger	179
J.3	Step 3 — stability, and the branch-switch template	180
K	The Page curve entropy calculation	181
K.1	Step 1 — the exact random-state average	182

K.2	Step 2 — the curve and its corner	182
K.3	Step 3 — the evaporation clock	183
K.4	Step 4 — what GDOS keeps, and what it flags	183
L	From state counting to holography: the modern era	184
L.1	Step 1 — Strominger–Vafa: the first microscopic count	185
L.2	Step 2 — the holographic principle and AdS/CFT	185
L.3	Step 3 — geometrizing entanglement: Ryu–Takayanagi and the quantum extremal surface	186
L.4	Step 4 — islands and replica wormholes: the Page curve from gravity . . .	186
L.5	Step 5 — the remaining problems, stated precisely	187
L.6	Step 6 — where GG-Theory enters	187
M	Comparison with the islands formula	189
M.1	Step 1 — the QES condition, derived, and the two-branch race	189
M.2	Step 2 — why the two formulations agree	190
N	Island/QES and GG-Theory language	191
O	Entropy balance, channels, and data processing	192
O.1	Step 1 — reduction is a channel: Kraus and Stinespring	192
O.2	Step 2 — the contraction: data processing	193
O.3	Step 3 — Spohn’s production identity, worked	193
O.4	Step 4 — the single-mode exchange ledger, in closed form	194
P	Extremal black holes and the third law	196
P.1	Step 1 — the Kerr–Newman horizon data	196
P.2	Step 2 — area, spin, and the Christodoulou mass	196
P.3	Step 3 — the first law and the Smarr formula	197
P.4	Step 4 — the third law, GDOS-read	197
P.5	Step 5 — superradiance, the worked instance	198
Q	De Sitter horizon and cosmology	199
Q.1	Step 1 — the static patch and its surface gravity	200
Q.2	Step 2 — temperature and entropy, with units	200
Q.3	Step 3 — the GDOS reading	201
R	The induced-gravity species sum	202
R.1	Step 1 — the heat-kernel toolkit: where the sum rule comes from	202
R.2	Step 2 — the closure condition as a sum rule	203
R.3	Step 3 — why the Kaluza–Klein tower makes it self-consistent	203
R.4	Step 4 — the role of GU-5	204
R.5	Step 5 — the Standard-Model sector, and the honest gap	204
R.6	Step 6 — honest verdict	205
S	Internal source ledger	206
T	Cross-paper inheritance table	208
U	Consolidated benchmark audit	208
V	Named-statement deck	209
W	Promotion roadmap	209
X	Falsification-test catalog	211
Y	An annotated reading guide	211

Part XIII. Common Objections and Responses 212

65	Questions from the general reader	212
66	The classical relativist	223
67	The curved-spacetime field theorist	224
68	The string theorist	226
69	The islands and replica-wormhole theorist	227
70	The canonical-quantization referees	228

71	The emergent-gravity theorist	229
72	The gravitational-wave observer	230
73	The horizon-imaging astronomer	231
74	The quantum-information theorist	232
75	The quantum-foundations referee	234
76	The historian and the philosopher	235
77	The student	235
78	The cohort-internal referee	236

Part XIV. References

238

Preface

This paper is written to be readable, end to end, by a gravitational physicist — the reader who has fit a ringdown, argued about firewalls at a blackboard, or waited out an observing run — and by the graduate student meeting Hawking radiation for the first time. It is structured so that such a reader loses nothing by knowing no GG-Theory in advance: everything imported is restated with its hypotheses, every object is defined in Part II before it is used, every theorem is proved or proof-sketched on the page with the physics said in words immediately after the mathematics, and the appendix machine room re-derives the standard material from scratch at the level of a first quantum-field-theory course.

Here is the honest pitch for spending a week with it. You already own the phenomena: the four laws from your first relativity course, Hawking’s 1974 calculation, Page’s curve, the firewall argument, the islands papers. What you likely do not own is a single statement connecting them: that a horizon is the one place where an observer’s partial trace is *forced* by causal structure rather than chosen; that Hawking thermality is that trace computed, with the analytic continuation behind it a theorem of a one-sided spectrum rather than a trick; that the Bekenstein–Hawking $1/4$ is the ratio of two normalizations fixed before entropy enters — $4\pi/16\pi$ — with nothing to tune; that the Page turnover is a free-energy branch switch whose clock is an audited counting premise; and that the generalized second law’s production is a relative entropy spent toward the horizon’s thermal face, computed exactly in its class. Those identifications are theorems or honestly graded conditionals here, each with its register stated, each checkable with tools you already have.

A suggested week. Day one: Part I for the century you already partly own, and Parts II–III for the arena and the open problems stated precisely — do not skip the runtime vocabulary; the whole information story is bookkeeping on it. Day two: Part IV slowly — the forced trace, the Bogoliubov calculus done by hand with Appendix H beside it, and the thermal clock face (T-BH-5), which retires three continuation prescriptions at once. Day three: Part V — the four-step chain to the exact $1/4$, with the conical replica and Appendices I and R beside it. Day four: Part VI — the Page curve, the freshness clause, the firewall dissolution, and the second law’s arrow, with Appendix O in hand. Day five: Parts VII–VIII — the worked numbers and the experiments record with its dated timeline. The weekend: the falsification catalog, the Conclusions, the thirty plain-English questions that open the objections Part, and as many of the fifty-three technical objections as you have appetite for; if you have spent time at an interferometer, a telescope, or a blackboard, your objection is in there under your own job title.

Two further notes. The register of the paper — what is derived, what is conditional, what is open, what is conceded — is stated at every claim and audited in the named-statement deck; the named gaps (the species sum above all) are the paper’s work orders, not its fine print. And the perspective is an experimentalist’s: the author has spent five decades in experimental physics watching black holes cross from thought experiment to data — the ringdowns now fit, the shadows now imaged — and the paper is written in the discipline that instrument campaigns teach: every claim priced, every kill criterion dated, nothing asserted that a future instrument or a future computation cannot audit.

Conventions, naming, and cross-references. The mathematical core of this paper is the induced-gravity derivation of the Bekenstein–Hawking entropy with its exact coefficient (§31.5, Theorem 1); a reader who wants only the exact $1/4$ may turn there directly. Numbered results carry the cohort global identifiers with the paper suffix -BH (the identifiers are permanent cohort keys: $T\text{-}BH\text{-}n$ is Theorem n wherever it stands in the Part sequence). This paper builds on the cohort versions of record posted at preprints.arisaka-gg.org — GG-A1-Constitution (Arisaka, 2026, A1-Constitution), GG-A2-GG6, GG-A3-BI6 (Arisaka, 2026, A3-BI6), GG-A5-GDOS, GG-A6-GRIP, GG-A8-QM, GG-A9-SM (Arisaka, 2026, A9-SM), and GG-A11-Origin — and in particular on Theorems 0 and 2 of GG-A2-GG6 (the forced bulk and its canonical metric), the open-system law of GG-A5-GDOS, and the hidden-fiber picture of GG-A8-QM. For the canonical Dictionary of program-defined acronyms, see GG-A2-GG6 Appendix B.

Part I

Black Holes, 1915–2026

A Century from Einstein’s Field Equations to the Event Horizon Telescope




Figure 1: A century of black-hole physics, 1915–2026. Each milestone sharpened a puzzle that the next had to answer; the present work (right) addresses the entropy coefficient and the information paradox in the six-dimensional GG-Theory bulk. The dates are the field’s and not this paper’s, and only the last entry is a claim rather than a milestone. Read left to right the century has a shape: each answer created the next question, and the two that were never closed — what the entropy coefficient is, and where the information goes — are the two this paper takes up.

Black holes were supposed to be the simplest objects in physics. A classical black hole is characterized by three numbers — mass, spin, charge — and nothing else. Two black holes with the same three numbers are *identical*, no matter what they were made of: a collapsed star and a collapsed library of the same mass make the same hole. This “no-hair” simplicity is a theorem of general relativity, and for half a century it looked like a virtue. Then, between 1972 and 1976, three discoveries turned the simplest object in physics into the sharpest crisis in physics — and the century since 1915 became the story of that reversal.

Figure 1, above, lays the century out as a single line, and it is worth pausing on how strange a line it is. It begins with an exact solution written down within weeks of the field equations, by an artillery officer at the Russian front. It passes through a golden age in which the solutions acquired theorems — singularities proved generic, the rotating solution found, four laws with an unmistakably thermodynamic shape. It turns, in the four years from 1972 to 1976, into the crisis just described: an entropy that should not exist, a radiation that should not be thermal, and an information paradox that set the agenda for the next fifty years. And it ends, remarkably, with instruments: interferometers that hear black holes collide and a telescope the size of the Earth that photographs their shadows. Few objects in physics have traveled so far — from a formula in wartime to data on a screen — in a single century.

The rest of Part I walks that line at reading pace, milestone by milestone, gathering on the way the exact statements the later Parts must answer: what precisely the four laws assert, what Hawking actually calculated, what Page’s curve demands, what the holographic decade established and what it left open. A reader who owns this history already can skim ahead to Part II, where the

paper’s own arena is built. But first, one section states the claim of the whole paper in the form the cohort shares.

1 The central claim: black-hole entropy and the information channel as an open-system observable

It is worth stating, in one place, the recognition this paper shares with the whole GG-Theory cohort. No physical system is closed — so the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection *is* observation itself; its dynamics is the open-system law GDOS (GG-A5-GDOS), forced by the geometric Gi-4 (GG-A6-GRIP) and quantum-information Qi-4 (GG-A7-Qi4, [Arisaka, 2026, A7-Qi4](#)) pillars. It follows that the black-hole entropy $S = A/4$ and the fate of information are *registered observables* of one projection (OGGEP/GDOS), fixed by the geometry rather than postulated — which is why this paper can derive it, and why a single clean measurement can refute it. The ontology this rests on — a real six-dimensional bulk, with the boundary’s complex structure derived — carries its own cohort grade: the selection of the real category is L-Constitution-1 (alias L-Real-1) of GG-A1-Constitution, held at *candidate* grade with its falsifier row (F-Constitution-10) registered there ([Arisaka, 2026, A1-Constitution](#)), and this paper inherits that grade wherever the reality of the bulk does load-bearing work. The funnel below is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs.

2 Einstein’s 1915 field equations

On November 25, 1915, after an eight-year struggle, Einstein presented to the Prussian Academy of Sciences the final form of the field equations of general relativity ([Einstein, 1915](#)):

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi G_N T_{\mu\nu}, \quad (1)$$

where $R_{\mu\nu}$ is the Ricci tensor, R the Ricci scalar, $g_{\mu\nu}$ the metric, G_N Newton’s constant, and $T_{\mu\nu}$ the stress-energy tensor. The equations relate the curvature of spacetime (left-hand side) to the matter and energy content (right-hand side). Three tests followed quickly: Mercury’s perihelion precession (43 arcsec/century, explained the same month by Einstein); gravitational deflection of starlight by the Sun (confirmed by Eddington in 1919 at 1.75 arcsec, matching GR); and gravitational redshift (verified progressively through 1960s).

By 1919, GR was the established theory of gravity.

3 Schwarzschild 1916: the first exact solution

Within two months of Einstein’s publication, Karl Schwarzschild — then serving as an artillery officer on the Russian front — solved the field equations exactly for the spherically symmetric, static,




Figure 2: The central claim, as one deductive funnel, shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars. The same funnel opens the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs — here, black-hole entropy and the information channel.

vacuum case (Schwarzschild, 1916). The Schwarzschild metric is

$$ds^2 = - \left(1 - \frac{2G_N M}{c^2 r} \right) c^2 dt^2 + \frac{dr^2}{1 - \frac{2G_N M}{c^2 r}} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2). \quad (2)$$

At $r = r_s \equiv 2G_N M/c^2$ (the Schwarzschild radius), the metric has a coordinate singularity: g_{tt} vanishes and g_{rr} diverges. At $r = 0$, there is a genuine curvature singularity: the Kretschmann scalar $R_{\mu\nu\rho\sigma} R^{\mu\nu\rho\sigma} \sim 1/r^6$ diverges.

Schwarzschild died four months after his paper, from an autoimmune skin condition contracted at the front. He did not live to see the full physical significance of his solution — particularly that $r = r_s$ is not a physical singularity but an *event horizon*, a one-way causal surface separating the interior from the exterior.

3.1 The Einstein-Rosen bridge

In 1935, Einstein and Rosen ([Einstein and Rosen, 1935](#)) discovered that the Schwarzschild metric, when analytically extended, has a maximally extended form in which the interior connects to a second asymptotic region via a non-traversable “bridge.” This is the *Einstein-Rosen bridge* or wormhole. Its importance was not appreciated until the 1960s, when it was recognized as a key feature of the Kruskal extension.

3.2 Physical status of the interior

From 1916 through about 1960, the Schwarzschild interior was considered unphysical: a mathematical curiosity not realized in nature, because stars do not collapse past their own Schwarzschild radius in typical astrophysical scenarios. The dominant view, articulated by Eddington and others, was that quantum effects or some other mechanism would prevent catastrophic collapse.

This view collapsed when, in 1939, Oppenheimer and Snyder ([Oppenheimer and Snyder, 1939](#)) performed the first detailed calculation of gravitational collapse of a pressureless dust cloud, showing that the collapse proceeds to a singularity in finite proper time (as seen by an infalling observer) though with infinite coordinate time (as seen by an external observer). The “unphysical” interior became, in principle, physical.

4 The golden age, 1958–1975

Between 1958 and 1975, a dramatic reformulation of black-hole physics took place, driven by roughly a dozen key contributions. We review the most important.

4.1 Finkelstein 1958: the horizon is a coordinate artifact

David Finkelstein ([Finkelstein, 1958](#)) introduced an ingoing-null coordinate system in which the Schwarzschild metric is manifestly regular at $r = r_s$:

$$ds^2 = - \left(1 - \frac{2G_N M}{c^2 r} \right) dv^2 + 2dv dr + r^2 d\Omega^2. \quad (3)$$

This is the Eddington-Finkelstein metric. It showed that the apparent singularity at $r = r_s$ is a coordinate artifact, not a physical one, and that the surface $r = r_s$ is a one-way causal boundary: matter and signals can enter it but cannot exit.

4.2 Kruskal 1960: the maximal extension

Martin Kruskal and (independently) Szekeres obtained the maximal analytic extension of the Schwarzschild solution ([Kruskal, 1960](#)). The Kruskal diagram shows the complete Schwarzschild geometry as four regions: two asymptotic regions (I and III), a black-hole interior (II), and a white-hole interior (IV), joined along the horizon. The Penrose conformal diagram makes this structure visually intuitive.

4.3 Kerr 1963: the rotating solution

Roy Kerr ([Kerr, 1963](#)) discovered the exact solution for a rotating, uncharged black hole:

$$ds^2 = - \left(1 - \frac{2G_N M r}{\rho^2} \right) dt^2 - \frac{4G_N M a r \sin^2 \theta}{\rho^2} dt d\phi + \frac{\rho^2}{\Delta} dr^2 + \rho^2 d\theta^2 + \frac{\Sigma \sin^2 \theta}{\rho^2} d\phi^2, \quad (4)$$

where $a = J/(Mc)$ is the spin parameter, $\rho^2 = r^2 + a^2 \cos^2 \theta$, $\Delta = r^2 - 2G_N M r/c^2 + a^2$, and $\Sigma = (r^2 + a^2)^2 - a^2 \Delta \sin^2 \theta$. The Kerr solution has two horizons (outer and inner), an ergosphere where frame-dragging is so extreme that no observer can remain stationary, and a ring singularity at $\rho = 0$.

The subsequent charged Kerr-Newman solution (1965) completed the family of stationary asymptotically-flat black-hole solutions characterized by only three parameters: mass M , angular momentum J , and charge Q . This is the *no-hair theorem*, conjectured by Wheeler and proved by Israel, Carter, Hawking, and Robinson in the 1970s.

4.4 Penrose-Hawking singularity theorems

Roger Penrose in 1965 ([Penrose, 1965](#)), and Stephen Hawking in 1966–1970 ([Hawking and Penrose, 1970](#)), proved a series of singularity theorems. The essential result: under mild assumptions (dominant energy condition, existence of a trapped surface, or equivalent), gravitational collapse necessarily produces a singularity. Singularities are a generic feature of GR, not an artifact of special symmetries.

This result won Penrose a share of the 2020 Nobel Prize in Physics.

4.5 The four laws of black-hole thermodynamics

In 1973, Bardeen, Carter, and Hawking ([Bardeen et al., 1973](#)) derived four laws of black-hole mechanics whose shape was unmistakably thermodynamic. The zeroth law states that the surface gravity κ is constant over the horizon of a stationary black hole, just as temperature is constant throughout a body in equilibrium. The first law, $dM = (\kappa/8\pi)dA + \Omega dJ + \Phi dQ$, has exactly the form of $dU = T dS + \text{work}$, with the surface gravity standing where temperature should stand and the area A standing where entropy should stand. The second law is the area theorem just met: the horizon area cannot decrease, $dA/dt \geq 0$, echoing $dS/dt \geq 0$. And the third law states that extremal black holes ($\kappa = 0$) cannot be reached in any finite process, as absolute zero cannot. The authors were careful about what they claimed: these were presented as *analogies* — formal coincidences between geometry and heat — not genuine thermodynamic statements. It took the next two years to show that nature meant them literally.

4.6 Bekenstein 1972–1973: entropy is real

Jacob Bekenstein, a student of Wheeler’s, argued (Bekenstein, 1973) that the area-non-decrease should be interpreted as a genuine entropy law. He proposed

$$S_{\text{BH}} = \frac{A}{4\ell_P^2} k_B = \frac{Ak_B c^3}{4\hbar G_N}, \quad (5)$$

with $\ell_P^2 = \hbar G_N / c^3$ the Planck length squared. The argument was largely thermodynamic: if a black hole had no entropy, you could throw a hot cup of tea into one and delete its entropy from the universe, violating the second law. Bekenstein’s proposal faced considerable initial skepticism, including from Hawking.

Two things about Eq. (5) deserve to be called scandalous, and both drive the rest of this paper. It is enormous: a single solar-mass black hole carries $\sim 10^{77} k_B$ — more entropy than all the ordinary matter in the galaxy. And in statistical mechanics entropy *counts* something: $S = k_B \ln(\text{number of microstates})$, Boltzmann’s tombstone equation. So Eq. (5) is a census, and it demands an answer to the question: *states of what?* For fifty years no one could say what was being counted, or why the count comes out to exactly one quarter of the horizon area in Planck units — not one half, not $1/4\pi$, but $1/4$. Part V returns that answer.

4.7 Hawking 1974: black holes radiate

In 1974, Stephen Hawking (Hawking, 1975) calculated what happens when quantum fields propagate on the Schwarzschild background. The result was sensational: a black hole emits *thermal radiation* at a temperature

$$T_{\text{H}} = \frac{\hbar c^3}{8\pi G_N M k_B} = \frac{\hbar \kappa}{2\pi c k_B}, \quad (6)$$

where $\kappa = c^4/(4G_N M)$ is the surface gravity. For a solar-mass black hole, this gives $T_{\text{H}} \sim 10^{-7}$ K, utterly negligible compared to the cosmic microwave background (2.7 K). For a primordial black hole of mass 5×10^{11} kg, $T_{\text{H}} \sim 2 \times 10^{11}$ K, and the black hole evaporates in about the present age of the universe; anything lighter is already gone.

Hawking’s calculation was a bombshell. It meant that (a) Bekenstein was right: black holes carry real entropy $S_{\text{BH}} = A/(4G_N \hbar)$; (b) the four analogies are genuine thermodynamic laws; (c) black holes are not entirely black — they radiate; (d) ultimately, they *evaporate*. And with this last point came the paradox.

5 The information paradox (Hawking 1976)

5.1 The paradox stated

Consider a pure quantum state $|\psi\rangle$ that collapses gravitationally to form a black hole. The black hole then evaporates via Hawking radiation, which is thermal to all orders in perturbation theory.

A thermal state is mixed. So the process is:

$$|\psi\rangle \rightarrow \text{Black hole} \rightarrow \text{Hawking radiation} \approx \rho_{\text{thermal}}. \quad (7)$$

If the final state is truly thermal, then the map $|\psi\rangle \rightarrow \rho_{\text{thermal}}$ is a pure-to-mixed evolution, which violates unitarity. If unitarity is preserved, the initial pure-state information must be somehow encoded in the radiation. Unitarity is a precise way of saying that information is never destroyed, only scrambled: burn a photograph and the information survives — uselessly, but really — in the smoke and the ash. Hawking’s calculation said a black hole is not like fire: the photograph goes in, and nothing — not even scrambled nothing — comes out.

In 1976, Hawking (Hawking, 1976) argued explicitly that the information was *lost*: the evolution was fundamentally non-unitary, and the quantum-mechanical description of gravity required revision. This was the *information paradox* in its original formulation. It set off a five-decade controversy.

5.2 Positions in the 1976–1995 debate

Through the 1980s and early 1990s the community sorted itself into four camps. Hawking himself — with Penrose initially alongside — held that information is simply *lost*: quantum gravity is fundamentally non-unitary, density matrices evolve into density matrices, and the S-matrix must be replaced by what Hawking called a $\$$ -matrix. A second camp, whose case Don Page built in 1980 and sharpened in 1993, held that information is *preserved in correlations*: the radiation is not exactly thermal; subtle correlations between the emitted quanta encode the missing information, and the total state remains pure — the position that implies the *Page curve* for the entanglement entropy. A third camp (Aharonov–Bohm; Banks–Susskind) proposed that information is *preserved in a remnant*: evaporation halts at the Planck scale, leaving a Planck-mass object holding everything the hole ever swallowed. The difficulty is that such a remnant would need unboundedly many internal states — $\sim e^{S_{\text{initial}}}$ for arbitrarily large initial black holes — and this infinite-species problem proved hard to live with. And a fourth camp — Susskind, ’t Hooft — bet that information is *preserved, somehow*: the specific mechanism unknown, but string theory or some other complete theory would rescue unitarity by a holographic reorganization of the degrees of freedom. By 1993, Page’s analysis had made the “information-preserved in subtle correlations” scenario the leading candidate. His paper gave the explicit entanglement-entropy curve that the radiation should follow.

5.3 Page curve

For a unitary evaporation process of a black hole initially in a pure state, the von Neumann entropy of the radiation grows then shrinks, following the *Page curve* (Page, 1993). The logic is elementary bookkeeping on entanglement. At early times ($t \ll t_P$) the emitted radiation is a small subsystem of the total $S + R$ Hilbert space, and its entanglement entropy grows as $\ln(d_R)$, where d_R is the dimension of the radiation subspace — roughly $S_{\text{BH}}^{(R)} = A_{\text{missing}}/(4G_N\hbar)$, the entropy of the area that has so far evaporated away. At the Page time ($t = t_P$) the radiation entropy catches up with the black-hole entropy; half the initial area has radiated away, $A(t_P) = A(0)/2$. At late times ($t > t_P$) the black hole is the smaller subsystem, so the radiation entropy is capped by the

black-hole entropy, $S_{\text{BH}}^{(\text{radiation})} = S_{\text{BH}}^{(\text{BH})}$, and both decrease together as the remaining area shrinks. At the end ($t = t_{\text{evap}}$) both entropies reach zero: the final state is pure, and every correlation has been returned to the radiation.

The Page curve is tent-shaped: it rises linearly from 0, peaks at t_P , and falls linearly back to 0. Its peak is at $S_{\text{max}} = A_0/(8G_N\hbar) = S_{\text{BH}}^{(0)}/2$.

The Page curve is what *unitarity* predicts. The naive Hawking calculation predicts monotonically increasing entropy. For forty years the question was: what mechanism changes the Hawking calculation at $t \sim t_P$ to produce the Page curve? The entire modern subject can be stated in three words: *derive the turnover*.

6 Holography and AdS/CFT, 1993–2000

6.1 ‘t Hooft and Susskind 1993–1995: the holographic principle

Gerard ‘t Hooft ([‘t Hooft, 1993](#)) and Leonard Susskind ([Susskind, 1995](#)) independently proposed the *holographic principle*: the physical content of a region of space is encoded entirely on its boundary, with one Planck-area unit of information per bit. This was motivated directly by the Bekenstein–Hawking entropy: if $S_{\text{BH}} = A/(4\ell_P^2)$ counts microstates, and the area is a 2D surface, then the microstate information “lives” on the surface, not in the bulk volume. Extrapolating, all physical systems might obey a holographic bound: $S \leq A/(4\ell_P^2)$ in appropriate units.

This was a radical claim. Ordinary field theories have volume- extensive entropy, not area- extensive. If the holographic principle is correct, then ordinary field theories are approximations that break down at the level of gravitational degrees of freedom.

6.2 Maldacena 1997: AdS/CFT

Juan Maldacena ([Maldacena, 1998](#)) proposed the AdS/CFT correspondence in 1997. In its original form:

Type IIB string theory on $\text{AdS}_5 \times S^5$ is equivalent to $\mathcal{N} = 4$ super-Yang-Mills in 4D Minkowski space.

The two theories are dual: every quantity in one can be translated to a quantity in the other via a specific dictionary. Bulk states correspond to boundary states; bulk operators to boundary operators; and, critically, bulk gravitational dynamics to boundary gauge dynamics. The holographic principle is realized concretely: the 5D+1 bulk theory (including gravity) is captured entirely by a 4D boundary theory (without gravity).

AdS/CFT became a dominant framework in theoretical physics for the next 25 years, and it earned the position three times over. It provided a concrete setting in which black-hole evaporation is manifestly unitary, because the boundary theory is a standard quantum field theory and standard quantum field theories do not lose information. It provided a computational tool: bulk gravitational problems can be mapped to boundary gauge-theory problems, and vice versa, often with dramatic simplifications in one direction or the other. And it provided a testing ground — for the holographic

principle, for the information paradox, and for quantum gravity in general — on which sharp calculations could replace verbal arguments.

The paradox, in AdS/CFT language, became: *How does the bulk gravitational picture (which suggests information loss) reconcile with the boundary gauge-theory picture (which is manifestly unitary)?*

6.3 Ryu-Takayanagi 2006: entanglement and geometry

Shinsei Ryu and Tadashi Takayanagi ([Ryu and Takayanagi, 2006](#)) in 2006 proposed a formula for entanglement entropy in AdS/CFT:

$$S(A) = \frac{\text{Area}(\Sigma_A)}{4G_N\hbar}, \quad (8)$$

where A is a boundary region, and Σ_A is the minimal bulk surface whose boundary coincides with ∂A . This is the *Ryu-Takayanagi formula*, often abbreviated RT.

RT makes quantitative the holographic principle at the level of entanglement: entanglement entropy of the boundary theory is the area of a geometric surface in the bulk. This formula has been checked in many examples and extended to time-dependent settings (HRT formula, 2007) and to quantum corrections (FLM formula, 2013).

6.4 The boundary reads the bulk

Between 2006 and 2019, the AdS/CFT program developed an increasingly sophisticated understanding of how bulk information is encoded on the boundary. The HKLL prescription (2006) showed how boundary operators reconstruct bulk states. In 2015, Almheiri, Dong, and Harlow recognized that the bulk-to-boundary map is a quantum error-correcting code — a discovery that reorganized the subject’s intuitions. Entanglement-wedge reconstruction identified exactly which portion of the bulk a given boundary region can access, and tensor-network models of AdS geometry (Swingle 2009; the HaPPY code, 2015) made the encoding concrete enough to draw on a blackboard. Each contribution clarified the bulk–boundary relationship; none yet resolved the information paradox at the level of the Page curve.

7 The firewall paradox, 2012

7.1 AMPS paradox

In 2012, Ahmed Almheiri, Donald Marolf, Joseph Polchinski, and James Sully ([Almheiri et al., 2013](#)) identified a sharp tension in the black-hole complementarity picture. The *AMPS argument*:

Consider a late-time Hawking quantum B emitted from an old black hole (past the Page time). Unitarity requires that B be entangled with the earlier radiation R . Smoothness of the horizon (no-drama postulate) requires that B be entangled with its interior partner A . But entanglement is *monogamous*: B cannot be maximally entangled with both R and A .

Something must give. AMPS argued that the most likely candidate is smoothness of the horizon — i.e., an infalling observer encounters a “firewall” of high-energy quanta at the horizon. This violates Einstein’s equivalence principle.

The firewall puzzle became the central open question in black-hole physics from 2012 to 2019. It sharpened the previous information-paradox discussions: unitary evaporation *requires* that something fail, and the natural candidate is the smoothness of the horizon — a foundational feature of GR.

7.2 ER=EPR proposal (Maldacena-Susskind 2013)

Juan Maldacena and Leonard Susskind (Maldacena and Susskind, 2013) proposed in 2013 that the solution to the firewall paradox is that *entangled black holes are connected by wormholes*: ER = EPR, where ER stands for Einstein-Rosen (bridge) and EPR for Einstein-Podolsky-Rosen (entangled pair).

In this picture, the Hawking radiation quanta are not separate from the black hole; they are the *far ends of wormholes* whose near ends are in the black hole interior. Entanglement between B and R is realized geometrically as a wormhole connecting them. The apparent monogamy violation is an artifact of treating the radiation and interior as separate subsystems; they are, in the ER=EPR picture, connected.

ER=EPR is conceptually bold but mathematically loose. It inspired further work (Cao, Carroll, Michalakis 2016; Van Raamsdonk 2010) but was not a complete resolution.

8 The islands formula, 2019–2020

8.1 Penington 2019, Almheiri et al. 2019

A breakthrough came in 2019–2020 when Geoff Penington (Penington, 2020), and independently Ahmed Almheiri, Netta Engelhardt, Donald Marolf, and Henry Maxfield (Almheiri et al., 2019), showed that the Page curve could be derived *directly from semiclassical gravity*, without needing any exotic assumptions, by a careful application of the quantum extremal surface (QES) prescription.

The *islands formula* states:

$$S(R) = \min \operatorname{ext} \left\{ \frac{\operatorname{Area}(\partial I)}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I) \right\}, \quad (9)$$

where R is the radiation region, I is the island (a bulk region disjoint from R), and one minimizes over all choices of I , then extremizes over its boundary.

At early times, the trivial choice $I = \emptyset$ gives $S(R) = S_{\text{semi-cl}}(R)$, which grows monotonically. At the Page time, a non-trivial island I deep inside the black hole becomes the minimum, and $S(R)$ starts to decrease, reproducing the Page curve.

8.2 Replica wormholes

The derivation of the islands formula requires evaluating the gravitational path integral with replicated boundary conditions, and summing over bulk geometries including wormhole-like configu-

rations that connect the replicas. These *replica wormholes* (Almheiri et al., 2020) are the mechanism by which the nonlinear contributions to the entanglement entropy produce the Page curve.

The physical picture: at sub-Page times, the radiation and the black hole are weakly entangled, and the trivial bulk geometry (no wormholes) dominates. At super-Page times, the wormhole topologies dominate, making the radiation a purification of the black-hole interior — the interior “lives inside” the radiation.

8.3 Status as of 2026

As of mid-2026, the islands formula is widely accepted as resolving the information paradox *within* the AdS/CFT framework. The resolution has two key features: (i) the black-hole evaporation is unitary, with the Page curve realized; (ii) the mechanism is subtle non-perturbative gravity, not new physics at the horizon.

The limitations, however, are real, and worth stating as carefully as the achievement. The derivations use Euclidean gravitational path integrals that are not rigorously defined. They require specific lower-dimensional models — JT gravity in two dimensions, hyperbolic quotients — and their extension to four dimensions remains partial. The physical interpretation of the replica wormholes is abstract: there is no clear “mechanism” in the operational sense, no account of what, physically, carries the information out. The connection to real astrophysical black holes in 4D asymptotically-flat space is indirect. And the connection to the real Standard-Model universe is obscure — AdS/CFT is not the Standard Model.

9 Observational developments, 2015–2026

9.1 LIGO and gravitational waves

In September 2015, the LIGO collaboration detected gravitational waves from the merger of two black holes (Abbott et al., 2016). The signal (GW150914) was a 200 ms chirp, followed by a ringdown phase during which the newly-formed black hole settled into its final state via quasi-normal-mode oscillations.

Subsequent observing runs turned the detection of mergers into a census. The fourth observing run (O4) of the LIGO–Virgo–KAGRA network ran from May 2023 to November 2025; the GWTC-4.0 catalog release (2025) added 128 significant events from its first segment to the 101 pre-O4 events of its reanalysis (the 90 of GWTC-3 plus reprocessed candidates), and the GWTC-5.0 release (2026) added a further 161 from the second, bringing the confirmed total to $101 + 128 + 161 = 390$ detections since 2015 — roughly three quarters of them from O4 alone, with the final O4 segment still under analysis. The census includes a neutron-star merger with electromagnetic counterpart (GW170817), intermediate-mass and mass-gap events (GW190521, and GW231123 at ~ 225 solar masses the most massive binary black-hole merger yet observed), and GW250114, the clearest signal on record (signal-to-noise ratio ~ 77 , a near-equal $32 + 34 M_\odot$ pair). These detections provide direct probes of the final-state geometry of coalescing black holes, accessible through the quasi-normal-mode spectrum in the ringdown phase; The area-theorem confirmations test the

classical half of the GSL story; the quantum half — entropy production toward the KMS face — is not directly tested by any interferometer, and no row below scores it as if it were. GW250114’s ringdown resolved the fundamental quadrupolar mode together with its first overtone, consistent with the Kerr spectrum, and its inspiral-to-ringdown horizon-area budget confirmed Hawking’s area theorem at high confidence.

9.2 Event Horizon Telescope, 2019–2026

In April 2019, the EHT collaboration published the first image of a black-hole shadow: M87*, the $6.5 \times 10^9 M_\odot$ supermassive black hole at the center of M87 ([Event Horizon Telescope Collaboration et al., 2019](#)). The shadow matched the Kerr prediction within errors. In 2022, the EHT released an image of Sgr A* (the $4 \times 10^6 M_\odot$ black hole at the Galactic center), also consistent with Kerr. Subsequent campaigns deepened both portraits: polarization-resolved imaging of Sgr A* (2024) revealed strong, ordered magnetic fields at the horizon scale; a second-epoch M87* image (2024, from 2018 data) reproduced the shadow diameter with an independent data set; and multi-epoch polarimetry of M87* (2025) tracked year-to-year reversals of the polarization pattern — evolving magnetized plasma around a shadow whose size stays put, exactly as a Kerr geometry demands.

EHT imaging constrains the horizon-scale geometry directly. Any deviation from Kerr ringdown or Kerr shadow at the 5–10% level would show up in these observations. Through mid-2026, no deviations from Kerr have been detected beyond astrophysical uncertainties.

9.3 Direct horizon probes are not yet sensitive enough

The current EHT and LIGO observations are consistent with classical Kerr to the precision achieved. They do not constitute a test of the Hawking radiation, the Bekenstein–Hawking entropy, or the information-paradox resolution mechanism. Those tests remain in principle possible but require much higher precision than available as of 2026.


This is important context: the black-hole information paradox is a *theoretical* problem, not an observational one. The theoretical constraints are unitarity, GR, QFT, and thermodynamics. No current observation has falsified any of these constraints; the paradox is purely one of consistency between them.

10 The cohort inheritance and the roadmap of the fourteen Parts

Before the open problems are posed in their modern form (Part III) and the GDOS reading is built (Parts IV–VI), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§10.1, Figure 3), and a panorama of the fourteen Parts of the present paper (§10.2, Figure 4). A reader of any later Part should be able to flip back to this Section and orient in both directions: which cohort companion carries which imported result, and where the present argument stands in the fourteen-Part arc.

10.1 The cohort inheritance

The program in one view. Figure 3 places the present paper among the twenty-two papers of the GG-Theory cohort. The reader meeting the program for the first time can read its architecture off the figure top to bottom — the constitutional root, the bulk trio, the open-system boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch on the map’s left flank — before continuing with the present paper’s own thread.



The Foundations Branch. The map’s left flank carries the cohort’s third branch — two member papers, a single column beside the spine, whose subject is the boundary’s complex structure and its mirrors. GG-P4-Real (Arisaka, 2026, P4-Real) is the origin-of-*i* member: its causal clock is the real structure beneath the analytic machinery this paper computes with — the positive-frequency splittings, thermal traces, and branch cuts of the Bogoliubov and collapse-channel sections (§27 closes entries R14–R15 of the real-bulk series it develops at full length). GG-P5-CPT (Arisaka, 2026, P5-CPT) defines time by its four endowments and fixes the law-level arrow as the dictionary’s one-way reading, derived over an exactly T-symmetric bulk — kin of the horizon thermodynamics derived here; its Wick-quartet theorem T-P5-B1 licenses the analytic continuation from which this paper proves its KMS thermal clock face (T-BH-5).

10.2 Roadmap of the fourteen Parts

The rest of the paper is organized into thirteen further Parts, and the walk through them is plotted in Figure 4. The present Part — the century, 1915–2026 — is nearly done: it has walked from Einstein’s field equations to the first images of a horizon and the ringdown decade, and it closes with the two orientation maps of this Section. Part II opens the toolbox: the six-dimensional arena with its canonical scale profile, and the runtime vocabulary — attractors, the two regimes — in which the black hole’s life cycle will be sorted; every object is named before it is used. Part III turns the history into a charge sheet: the information paradox stated precisely, the entropy-counting bar, the firewall argument, the Page curve, and the state of play — the open problems posed precisely enough to be answered or refuted.

Parts IV–VI are the engine room, and they are highlighted in the roadmap because together they carry the paper’s central claims. Part IV reads the horizon as an OGGEp port of the six-dimensional bulk: the geometrically forced trace, Hawking radiation as its output, the thermal clock face proven rather than assumed, the rotating anchor and the collapse channel, and the real-bulk register. Part V derives the central result — the Bekenstein–Hawking entropy with its exact coefficient, $S = A/4G_N$ and $1/4 = 4\pi/16\pi$, from induced gravity (Theorem 1) — together with its honest accounting. Part VI assembles the information story: the free-energy ledger, the explicit channel, the Page curve as the scheduled expiry of port freshness, the direction of the second law derived with its declared ceiling, no firewalls, the interior, the four laws, and the black hole as a way-station attractor.

Parts VII and VIII confront the sky: the quantitative predictions, laid out where the instruments actually point, and the experiments record — the kill criteria stated before the instruments report, with the cumulative evidence base. Parts IX–XI weigh the case and close it: the scorecards and framework comparisons; the Discussions — what is new, what remains open, and the falsification program walked row by row; and the Conclusions — the view from the six-dimensional bulk. The back matter fills three Parts of its own: the technical appendices (Part XII), the objections answered in the open (Part XIII), and the annotated references, each carrying its Relevance note (Part XIV). Read in order, the fourteen Parts compose a single self-contained argument; read by need, each Part names its imports and can stand a visit on its own.




Figure 4: Roadmap of the fourteen Parts of GG-C3-BH. The argument develops in five phases. *The paradox posed* (Parts I–III): the century of black-hole physics, the inherited framework, and the open problems stated precisely enough to be answered or refuted. *The derivation* (Parts IV–VI, highlighted): the horizon as a geometrically forced OGGEF port with its proven thermal clock; the Bekenstein–Hawking entropy with its exact coefficient, $S = A/4G_N$, from induced gravity (Theorem 1) — the result on which every later Part stands, and the strongest box in the column; and the information story, from the scheduled expiry of port freshness to the direction of the second law. *Confrontation with the sky* (Parts VII–VIII): the quantitative predictions and the experiments record with its kill criteria. *Weighing and summit* (Parts IX–XI): scorecards and comparisons, the Discussions, and the Conclusions. *The machine room and the record* (Parts XII–XIV): the technical appendices, the objections answered in the open, and the annotated references.

11 Part I in plain English: a century of the field, and the two questions it left open

The simplest objects in physics have the longest bibliography. A black hole is described by three numbers — mass, spin, charge — and a hundred years of work has gone into two questions those three numbers do not answer.

The century divides cleanly. From 1915 to about 1975 the work was geometric: Schwarzschild’s exact solution the year after the field equations, Kerr’s rotating one in 1963, the theorems that showed singularities were generic rather than artifacts of a convenient symmetry. By the middle of the 1970s the object was understood as geometry. Then two results changed what kind of object it was. Bekenstein said a black hole has an entropy proportional to the area of its horizon. Hawking computed the temperature that entropy implies, and found that the hole radiates.

Both open questions come from those two results. The entropy formula carries a coefficient — one quarter — and nobody derived it from a count of anything; it was read off a calculation and has sat there, exact and unexplained, for fifty years. And if the hole radiates it eventually evaporates, and the information that fell in has nowhere left to go. The debate that followed ran from 1976 to the present through holography, the Page curve, firewalls and islands, and its state today is a formula that works in controlled models and a mechanism nobody has grounded in a theory of the actual universe.

This paper takes up both, from a geometry fixed elsewhere in this cohort for reasons that had nothing to do with black holes. The last entry on the timeline is a claim rather than a milestone, and the Parts that follow are the argument for it.

Part II

The Inherited Framework: GG-Theory in One Part

What the Black-Hole Module Inherits — the Six-Dimensional Bulk, the Boundary Projection, the Zero-Parameter Integer Ledger, and the Open-System Runtime

This Part answers a single question: *what is GG-Theory, and what does a black-hole module built on it inherit?* Part I established the century’s baseline and ended with the claim the whole paper must earn; before the paper can claim geometric answers to the horizon’s open problems, the reader deserves a clear and unhurried picture of the machine the answers come out of. Nothing in this Part is new physics. Everything here was proved upstream, in the companion papers of the cohort, and is imported wholesale; the job of this Part is to make that inheritance *intuitive* — to hand a reader who has seen only Part I the four ideas the rest of the paper leans on, in the order in which they build on one another. The from-scratch derivation chapters a graduate reader may want — quantum fields in curved spacetime, the Bogoliubov calculus, heat kernels, the replica method — remain the business of the technical appendices (Part XII); this Part is the frame those tools hang on. The open problems themselves are restated precisely in the next Part, once the framework is in hand.

12 One geometry, no dials: the six-dimensional bulk

At the center of the whole program is a single geometric object — the constitution of GG-A1-Constitution ([Arisaka, 2026, A1-Constitution](#)) fixes what may count as an admissible claim, and GG-A2-GG6 ([Arisaka, 2026, A2-GG6](#)) derives the geometry itself — and it is worth meeting it in plain language before meeting it in equations. GG-Theory posits that reality is a *six-dimensional bulk*, a tapering “throat” of the form $M_4 \times I_r \times S_\theta^1$. Three pieces make it up. The familiar four dimensions M_4 are the spacetime we live in. A single radial direction I_r runs along the throat from a wide ultraviolet mouth down to a narrow infrared end; this is the energy axis, and it is best read as the renormalization-group scale, $r = \ln(E_{\text{GUT}}/\mu)$ — moving along r is nothing other than running a coupling from high energy to low. And a small compact angle S_θ^1 , the phase fiber, carries the locked gauge structure: its winding is where the cohort’s dark sector lives — and, at a horizon, where this paper’s microstates will be counted.

The word “warped” deserves a warning, because it is borrowed from a different picture. Here it does *not* mean an extra spatial dimension with a curved metric in the Randall–Sundrum sense. It means RG running along r : the warp factor is the field-theory / holographic-RG reading of scale dependence, exactly as the rest of the cohort reads it. With that one substitution the geometry stops being exotic and becomes familiar — “deeper into the throat” simply means “lower in energy.”

The single most important intuition to carry forward is this: *our four-dimensional world is the boundary of this bulk*. We do not live inside the throat; we live on its edge, and what we call physics is the shadow the bulk casts there. How that shadow is cast — the boundary projection itself — is the subject of the next section, where Figure 5 draws the throat, its radial scale axis, its compact fiber, and the four-dimensional boundary slab on which the black holes of Part I live.

The geometry, made explicit. Behind the picture is a definite, parameter-free metric. The bulk is

$$M_6 = M_4 \times I_r \times S_\theta^1, \quad ds_6^2 = e^{-2r} g_{\mu\nu}(x) dx^\mu dx^\nu + dr^2 + d\theta^2, \quad (10)$$

the canonical scale-profile metric of GG-A2-GG6 (Theorem 2: no free parameters k , R_θ , or r_c survive)(Arisaka, 2026, A2-GG6). The warp profile $A(r) = r$ is not chosen but forced, by demanding constant bulk curvature with a flat boundary ($A''(r) = 0$); there is no curvature scale, no compactification radius, and no brane tension to tune. Here I_r is the non-spatial *scale dimension* — the renormalization coordinate with range $r \in [0, K]$, $r = 0$ the UV boundary at E_{GUT} and $r = K$ the IR boundary at M_{lock} — and S_θ^1 is the non-spatial *phase dimension*, a circle of fixed coordinate period 2π whose proper size at the IR boundary is $\ell_{\text{lock}} \sim 1/M_{\text{lock}}$. The Planck–electroweak hierarchy is then the topological number e^{-K} , not a fine-tuned ratio. That this six-dimensional arena is *forced* by LC (AX-1) and GGEP (AX-2) rather than postulated is Theorem 0 of GG-A2-GG6.

Six dimensions are four plus one plus one. The count is forced role by role (Arisaka, 2026, A2-GG6): four Lorentzian directions M_4 carry registered observations and the only length dimension; one non-compact *scale* direction I_r is the geometrized RG flow (a dimensionless energy logarithm, not space); and one compact *phase* circle S_θ^1 carries integer-quantized Wilson-line holonomy. The three bulk directions are exactly what is needed to absorb the three dimensionful constants c , G_N , and $\hbar$; the boundary supplies the remaining two, k_B and the energy–information quantum Q_B — together the five universal constants GU-5 = $\{c, G_N, \hbar, k_B, Q_B\}$, of which the boundary-thermodynamic member k_B is the conversion factor that will carry the Bekenstein–Hawking entropy below. This “three-plus-two” architecture is the spine of the whole ledger.

One geometry carries all six forces. On the bulk the gauge and gravitational structure is a single group,

$$G_{\text{Bulk}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i},$$

with the Lorentz/spin factor $\text{Spin}(1, 5)$ (the gravity sector) on equal footing with the internal factors (Arisaka, 2026, A2-GG6). Here $\text{SU}(5)$ is the minimal simple unification containing the Standard Model; $\text{U}(1)_B$ is gauged baryon number, which turns proton stability into a theorem; and $\text{U}(1)_{Q_i}$ is the informational gauge force that the Qi-4 pillars later use. Projected onto the four-dimensional boundary this becomes the six forces an observer actually measures — strong, weak, electromagnetic, baryonic, informational, and gravitational — with the BI-6 modified Bianchi identity of the next-but-one section guarding the closure. Crucially for this paper, the same compact circle S_θ^1 that quantizes these holonomies is where the horizon microstates of Part V are counted and where the companion papers read the Kaxion dark matter: the black hole and the particle ledger are two projections of this one object, one story apart.

For a black hole, the 4D metric $g_{\mu\nu}(x)$ in Eq. (10) is the Schwarzschild (or Kerr, or Kerr–Newman)

geometry. The 6D bulk accommodates this by having the 4D metric vary over x^μ while the compact dimensions remain (at least in the low-energy effective description) unperturbed. The horizon in the 4D metric is at $r_s = 2G_N M/c^2$ for Schwarzschild.

13 The boundary projection: how four-dimensional physics arises

If the bulk is the whole of reality, why do we only ever measure four-dimensional physics? The answer is the *boundary law*, and it is the deepest single result the black-hole module inherits. An observer outside a horizon can never access what lies behind it, and no boundary observer at all can access the closed bulk; what an instrument registers is the projection of the bulk onto the four-dimensional edge the observer inhabits. GG-A5-GDOS (Arisaka, 2026, A5-GDOS) raises this to a principle — the open-system geometric–gauge equivalence principle (OGGEP), the third axiom of the constitution — which requires every admissible theory to supply an explicit dictionary

$$\mathcal{D}_{\text{OGGEP}} : (\mathcal{H}_{\text{Bulk}}, \rho_{\text{Bulk}}, U_{\text{Bulk}}) \longmapsto (\mathcal{H}_{\text{Bdy}}, \rho_{\text{Bdy}}, \Phi_{t,s}, \{\mathcal{I}_\alpha\}),$$

carrying the inaccessible bulk data to a boundary state, an admissible (completely positive, trace-preserving) evolution channel $\Phi_{t,s}$, and the instruments that read it. The decisive recognition is that this bulk-to-boundary projection *is observation itself*: measurement is not a postulate appended to the theory but what the projection does.

GG-A5-GDOS then makes the projection dynamical and unique. Its boundary law, the geometric dynamics of open systems (GDOS), is the Lindblad/GKLS generator

$$\mathcal{L}_{\text{GDOS}}(\rho) = -i[H_{\text{eff}}, \rho] + \sum_{\alpha} \gamma_{\alpha} \left(L_{\alpha} \rho L_{\alpha}^{\dagger} - \frac{1}{2} \{L_{\alpha}^{\dagger} L_{\alpha}, \rho\} \right), \quad \gamma_{\alpha} \geq 0,$$

derived — not postulated — from the exact bulk dynamics by a Born–Markov reduction (Arisaka, 2026, A5-GDOS). The ordinary Schrödinger equation is recovered as its closed-system limit; openness is primary and closure the special case. In this language an *observable* is the value an open system reads off at a stationary seam, an attractor of the GDOS flow — and from this one law GG-A5-GDOS recovers four-dimensional quantum mechanics, general relativity, and quantum field theory as boundary theorems rather than separate assumptions.

The warm consequence, the one to hold onto, is that *ordinary general relativity is recovered, not discarded*. Because four-dimensional general relativity is already a theorem of GG-A5-GDOS, the Schwarzschild and Kerr geometries of Part I live inside the framework automatically, with the four-dimensional Newton constant emerging as a scale-overlap integral over the throat. Part V sharpens that emergence into the paper’s central result: because no bare boundary G_N exists to be corrected, the entanglement entropy across a horizon is not a correction sitting on top of $A/4G_N$ — it *is* $A/4G_N$. We do not choose between GG-Theory and the relativity textbooks: the textbook geometry is exactly what the boundary law produces, and what the framework adds is a fixed, parameter-free account of the horizon’s thermodynamic face.

A boundary law fixes *what* the four-dimensional theory is; a black-hole module also needs to know *how* its lawful possibilities are explored and stabilized in time. That is the geometric runtime,




Figure 5: Schematic of the six-dimensional bulk $M_4 \times I_r \times S_\theta^1$ and its OGGEP projection to the four-dimensional boundary. The radial axis $r = \ln(E_{\text{GUT}}/\mu)$ is the renormalization-group / scale axis: the ultraviolet GUT edge sits at the top ($r = 0$, $E_{\text{GUT}} \sim 10^{16} \text{ GeV}$) and low-energy physics lies toward the boundary, with the inherited scale chain $E_{\text{Pl}} \rightarrow E_{\text{GUT}} \rightarrow M_{\text{lock}}$ ($M_{\text{lock}} \sim 3.85 \text{ TeV}$, the Kaluza–Klein lock scale) marked along it. The compact angle S_θ^1 is the phase fiber whose winding carries the dark sector — and, at a horizon, the microstates this paper counts in Part V. The bold arrow is the OGGEP bulk-to-boundary map: black-hole observables on M_4 (ordinary general relativity: Schwarzschild, Kerr) are boundary shadows of bulk structure, with the boundary governed by the GDOS boundary law. “Warped” denotes RG running along r , *not* a Randall–Sundrum spatial dimension.

the four invariance pillars Gi-4 of GG-A6-GRIP (Arisaka, 2026, A6-GRIP), organized as a single cascade

$$\text{TGG} \implies \text{CCI} \implies \text{GFF} \implies \text{GRIP}.$$

The total geometric–gauge object TGG packages the closed bulk ledger as one composite connection on the G_{Bulk} bundle; the covariant conservation identity CCI hands the boundary observer a local conservation law and, under closure, forces the integer ledger of the next section; the geometric free-energy functional $\text{GFF} = \text{GEP} + \text{CEP}$ ranks the admissible boundary configurations by a single scalar; and GRIP — the lifecycle $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, lawful exploration to symmetry-breaking to stabilized attractor — fixes the order in which a branch is explored, committed to, and frozen. The next two sections put this machinery to work: first the fixed ledger it runs on, then the way it selects the branches of a black hole’s life.

13.1 Which particles? The observer-indexed vacuum, and horizon thermality as openness

One further consequence of the projection deserves a named place before the ledger, because the black-hole Parts use it on every page that mentions particle production. In a curved or horizon-bearing spacetime there is famously no preferred vacuum: inequivalent positive-frequency conventions define inequivalent particle notions, related by Bogoliubov mixing. In the present framework this ambiguity has a fixed address. On the canonical bulk geometry inherited from (Arisaka, 2026, A2-GG6) the constitutional time is globally static, so the bulk-level positive-frequency structure is unique; what varies from observer to observer is the *dictionary*: which modes an instrument counts is fixed by the observer’s own time and horizon structure, and the instruments are part of the OGGE boundary package by definition. “How many particles” is therefore an observer-indexed boundary observable, not a bulk quantity — the same status the framework already assigns to boundary states themselves. The physical output of the ambiguity is produced by the GDOS projection itself: tracing the inaccessible partner of a horizon’s squeezed pairing yields an exactly thermal state (at $T = a/2\pi$ for the Rindler pairing), so horizon thermality is constitutional openness doing its ordinary work rather than new physics. The full statements are carried in the real-bulk treatment of (Arisaka, 2026, A8-QM), whose standalone, full-length development is the Foundations Branch member paper GG-P4-Real (Arisaka, 2026, P4-Real). Two boundary cases are closed there at the model level, as entries R14–R15 of the series — the rotating case, where the ergoregion defeats the co-rotating selection and the uniqueness anchor moves to the asymptotic frame, superradiance becoming a symplectic pumping channel in the same GDOS slot; and the collapse epoch, which factorizes into a universal thermal channel times a decaying transient — and the present paper is their full-length home: they are Theorems 3 and 4 (T-BH-3, T-BH-4; §27), displayed with the whole series at §29. What remains conjectural is the exact curved-space existence theory — the field’s own, imported where used. The Parts that follow use both facts on nearly every page: the bulk structure is unique, and the thermal content of horizons is the partial trace.

14 Zero free parameters: BI-6 and the integer ledger

A framework can be elegant and still be useless if it leaves the numbers free. What makes GG-Theory predictive rather than merely descriptive is that it has *no continuous dials* to turn, and two inherited results enforce this.

The first is the anomaly-locked bulk. Global consistency of the six-force closure requires a modified six-dimensional Bianchi identity, BI-6, for the antisymmetric tensor $H_3 = dB_2$ (Arisaka, 2026, A3-BI6; Arisaka, 2026, A4-Euclid):

$$dH_3 = \text{tr } F_{SU(5)}^2 - F_B^2 - \frac{2}{5} F_B \wedge F_{Qi} - \frac{2}{5} F_{Qi}^2 + \text{tr } R^2,$$

whose coefficient vector $(a, b, c, d) = (1, -1, -\frac{2}{5}, -\frac{2}{5})$ is fixed uniquely by three independent conditions — local closure ($dX_4 = 0$), global integrality of $[X_4]$, and anomaly cancellation by inflow through the single channel $I_8 = \frac{1}{2}X_4 \wedge X_4$. There is no continuous deformation: the trace-ratio lock $|a/c| = |a/d| = 5/2$ is rigid. The ultraviolet is legal, and it is legal in exactly one way.

The second is the integer ledger. GG-A4-Euclid (Arisaka, 2026, A4-Euclid) proves that the same closure forces a unique tuple of whole numbers, the OGN-5 ledger

$$(M, N, L; J, K) = (1, 3, 5; 7, 35),$$

through a five-step Diophantine cascade in which each integer is fixed by the one before: $M = 1$ (a single Green–Schwarz tensor channel), $L = 5$ (the minimal simple unification $SU(5)$), $N = 3$ (the three colors, forced by the dual-lock closure, with the BI-6 trace ratio $-(L - N)/L = -2/5$ as GG-A3-BI6’s independent check; the generation count $N_g = 3$ is a separate quantity the ledger does not fix, GG-P2-N3), $J = 2N + 1 = 7$, and the closure integer

$$K = LJ = M^2 + N^2 + L^2 = 35.$$

That two independent arithmetic expressions for K agree — the *dual lock* — is what closes the ledger; it is generically false, and true here only because both are forced by the same axioms. These integers are not fitted; they are the arithmetic of the geometry. Together with the structural constants π and e and the dimensionful conversion factors $\{c, G_N, \hbar, k_B, Q_B\}$, they are the entire stock from which every dimensionless number in the cohort is built: the same ledger that fixes the twenty-eight Standard-Model parameters of GG-A9-SM (Arisaka, 2026, A9-SM) and the π -locked cosmic inventory of GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) fixes what this paper needs, with no new freedom in between; the discipline that admits a number only if it is a cost-minimum element over this ledger is the MSMF/OGN principle.

Two entries matter most at a horizon. The closure integer $K = 35$ sets the depth of the throat and with it the lock scale $M_{\text{lock}} \approx 3.85 \text{ TeV}$ — and it is the square of the corresponding lock length against an astrophysical horizon radius, $(\ell_{\text{lock}}/r_s)^2 \sim 10^{-46}$, that makes every GG-Theory correction to Kerr geometry invisible to any foreseeable instrument (Part VII): the framework’s observational signature is clean general relativity, and the ledger is why. And the integer $J = 7$ is the conjectured step of the primordial-black-hole mass ladder — a model-level conjecture whose exponent is not

derived here (§49) and whose kill row is F-BH-4 in the falsification table. The Planck–electroweak hierarchy, the silence of the ringdowns, and the dark-matter signature of the companion papers are literally the same integers, read on different strata.

15 The runtime: how a black hole’s branches are selected

The framework so far is static: a bulk, a boundary law, a fixed ledger. But a black hole has a *history* — it forms, accretes, merges, radiates, crosses the Hawking–Page and Page thresholds, and (if the universe permits) evaporates; something has to choose, event by event, which branch of the lawful possibilities actually happens. That choosing is what GG-Theory calls the runtime, and it is the last piece the module inherits. The runtime has two components, both proved in the pillar papers GG-A6-GRIP and GG-A7-Qi4 (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A7-Qi4): the GRIP cascade, which fixes the order in which lawful branches may be explored, committed to, and stabilized; and the GFF ledger, the scalar ranking functional $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ that says *which* branch is favored. The GFF ledger is, in form, a free-energy functional; its horizon instance — the generalized entropy, with the area term and the bulk-entanglement term assigned to their owners — is built in §33, where the information-theoretic lineage is recalled.

The two regimes, and where a descent may halt, are the vocabulary the reader will use most. In GG-Theory, every open system descends the single scalar cost functional just named, and the descent operates in exactly two regimes, split by one question: *does the landscape change?* GRIP-In is all dynamics on a *fixed* landscape: relaxation within a stable basin (a *Geometric Attractor*, GA), and passage *between* basins through whatever physical channel connects them. GRIP-Out is the complementary regime: a transient event that *reorganizes the landscape itself*. A descent halts only at the basin with no exit channel at all — the *Ground Geometric Attractor* (GGA), the unique deepest basin (the No-Exit theorem — carried at GG-A11-Origin’s own grade, with its above-particle-strata reach conjectural there — of (Arisaka, 2026, A11-Origin)). Among the many stable basins of a stratum, exactly one is this global minimum, $\text{GGA} := \arg \min F_{\text{GFF}}$; the distinction between a GA (a stable record, not eternal) and the GGA (the deepest state, with nothing to decay to) is load-bearing at a horizon, because this paper will show that today’s astrophysical black holes are *conditional* ground states — exitless only while $T_H < T_{\text{CMB}}$, with the condition expiring on a calculable cosmological schedule (§42.2).

Which branch, among the many lawful attractors of a stratum, is the one actually realized? The selector is the set of four quantum-information pillars Qi-4 of GG-A7-Qi4 (Arisaka, 2026, A7-Qi4). A configuration is the observed *ground* geometric attractor only if it passes all four admissibility filters at once: *Qi-QUA*, that it carry the locked four-current identity $J_E/Q_B = J_e = J_B = J_{Q_i}$ (the Arisaka equation); *Qi-EP*, that it be a genuine extremum of the one projected variational principle, not merely a local minimum of a coarse-grained potential; *Qi-CON*, that its scale signatures be those of a fixed point, surviving transport along r ; and *Qi-HAL*, that the runtime can actually halt there, with no admissible descent remaining. The observed value is the unique four-certificate survivor of its stratum — the same criterion that makes every electron, proton (Arisaka, 2026, P6-Proton), and Kaxion the identical object everywhere.

This is not an abstraction without a worked precedent: it is exactly how GG-A9-SM (Arisaka, 2026, A9-SM) fixed the Standard Model. There each parameter is a coordinate of a geometric attractor reached by this cascade — the electroweak vacuum selected as the stationary point of a Wilson–Hosotani holonomy potential (the GSSB step), the constants read off only after the system has descended onto that branch. GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) runs the identical machinery one story up, on the cosmological boundary. The present paper runs it at the horizon: formation, accretion, merger, the Hawking–Page transition, the Page turnover, and the evaporation endpoint are the GRIP readouts of the same runtime, and the reader will watch the two-regime grammar sort every one of them into its place in Parts IV–VI (§42, §36) — the sorting itself is new.

The electroweak paper also identified Wilson-line holonomy as the cleanest prototype of geometric spontaneous symmetry breaking (Arisaka, 2026, A9-SM; the runtime form is Arisaka, 2026, A6-GRIP). The black-hole paper sharpens that lesson into an identity: the same Wilson-line holonomies on S_θ^1 that serve the runtime as branch variables are, at a horizon, the very degrees of freedom whose traced entanglement is the Bekenstein–Hawking entropy (Part V). The branch variable and the state count are one geometry.

The whole inheritance can be stated once, as a single descent ladder:

$$\text{BI-6} \implies \text{TGG} \implies \text{CCI} \implies \text{GFF} \implies \text{GRIP}.$$

Read left to right: the anomaly-locked bulk fixes the legal ultraviolet stage; TGG packages the total closed ledger; CCI hands the four-dimensional observer a local boundary conservation law; GFF ranks the admissible branches of that boundary theory; and GRIP fixes the order in which they are explored, selected, and stabilized. The rule this ladder enforces is stated and proved as the inheritance theorems of the cosmology companion (T-Cosmos-1 and T-Cosmos-2 of Arisaka, 2026, A10-Cosmos); the present paper imports the rule rather than restating the theorems: *a black-hole source term is admissible only if it survives each layer of the ladder in order, and no continuous dial may be introduced along the way*. Figure 6 traces this descent for the black-hole observables of this paper, with each output carrying its register.

16 What this buys the black hole

Pulling the four ideas together gives the discipline that governs the rest of the paper. The bulk supplies one geometry; the boundary law makes ordinary general relativity — and, at a horizon, the forced trace — a projection of it; the integer ledger leaves no continuous dials; and the runtime selects which lawful history happens. The single operative consequence is a rule with no exceptions: *every black-hole source term must descend the inheritance chain* — bulk, projection, ledger, runtime — and no new dial may be introduced along the way. A quantity that cannot be placed on this chain is not derived; it is merely appended, and is inadmissible by construction.

Table 1 makes the discipline concrete, mapping the open problems of the next Part to the inherited components most relevant to each. It is a planning map, not a victory lap: it shows where each answer must come from, not that each answer is already complete.




Figure 6: The zero-parameter derivation tree, black-hole instance. Every black-hole observable of this paper (bottom tier) is traced back, through the GG-Theory machinery (middle tier), to a *fixed* input ledger (top tier) containing no continuous free parameters: the GU-5 dimensional constants, the proved OGN-5 integer ledger $(1, 3, 5; 7, 35)$, and the structural constants π and e . Output nodes are border-coded by the paper’s register, and the colors are not cosmetic. *Solid green (D)*: the thermal clock face T_H , proven at quasi-free grade (T-BH-5). *Dashed amber (D*)*: the exact coefficient (T-BH-1, conditional on its three stated premises), the Page time (L-BH-Fresh), the second law’s production (L-BH-Arrow, exact in its exchange class; the Wall class open), the de Sitter same-theorem entropy (C-BH-2), and the lock-scale suppression of Kerr corrections. *Dotted gray (open)*: the numerical species closure — the module’s registered internal kill row (F-BH-9), a work order, not a result. The full grades, premises, and exposures are audited row by row in the named-statement deck of the appendices.

Table 1: Working map from the horizon’s open problems (Part III) to the most relevant inherited components from the upstream papers. A planning table for the Parts that follow, not a claim that every problem is already solved. Nothing in the second column is derived here. Every entry names a component that already exists in an upstream paper, at the grade that paper assigns it, and the table’s whole content is the pairing. A reader who wants to know whether a given problem is closed should read the Part named in the first column rather than trust this row.

Problem (Part III)	Most relevant inherited inputs
The information paradox	OGGEP port semantics; GDOS open-system evolution; bulk unitarity (AX-1)
The Bekenstein–Hawking $1/4$	fully induced G_N (GGEP); canonical scale profile; heat-kernel/replica machinery (Part XII)
What the microstates are	compact fiber S^1_∂ ; Wilson-line holonomy; the GG-A8-QM hidden-fiber picture
The Page curve mechanism	GFF ledger and branch competition; Port-Freshness Clause (GG-A5-GDOS); GRIP two-regime grammar
Firewalls and the interior	bulk custody of entanglement; the observer-indexed dictionary (§13.1)
The second law’s direction	the constitutional arrow (GG-A5-GDOS); relative-entropy accounting; KMS clock face (T-BH-5)
The singularity	6D smoothness of the projection’s coordinate pathology (interpretive grade)
The universe’s own horizon	observer-anchored port; same-arena clause; the GG-A10-Cosmos cosmological boundary

Inheritance constrains, but it does not by itself prove; so every entry in what follows carries its register — derived, derived-conditional, model-level, open, conceded — said in those words at the claim and audited in the named-statement deck of the appendices. The border colors of Figure 6 are exactly these registers, and they are not cosmetic: they are the difference between a result and an intention.

There is a discipline hidden in this inheritance, and it deserves a plain statement. Because every tool used here was built — and paid for — in the particle-physics and cosmology companions, the black-hole module is never free to invent what it needs: whatever explains the horizon must already have explained the electron. That is why the module can be killed from outside black-hole physics (Part VIII), and why its cleanest prediction is silence — clean Kerr exteriors wherever instruments can see.

The framework is now in hand. The bulk, its boundary projection, the zero-parameter ledger, and the open-system runtime are everything the black-hole module inherits and everything it is allowed to use. The open problems of the horizon — stated first in the field’s own terms, then translated into this vocabulary — are the business of the next Part.

17 Part II in plain English: what the black-hole module may use, and what it is forbidden to invent

There is almost nothing about black holes in this Part, and the omission is the point.

Any framework can be made to explain a horizon if it is allowed to add whatever the horizon needs. This one is not allowed. The six-dimensional bulk, the projection that produces four-dimensional physics, the integer ledger that fixes the field content, and the open-system runtime that selects among branches were all fixed in the companion papers — for the electron, for the Standard Model, for the cosmos. The black-hole module inherits them entire and adds no postulate and no parameter of its own.

That constraint cuts both ways, and the Part says so rather than only claiming the flattering half. It is what makes the later results worth something: a coefficient derived from machinery built for other purposes is a prediction, while the same coefficient derived from machinery built to produce it is a restatement. It is also what makes the module fragile in an unusual way, because it can be killed from outside black-hole physics altogether. A collider result or a single proton decay would retire it without any telescope being involved.

The cleanest thing the inheritance buys is negative. Nothing in the framework permits a departure from Kerr outside the horizon, so wherever instruments can currently look, the module's prediction is silence. That is a weak prediction and a badly exposed one, and the Parts that follow do not try to make it sound like more than it is.

Part III

The Open Problems

Information Paradox, Page Curve, Firewalls, and the Entropy-Counting Bar

18 What the information paradox really asks

The information paradox, in its most refined form, is the following question: Given two theories — gravity (classical GR) and quantum mechanics — that disagree about the long-time behavior of a pure state subjected to gravitational collapse, which theory is correct, and what does the resolution tell us about quantum gravity?

Classical GR predicts information loss: the pure initial state collapses into a black hole whose subsequent evolution is determined only by mass, charge, and angular momentum (no-hair). Any internal structure is invisible from outside. When the black hole evaporates via Hawking radiation, there is no mechanism in GR by which the internal information could be recovered.

Quantum mechanics predicts information preservation: closed systems evolve unitarily; pure states remain pure; the entanglement entropy follows the Page curve.

The question is: which theory breaks down, and where?

18.1 Possible resolutions

Four classes of resolution have been seriously proposed over fifty years, and it is worth walking them in the order history did. The first is Hawking’s own 1976 position: information is simply *lost* — quantum gravity is fundamentally non-unitary, and unitarity must be abandoned. The second holds that information is *preserved through small corrections* to Hawking’s calculation: subtle quantum-field-theory-in-curved-spacetime effects that encode the information in radiation correlations. The third is *holography*: a complete quantum-gravity theory stores the bulk degrees of freedom on a boundary, so that unitarity is manifest — as it is in AdS/CFT. And the fourth invokes *new ultraviolet physics at the horizon itself*: firewalls, fuzzballs, or remnants that modify the standard picture exactly where it was thought safest.

Fifty years have sorted the four unevenly. The first position has been abandoned: evidence from AdS/CFT and related frameworks strongly suggests unitarity holds. The second is demonstrably insufficient — Mathur’s 2009 theorem, taken up just below, shows that no small correction to Hawking’s calculation can produce the Page curve (Mathur, 2009). The fourth is logically possible but violates the equivalence principle, and is widely considered unattractive for exactly that reason. The third — holography — is the current consensus, with the islands formula providing a concrete mechanism.

GG-Theory provides a specific realization of the third position, grounded in a complete 6D bulk theory rather than an abstract AdS/CFT duality.

18.2 Mathur’s theorem

Samir Mathur’s 2009 paper (Mathur, 2009) proved that *any* small perturbation of the Hawking calculation — including corrections from standard QFT in curved spacetime, arbitrary local modifications at the horizon, etc. — cannot produce the Page curve unless the corrections are of order unity. This is Mathur’s “small corrections theorem.”

The implication: whatever resolves the information paradox must involve *large* deviations from the standard Hawking picture. These deviations must be non-local, or non-perturbative, or geometric-at-the-horizon, or some combination. The islands formula realizes the “non-perturbative geometric at the horizon” option. GG-Theory realizes a “non-local from the bulk side of OGGE” option, which is equivalent to the islands formula but physically more concrete.

19 The Bekenstein–Hawking entropy counting problem

19.1 How many microstates?

If $S_{\text{BH}} = A/(4G_N\hbar)$ is the von Neumann entropy of the black hole, then the number of microstates is e^{S_{BH}/k_B} . For a solar-mass black hole ($M = M_\odot = 2 \times 10^{30}$ kg), this is approximately $e^{10^{77}}$ microstates. That is an enormous number.

What are these microstates? In ordinary thermodynamics the question has a mundane answer: microstates are specific configurations of molecules, atoms, or spins, and the entropy is the logarithm of the number of configurations compatible with a given macrostate — a gas has entropy because there are many ways to arrange its molecules. But a black hole seems to offer nothing to arrange. By no-hair, the macroscopic observables are (M, J, Q) ; and the same theorem says the horizon geometry depends on these parameters and nothing else. Where, then, do the microstates live? The tension is stark: the no-hair theorem says a black hole has no features; the entropy formula says it has $e^{10^{77}}$ of them.

This is the *black-hole microstate counting problem*. Different approaches give different answers:

Strominger-Vafa 1996. Andrew Strominger and Cumrun Vafa (Strominger and Vafa, 1996) showed in 1996 that for certain *extremal* string-theory black holes, the microstates can be counted as string configurations with specific brane charges. For the 5D extremal black hole they considered, the count was e^S microstates with $S = 2\pi\sqrt{Q_1 Q_5 n}$, matching the Bekenstein–Hawking area law exactly. This was celebrated as the first microscopic derivation of the Bekenstein–Hawking formula and became a milestone in string theory.

Limitations: (i) only extremal black holes (zero temperature); (ii) only in $d > 4$ dimensions; (iii) only in specific string-theory backgrounds; (iv) the microstates are *configurations in the dual theory*, not configurations of the black hole itself. The pattern is worth naming: the count works in a corner protected by supersymmetry, which guarantees the number of states cannot change as the coupling

is dialed from the calculable regime to the black-hole regime. Remove the protection — ask about the Schwarzschild or Kerr holes actually in the sky — and the method loses its footing; a quarter century of effort has extended it to *near*-extremal cousins, never to the generic astrophysical hole.

Loop quantum gravity counting. LQG (Ashtekar-Baez, Rovelli-Smolín 1996-1998) counts microstates as spin-network states on the horizon. The count gives $S_{\text{BH}} = A/(4G_N\hbar)$ up to a Barbero-Immirzi parameter γ , which is fixed by matching the LQG calculation to the Bekenstein formula — but this is a consistency check, not a derivation.

Fuzzball proposal. Samir Mathur has argued for nearly two decades that black-hole microstates are specific horizonless geometries — “fuzzballs” — that differ from each other on scales down to the Planck length. In this picture, there is no horizon at all; the region inside the “fuzzball radius” is a quantum superposition of horizonless configurations.

GG-Theory counting. In GG-Theory, microstates are 6D bulk configurations compatible with fixed 4D boundary data (M, J, Q) . Explicitly, a microstate is specified by: (i) the 4D boundary values of the metric, the gauge field, and the Higgs field; (ii) the radial profile of these fields in I_r ; (iii) the Wilson-line holonomy Θ_H on S_θ^1 . The 4D observer sees only the zero-mode components, which are fixed by (M, J, Q) . The microstate count is the entanglement entropy of these traced fiber degrees of freedom; by the induced-gravity argument of §31.5 it equals $A/4G_N$ exactly — about $e^{A/(4G_N\hbar)/k_B}$ states at leading order.

This is derived in full in §31.2–31.10.

19.2 The species problem

If there are N species of light fields in nature, a simple counting argument suggests that black-hole evaporation should produce them in proportion, giving more entropy per unit area than $1/(4G_N\hbar)$. This is the *species problem*: why is the coefficient $1/4$ and not $N/4$?

In the standard treatment the resolution is species-bound reasoning (Dvali and Redi, 2008): N light species lower the gravitational cutoff to $\Lambda_{\text{sp}} = M_{\text{Pl}}/\sqrt{N}$ and renormalize Newton’s constant, so the area coefficient stays $1/4$ once expressed in the infrared G_N .

In GG-Theory the species problem does not arise in the first place, because the coefficient is *not* a species count at all: Theorem 1 fixes $1/4 = 4\pi/16\pi$ as the regulator- and species-*independent* ratio of the conical-deficit curvature normalization to the Einstein–Hilbert normalization (§31.5). The OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ — with $N = 3$ its second entry and $L = 5$ fixing the $\text{SU}(5)$ rank — governs the field *content* (and hence the induced value of G_N through the species sum of Appendix R), but it does *not* enter the $1/4$. What the older species-bound argument had to engineer, the induced-gravity derivation makes automatic.

20 The singularity: what’s inside?

The classical black-hole interior contains a curvature singularity: a locus where the Kretschmann scalar diverges, tidal forces become infinite, and the classical theory breaks down. This is where quantum gravity is needed — the Planck-scale physics that GR cannot resolve.

The singularity inside the black hole is not observable from outside (causal disconnection by the horizon), but it affects the interpretation of the information paradox. If the singularity is where information “ends up,” and information is preserved, then the singularity must either: (i) not be a true endpoint, and information escapes via some non-classical channel; or (ii) be replaced by a smooth geometry in the full theory.

20.1 Singularity replacements

The proposals on the table span half a century of styles, and each removes the singularity by removing something different. A *big bounce* removes the endpoint: the singularity is replaced by a non-singular transition that connects the black-hole interior to another universe or region (Penrose 1965; Ashtekar–Bojowald 2006). The *fuzzball* removes the interior altogether: a quantum superposition of horizonless geometries extends up to a fuzzball radius $r_f > r_s$, and there is no singularity because there is no inside (Mathur 2008). The *firewall* removes the crossing: the horizon itself becomes a high-energy quantum wall, and the interior is never reached because infalling observers burn up at the wall (AMPS 2012). The *quantum-extremal-surface* reading of the islands program removes the interior’s independence: for a boundary observer the apparent interior is a bulk geometric structure redundantly encoded on the boundary (2019–2020). And the GG-Theory reading, developed in this paper, removes only the *observability*: the 4D interior is a radial-coordinate region where the scale profile e^{-r} is small, and the “singularity” is the limit in which that region is traced over by OGGEF, giving the boundary description an apparent termination.

In GG-Theory, the black-hole interior is *real*, unitary, and finite — it just cannot be observed from outside by the 4D observer because those directions are traced over by OGGEF. The apparent singularity is a coordinate artifact of the 4D projection, analogous to the apparent singularity at $r = r_s$ in Schwarzschild coordinates before Finkelstein regularized it.

21 Firewalls revisited

21.1 The AMPS argument in detail

Let B be a late-time Hawking quantum, A its interior partner, and R the earlier radiation. The postulates of black-hole complementarity (Susskind–Thorlacius–Uglum 1993) are four. The first is *unitarity*: Hawking radiation is the unitarily evolved pure state, so B is entangled with R . The second is that *quantum field theory in curved spacetime works at the horizon*: so B is entangled with its partner A on the other side. The third is *no-drama*: an infalling observer crossing the horizon sees nothing unusual. And the fourth is that the *black-hole entropy equals the Bekenstein–Hawking value*.

At late times (past Page), postulate 1 says $B \otimes R$ is an approximate Bell pair. Postulate 2 says $A \otimes B$ is a Bell pair. But B cannot be maximally entangled with both A and R (monogamy of entanglement). One of the four postulates must fail.

AMPS argued that postulate 3 (no-drama) is the weakest and most likely to fail. Infalling observers crossing the horizon would encounter a *firewall*: highly-energetic quanta that burn them up. This violates the equivalence principle (which is what no-drama is based on) but preserves unitarity, QFT-in-curved-spacetime, and the Bekenstein entropy.

21.2 The GG-Theory resolution

In GG-Theory, postulate 2 is what fails — but in a subtle way that preserves physical smoothness at the horizon.

The point is: B is not entangled with a partner A that exists independently inside the horizon. The “partner A ” is part of the bulk degrees of freedom that are traced over by OGGEF. Once we recognize that 4D “interior” is a partial description of the 6D bulk, the apparent bipartite A - B entanglement dissolves: B is entangled with the full bulk state, which is not separately identified with a 4D “interior observer.”

The infalling observer crosses the horizon and accesses more of the 6D bulk — specifically, regions with smaller scale profile. Nothing forces a firewall on them; they simply move to a region where the OGGEF projection would apply to a larger bulk sector.

In the language of the four postulates: unitarity is preserved, at the 6D level; quantum field theory near the horizon is preserved, with the one amendment that the “interior” is not an independent 4D sector; no-drama is not forced to fail — nothing in the entropy bookkeeping requires the horizon to be other than smooth; and the entropy remains the Bekenstein–Hawking value. Nothing on the list is sacrificed. The apparent violation of monogamy is an artifact of the 4D-only description; at the 6D bulk level, everything is consistent.

22 The Page curve derivation

The derivation of the Page curve remains one of the most important unsolved problems in standard AdS/CFT for *asymptotically-flat Schwarzschild black holes*. In 2D JT gravity coupled to matter, the Penington-Almheiri-et-al. derivation succeeds. In 4D asymptotically-flat Schwarzschild — the actual astrophysical case — the derivation is incomplete.

In GG-Theory, the Page curve’s shape follows from bulk unitarity plus CPTP projection, its turnover is named as a branch switch and dated by the expiry of port freshness, and the quantitative curve — a replica or island computation with a Page time — is not performed here. Part VI develops this explicitly.

23 The state of play

As of this writing — a dated snapshot, refreshed at each version of record — the consensus position in the theoretical physics community can be stated in five sentences. *Information is*

preserved: Hawking’s 1976 proposal that it is fundamentally lost is now considered incorrect, and the radiation must encode the initial information in correlations. *The Page curve is the correct entanglement-entropy curve*: pure-state evolution requires the tent-shaped rise and fall. *AdS/CFT remains the best-understood setting* in which to study the problem, with the islands formula delivering the Page curve in two-dimensional JT gravity. *The real-universe case — four-dimensional, asymptotically flat, Schwarzschild — is less well understood*, because its AdS/CFT dual is not known. And *no observational test of the information paradox has been performed or is currently feasible*: the Event Horizon Telescope and LIGO see the classical geometry; they do not constrain the entanglement structure of Hawking radiation.

GG-Theory enters this situation with a clear claim: the standard resolution *is essentially correct* but is couched in an abstract formalism (AdS/CFT, replica wormholes) that is not grounded in a physically realistic ultraviolet theory. GG-Theory provides the missing UV grounding — replacing the abstract AdS/CFT duality with a concrete 6D bulk for the *real* universe — and with it, the information paradox is resolved at the level of physical geometry, not formal duality.

The rest of this paper develops that claim explicitly.

24 Part III in plain English: the four things nobody has settled, stated in the field's own words first

Four problems, and this Part states each of them the way the field states it before translating any of it into this framework's vocabulary.

Where does the information go. If the hole radiates thermally and then evaporates, a pure state has become a mixed one, and quantum mechanics does not permit that. Mathur's theorem sharpened the puzzle into a dilemma: small corrections to Hawking's calculation cannot repair it, so something has to change by order one.

What are the microstates. The entropy formula counts something — the exponential of a quarter of the horizon area — and no calculation in the real universe has said what. String theory counts them in supersymmetric cases. Nothing counts them for the hole at the center of this galaxy.

What is inside. The singularity is the place where general relativity says of itself that it has stopped applying, and every proposal to replace it is a proposal about physics nobody has tested.

Is the horizon smooth. The firewall argument says you cannot have unitarity, an undramatic horizon and ordinary field theory outside it all at once. Something has to give, and fourteen years of argument have not settled which.

The Part's own verdict is worth restating plainly, because it is not the verdict a new framework usually offers. The standard resolution is judged essentially correct. What it is judged to lack is a grounding in the physics of the actual universe rather than in a formal duality. That is a smaller claim than replacing the standard account, and it is the claim the rest of the paper has to defend.

Part IV

The Horizon as Port

The Geometrically Forced Trace, Its Thermal Clock, and the Custody of the Anchor

This is the Part where the machinery meets its object. Parts I and III posed the horizon’s problems in the field’s own terms; Part II handed us the inherited framework — a unitary six-dimensional bulk, a boundary law under which every observer’s physics is an open-system projection, and a runtime that selects among lawful branches. Now we apply that framework to the black hole itself, and the reason for doing so is not that the black hole is one more application among many. It is that the black hole is the framework’s *native* object. The constitution’s central claim is that no physical system is closed and that observation *is* the bulk-to-boundary projection — a partial trace over what the observer cannot access. Everywhere else in physics, what gets traced out is a matter of circumstance: an experimenter decides what to measure and what to ignore, an environment couples to whatever it happens to touch. A horizon is different in exactly one respect, and it is the respect that matters. Behind a horizon, the inaccessible sector is not chosen, ignored, or accidental; it is fixed by the causal structure of spacetime itself. *The laboratory trace is chosen; the horizon trace is forced.* The black hole is therefore the one place in nature where the constitution’s central operation is performed by geometry rather than by an observer’s convention — which makes it not merely a consumer of the OGGEF/GDOS machinery but its cleanest test.

The Part unfolds in five steps. The opening section makes the claim precise: the horizon is an OGGEF *port*, and the claim is registered as a named key (L-BH-Port) with its grade stated. The next section computes what flows out of the port — Hawking radiation as the output of a completely positive channel, with the Bogoliubov calculus done by hand. Two theorems then close the cases the textbook treatment leaves loose, the rotating anchor and the collapse channel, at the model level and with their verifications printed. The thermal clock face turns the field’s analytic-continuation “tricks” into a theorem. And the closing section displays the real-bulk standing of everything the Part has built. Figure 7, below, is the whole argument in one picture; it is worth a long look before the equations begin.

25 The horizon as an OGGEF port

Figure 7 draws the contrast just stated: the laboratory trace is chosen; the horizon trace is forced. The arena — the canonical six-dimensional bulk and its scale profile — is fixed once, in Part II (§12); we work on it throughout.




Figure 7: The horizon as an OGGEp port. The six-dimensional bulk state is pure and unitary; the horizon is the locus at which OGGEp traces out the inaccessible sector — the 4D interior together with the non-spatial fiber (I_r, S_θ^1) — producing the open, completely positive boundary state ρ_∂ whose output is the Hawking radiation. The Bekenstein–Hawking entropy is the von Neumann entropy of ρ_∂ : the entanglement of the traced fiber, evaluated over the horizon.

25.1 The horizon is a special locus

At the horizon, four things happen at once in the 4D description. The metric has a coordinate singularity: the g^{rr} factor vanishes in Schwarzschild coordinates. The effective temperature, as measured by a static observer at infinity, diverges on approach, $T(r)/T_\infty \rightarrow \infty$ as $r \rightarrow r_s^+$. The scale profile in the bulk is strongly modified by the gravitational field. And — the decisive one — the 4D observer cannot access the interior $r < r_s$ at all: it is causally disconnected. The first three are symptoms; the fourth is the diagnosis.

In GDOS language, the horizon is a *port*: a specific locus at which the 4D observer’s accessible Hilbert space $\mathcal{H}_{\text{out}}$ is separated from an inaccessible Hilbert space $\mathcal{H}_{\text{in}}$. The OGGEp functor takes the full 6D state ρ_{bulk} and partial-traces over the inaccessible portion:

$$\rho_\partial^{\text{out}} = \text{Tr}_{\text{in}}(\rho_{\text{bulk}}), \quad (11)$$

where Tr_{in} denotes tracing over (a) the bulk compact directions S_θ^1 and the inaccessible radial tail, and (b) the 4D interior $r < r_s$. In the laboratory, *which* degrees of freedom are traced out is a choice — an experimenter decides what to measure and what to ignore. Behind a horizon, the bulk partners of the modes one sees are not merely ignored; they are unreachable as a matter of geometry. The partial trace stops being the observer’s bookkeeping convention and becomes a fact about the world’s causal wiring.

This is the central structural claim of GDOS for black holes: *the horizon plays the same role as Tr_{env} in open-system quantum dynamics. It is a one-way causal boundary that implements a CPTP projection on the 4D boundary Hilbert space.*

25.2 The geometrically visible port: black holes as the cleanest test of OGGE

The horizon is not merely *a* port; it is the one place in physics where the OGGE partial trace becomes *geometrically visible*. This deserves emphasis, because it is what makes the black hole the sharpest possible test of the open-system constitution (Arisaka, 2026, A5-GDOS).


In ordinary laboratory quantum mechanics the OGGE coarse-graining operates everywhere on the boundary $\mathcal{M}_4$ at every moment (AX-3), but the trace is *operationally undirected* — always on, but pointing nowhere in particular. A skeptic can always object: *in principle*, a sufficiently clever observer could extend her resolution into the bulk modes and recover the full state, so the coarse-graining is a choice, not a law. At a black-hole horizon that evasion is unavailable. The interior is causally disconnected from the exterior by the event horizon (AX-1, LC); no cleverness, no future technology, no idealized super-observer can recover the interior modes from outside. The inaccessibility is built into the causal structure of the spacetime.

The trace becomes ontological. The combined boundary state of an exterior observer is therefore


$$\rho_{\partial}^{\text{out}}(t) = \text{Tr}_{\text{in}}(\rho_{\text{bulk}}(t)), \quad (12)$$

where Tr_{in} now traces over *both* the non-spatial fiber (I_r, S_{θ}^1) that OGGE discards everywhere *and* the causally disconnected interior $r < r_s$. The first trace is the operational coarse-graining of AX-3; the second is geometrically forced. At the horizon the OGGE coarse-graining stops being merely operational and becomes *ontological* — a structural feature one can point to. This is the precise sense in which the black hole is the cleanest empirical test of AX-3: were AX-3 wrong, black-hole physics would have no constitutional account and the information paradox, the firewall, and the Page curve would be permanent puzzles; because AX-3 holds, all three dissolve at once, as the remainder of this Part shows (Arisaka, 2026, A5-GDOS). This paper registers the claim of this subsection under the key L-BH-Port: *at a black-hole horizon the OGGE trace is forced by the causal structure rather than chosen by the observer — the port is ontological, not operational*. Grade: Derived, given AX-1 and AX-3; it is the premise every information-channel result of this Part consumes.


Three settings, one trace. The same structure appears in three settings that differ only in *why* the trace is forced (Figure 8). In the laboratory the trace is operative but undirected. At a black-hole horizon it is forced by the causal structure. At the cosmological (de Sitter) horizon of radius c/H it is forced by cosmic expansion: modes beyond it recede faster than light and are permanently inaccessible. Black holes and the cosmological horizon are thus the two geometrically-forced strata of one constitutional move — which is why the same generalized-entropy ledger $\text{GFF} = F_{\text{GEP}} + F_{\text{CEP}}$ (§33) governs both, and why the present black-hole results carry over to de Sitter. In this sense

Lab: OGGEF invisible

Trace operative but undirected; observer could “in principle” resolve further

Black hole: forced

Trace forced by causal structure; no technology can recover the interior

Cosmos: forced

Trace forced by cosmic expansion; outside recedes faster than light

Figure 8: The OGGEF partial trace in three settings, adapted from (Arisaka, 2026, A5-GDOS). **Left:** ordinary laboratory — the trace over the non-spatial fiber operates everywhere but is operationally undirected (AX-3 operative but invisible). **Middle:** black hole — the trace over interior modes is geometrically forced by the causal structure of the Schwarzschild horizon. **Right:** cosmological horizon — the trace over modes beyond the de Sitter horizon ($r = c/H$) is forced by cosmic expansion. In the middle and right panels the constitutional content of AX-3 becomes ontologically visible; the black hole is the cleanest such test.

black holes are the *constitutional Rosetta stones* of the program: the place where an abstract axiom becomes a geometric inevitability.

26 Hawking radiation as the OGGEF output

26.1 The standard derivation

In Hawking’s 1974 calculation, the emission rate of quanta of frequency ω from a Schwarzschild black hole is

$$\frac{dN}{dt d\omega} = \frac{\Gamma_\ell(\omega)}{e^{\hbar\omega/k_B T_H} \mp 1}, \quad (13)$$

where $\Gamma_\ell(\omega)$ is the graybody factor for partial wave ℓ , and the $\mp$ is Bose or Fermi for each particle type. The spectrum is thermal to all orders in the saddle-point approximation of the gravitational path integral.

Hawking’s derivation uses QFT-in-curved-spacetime: the vacuum state of a quantum field near the horizon is not the Minkowski vacuum; it is thermally populated in the observer’s frame. This is closely analogous to the Unruh effect for accelerated observers.

26.2 GDOS re-derivation

In GDOS, the Hawking radiation is the output of the OGGEF channel $\mathcal{D}_{\text{BH}} : \rho_{\text{bulk}} \mapsto \rho_{\partial}^{\text{out}}$, which is CPTP, and the thermal spectrum follows in three moves. The bulk state near the horizon is approximately the Hartle–Hawking vacuum, which is thermal in the local-observer frame at inverse temperature $\beta_H = 2\pi/\kappa$, with κ the surface gravity. The OGGEF functor partial-traces over the interior, leaving a reduced state on the exterior. And for a canonical pair of bulk modes $(\hat{a}_{\text{in}}, \hat{a}_{\text{out}})$

related by a Bogoliubov transformation, that reduced state on $\hat{a}_{\text{out}}$ is thermal at temperature T_{H} .

This is essentially Hawking’s argument, reformulated in the GDOS language (Arisaka, 2026, A5-GDOS). The new content is that *the thermal character is automatic from the CPTP structure*; it does not need to be derived anew for each metric. Any bulk geometry with a horizon (Schwarzschild, Kerr, Reissner-Nordström, de Sitter, etc.) produces a thermal output spectrum at the appropriate temperature.

26.3 The Hawking temperature

The Hawking temperature is the scale-profile-determined temperature at the horizon:

$$k_B T_{\text{H}} = \frac{\hbar \kappa}{2\pi c}. \quad (14)$$

In GG-Theory, this follows from the Unruh–Tolman relation (developed in the companion GG-Theory thermodynamics paper): $T(r_*) = T_{\text{bulk}} e^{r_*}$, the gravitational blueshift set by the canonical scale profile e^{-2r} of the bulk. For a Schwarzschild black hole, the scale profile near the horizon is $\sqrt{g_{tt}} = \sqrt{1 - r_s/r}$, and the divergence at r_s is compensated by the acceleration of the proper-distance-from- horizon observer, giving a finite Hawking temperature at infinity.

The numerical match is automatic; GG-Theory does not need to add new physics at the horizon. The standard Hawking calculation is the zeroth-order GDOS prediction.

26.4 The explicit Bogoliubov derivation: why the trace is thermal

The preceding subsections established that the thermal spectrum is the automatic output of the OGGEp channel. Because a black-hole referee will want the mode-level computation, we give it explicitly: the thermality is *derived*, not posited, and the derivation makes plain that it is the same partial trace that produces mixedness everywhere in the framework.

Two mode bases and their mismatch. Expand the bulk free field on $\mathcal{M}_6$ near the horizon in two natural bases — the “in” modes $\{u_{\omega}^{\text{in}}\}$ regular on past null infinity, and the “out”/horizon modes $\{u_{\omega}^{\text{out}}\}$ natural to the static exterior. Because the horizon is a Killing bifurcation surface the two bases differ, related by a Bogoliubov transformation

$$a_{\omega}^{\text{out}} = \int d\omega' (\alpha_{\omega\omega'}^* a_{\omega'}^{\text{in}} - \beta_{\omega\omega'}^* a_{\omega'}^{\text{in}\dagger}), \quad (15)$$

and a non-zero β means the in-vacuum is populated in the out-basis.

The Planck factor. Matching positive-frequency modes by analytic continuation through the complexified horizon (the Damour–Ruffini/Wald prescription (Wald, 1984) — a step that is licensed, not assumed: Theorem 5(ii)) gives the universal ratio

$$\left| \frac{\beta_{\omega}}{\alpha_{\omega}} \right| = e^{-\pi\omega/\kappa} \quad \implies \quad \langle N_{\omega} \rangle = |\beta_{\omega}|^2 = \frac{1}{e^{2\pi\omega/\kappa} - 1} = \frac{1}{e^{\hbar\omega/k_B T_{\text{H}}} - 1}, \quad (16)$$

with surface gravity κ and $k_B T_H = \hbar\kappa/2\pi c$ — the Planck spectrum (Hawking, 1975; Unruh, 1976). (Mode exponents throughout this Part are written in geometric units, $e^{2\pi\omega/\kappa}$; SI factors of c are restored only in temperature statements, via $\kappa_{\text{acc}} = c^2\kappa_{\text{geom}}$.)

The squeezed-state structure and the OGGEF trace. The content of the Bogoliubov mismatch is that the global Unruh vacuum is a two-mode squeezed state pairing each exterior mode with an interior partner across the horizon,

$$|0\rangle_U = \prod_{\omega} \frac{1}{\cosh r_{\omega}} \sum_{n=0}^{\infty} (\tanh r_{\omega})^n |n\rangle_{\omega}^{\text{out}} \otimes |n\rangle_{\omega}^{\text{in}}, \quad \tanh r_{\omega} = e^{-\pi\omega/\kappa}. \quad (17)$$

The OGGEF partial trace over the causally inaccessible interior partners then produces the exterior state by *direct computation*:

$$\rho_{\text{out}} = \text{Tr}_{\text{in}} |0\rangle_U \langle 0|_U = \prod_{\omega} (1 - e^{-\hbar\omega/k_B T_H}) e^{-\hbar\omega N_{\omega}/k_B T_H} = \frac{e^{-H_{\text{out}}/k_B T_H}}{Z}. \quad (18)$$

This is the explicit statement that Hawking thermality *is* the OGGEF partial trace of a pure two-mode-squeezed bulk state — the same operation that yields a mixed boundary state from a pure bulk everywhere in the cohort (§25.2), here with the trace geometrically forced. The observed flux carries the graybody transmission $\Gamma_{\ell}(\omega)$ of the exterior curvature barrier,

$$\frac{dN_{\omega}}{d\omega dt} = \frac{\Gamma_{\ell}(\omega)}{e^{\hbar\omega/k_B T_H} \mp 1}. \quad (19)$$

Why every horizon radiates. Nothing in this computation used the Schwarzschild metric; it used only that a horizon induces a Bogoliubov mismatch with a single scale κ . Hence Schwarzschild, Kerr, Reissner–Nordström, and the de Sitter cosmological horizon all radiate thermally at $T = \hbar\kappa/2\pi c k_B$, each with its own surface gravity — a fact we use in the next subsection and in the cosmological-horizon treatment. The pedigree of every analytic step used here — why the positive-frequency selection is a derived clock, why the continuation through the complexified horizon is lawful, and why the trace is exactly the thermal state of that clock — is recorded as Theorem 5 (T-BH-5), once the rotating and collapse cases have been brought into the same frame (§28).

27 The rotating anchor and the collapse channel: two theorems at the model level

The vacuum story of this Part has so far been told for static exteriors: the Killing frame is globally timelike, the positivity selection fixes the exterior complex structure uniquely, the vacuum multiplicity (Boulware, Unruh, Hartle–Hawking) is a port assignment at the horizon, and thermality is the GDOS partial trace itself — the same statement the cosmology companion makes in general form: observer-adapted vacua are *dictionary* data, and a change of dictionary is a Gaussian channel acting on the boundary description, not an ambiguity in the bulk state (Arisaka, 2026, A10-Cosmos).

Two loose ends remain, and both are genuine: *rotation*, where the ergoregion makes the co-rotating Killing energy indefinite and the literature’s exact result is that the Kerr exterior admits no Hadamard state invariant under the horizon Killing flow — no Hartle–Hawking analogue (Kay and Wald, 1991); and *collapse*, where nothing is stationary and no Killing selection exists during infall. This section closes both at the model level — in the same custody pattern as the horizon port — as entries R14 and R15 of the cohort’s real-bulk series (Arisaka, 2026, A8-QM; the full-length member-paper treatment of the series is GG-P4-Real, Arisaka, 2026, P4-Real). The two statements are verified numerically in the program’s verification suite, available with the source through the program’s site (www.arisaka-gg.org), and the honest regime declarations are part of the statements.

Theorem 3 (The rotating anchor (T-BH-3)). *In the superradiance model class — asymptotic modes of positive energy coupled to ergoregion modes of negative co-rotating energy through the standard quadratic (two-mode-squeezing) interaction: (i) the co-rotating quadratic form is indefinite, so the positivity selection of the exterior complex structure is unavailable in the rotating frame — the model form of the Kay–Wald obstruction; (ii) the asymptotic frame’s quadratic form is positive, and its polar phase selects the asymptotic J uniquely — the anchor moves to spatial infinity, where every instrument that will ever report a measurement actually sits; (iii) read through that fixed asymptotic J , the rotating dynamics is an exactly symplectic Gaussian channel which amplifies populations at the superradiant rate ($\langle n \rangle = \sinh^2(gt)$ at resonance) while leaving J untouched: superradiant energy extraction is GDOS pumping, exactly as horizon thermality is GDOS tracing, and the rotating vacuum multiplicity is port assignment relative to the asymptotic anchor.*

Proof. (i) In the rotating frame the superradiant band contributes modes of co-rotating energy $\omega - m\Omega_H < 0$ (Starobinsky, 1973; Unruh, 1974), so the quadratic form carries both signs; positivity — the hypothesis of the uniqueness selection — fails, and no invariant complex structure with positive energy exists there. In the two-mode model the form is $\omega_1(x_1^2 + p_1^2)/2 - \omega_2(x_2^2 + p_2^2)/2 + g(x_1x_2 - p_1p_2)$, indefinite for all g . (ii) At spatial infinity the frame is non-rotating and the form $\omega_1(x_1^2 + p_1^2)/2 + \omega_2(x_2^2 + p_2^2)/2$ is positive; the polar phase of its generator is the standard J , unique by the positivity selection. (iii) The Heisenberg flow of the full generator is symplectic by construction; the interaction term $g(a^\dagger b^\dagger + ab)$ generates two-mode squeezing, so the evolution read on the asymptotic J -space is a Gaussian completely-positive channel that grows $\langle a^\dagger a \rangle = \sinh^2(gt)$ at resonance while the covariance-independent reading structure — the asymptotic J — is fixed by definition of the frame. Numerically: the form’s indefiniteness, the polar-phase uniqueness at infinity, the exact symplecticity (10^{-9}), the state validity (uncertainty relation), and the $\sinh^2(gt)$ law (4×10^{-14}) are all verified. What is *not* claimed: the Hadamard analysis of the exact Kerr exterior — which asymptotic constructions are physically admissible there — is the field’s own active area (Kay and Wald, 1991), imported where used. The custody clauses are insensitive to that choice: the choice *is* the port. ■

Theorem 4 (The collapse channel (T-BH-4)). *For quadratic field dynamics static in the far past and carrying the exponential horizon redshift $x \sim e^{-\kappa u}$ at late times — the moving-mirror class (Davies and Fulling, 1976), the model form of gravitational collapse: (i) the in-frame complex structure is selected uniquely at past infinity, where the dynamics is static; (ii) the out/in Bogoliubov*

overlaps reduce to $I_{\pm} = \int_0^{\infty} x^{i\omega/\kappa-1} e^{\pm i\omega'x} dx$, and $|I_+/I_-|^2 = e^{-2\pi\omega/\kappa}$ exactly: the Planck factor at $T_H = \kappa/2\pi$ is the branch-cut asymmetry of the exponential redshift — a real, classical, kinematic datum; (iii) for a trajectory static early and exponential late, the same ratio converges to the Planck value as ω' grows: the late-time out-field sees only the exponential region, so the out-state factorizes into the universal thermal channel (at T_H , independent of the early trajectory) times a transient carrying the collapse details, whose contribution decays. The collapse epoch populates states through a Gaussian channel and never touches the origin of the complex structure.

Proof. (i) is the static selection applied at past infinity. (ii) is the classical evaluation underlying Davies and Fulling (1976) and Hawking (1975): with $\nu = \omega/\kappa$, $I_{\pm} = \Gamma(i\nu) (\mp i\omega')^{-i\nu}$ up to a common factor (principal branch), and $|(\mp i\omega')^{-i\nu}| = e^{\mp\nu\pi/2}$, so the ratio of moduli squared is $e^{-2\pi\nu}$ — verified by direct quadrature (with the oscillatory endpoint handled analytically) to better than 10^{-3} relative at $\nu = \frac{1}{2}, 1, 2$. (iii) The factor $e^{\pm i\omega'x}$ confines the integral to $x \lesssim 1/\omega'$; for ω' large this lies entirely in the exponential region, and the early-trajectory contribution decays as an endpoint term — verified on a static-then-exponential trajectory, where the relative deviation from the Planck ratio falls from 0.13 at $\omega' = 3\kappa$ to below 10^{-2} at $\omega' = 30\kappa$. The full gravitating statement — that a collapsing star’s late-time radiation carries exactly this ratio — is Hawking’s derivation (Hawking, 1975), imported with attribution; what the model class proves is the factorization *structure*: universal late-time channel, decaying transient, custody preserved. ■

The custody close. With these two statements the custody pattern of the cohort — *the dictionary owns the sectors; the geometry owns the i* — holds in five instances: the charge angle, the horizon port, the infrared cloud, the rotating anchor, and the collapse channel — the first three superselection structures proper (an unreadable relative phase between sectors), the last two anchor-and-channel instances of the same division of custody. In every one, what varies between admissible descriptions is a c-number port datum owned by the operational dictionary — a winding class, a horizon boundary condition, a soft cloud, a stationary congruence, an in-vacuum — and in every one the derived complex structure passes through untouched. Black holes, on this reading, are not where the quantum formalism breaks; they are where its custody division shows most plainly.

Physics intuition. A rotating black hole is a harbor with a spinning quay: no clock mounted on the quay keeps honest time everywhere (that is the ergoregion), so the ship’s chronometer is set out at sea, where the water is calm — the asymptotic frame — and the spinning quay can then do only one thing to the cargo: pump it, which is superradiance. Collapse is the harbor being built while the ship watches: however violent the construction, what reaches the ship at late times is a universal thermal wake at the horizon temperature, plus a transient that carries the construction noise and dies away. In both cases the pumping and the wake are freight — states, populations, energy — and never touch the chronometer. The dictionary owns the freight manifests; the geometry owns the clock.

28 The thermal clock face as a theorem

Every thermality statement in this Part has run on one analytic device, bought in three pieces: a positive-frequency selection (which modes are “particles”), an analytic continuation through the complexified horizon (the Planck factor of §26.4), and a periodicity in imaginary time (the Euclidean arguments of Appendices H and I). In the older literature each piece is a prescription — serviceable, correct, and unexplained: the thermal character appears to be a statement about a cleverly chosen vacuum. It is not. It is a statement about the partial trace, made with a clock the bulk already owns — and the cohort now holds the theorem that says so. The Wick-quartet theorem T-P5-B1 of the CPT companion (Arisaka, 2026, P5-CPT) proves that for quasi-free fields on a static arena four data are one datum, each determining the others with orientations included: the positive-energy complex structure (the clock’s rotor), the harmonic-conjugate splitting (positive-frequency selection), the imaginary-time contraction semigroup (Wick rotation), and the holomorphic time tube (the lawful meaning of *it*) — all four notations of a single analytic fact, the one-sidedness of the clock’s spectrum. This section consumes that theorem at the horizon.

Theorem 5 (The thermal clock face (T-BH-5)). *For quasi-free bulk fields on the static exterior of a bifurcate Killing horizon of surface gravity κ — the arena of §26.4, with rotation handled by the asymptotic anchor of Theorem 3: (i) the exterior positive-frequency structure is the polar rotor of the Killing generator, the derived complex structure of the causal clock (Arisaka, 2026, P4-Real), selected uniquely by the positivity clause — an output of the constitution, not an input to the calculation; (ii) by T-P5-B1 the exterior evolution extends to a contraction semigroup in imaginary Killing time, so the analytic continuation through the complexified horizon used in Eq. (16) and the imaginary-time section used in Appendices H and I exist as theorems of the one-sided spectrum; no continuation prescription is imported; (iii) the regular quasi-free state of the exterior algebra — where such a state exists — is the thermal (KMS) state of that clock at the period fixed by the horizon’s smooth Euclidean cap, $\beta = 2\pi/\kappa$ (Kay and Wald, 1991); the Bogoliubov ratio $|\beta_\omega/\alpha_\omega| = e^{-\pi\omega/\kappa}$ and the thermal trace of Eq. (18) are that KMS condition displayed mode by mode. The thermal clock face on which this paper’s Hawking temperature runs is therefore proven, not assumed: the quartet supplies the face; the geometry contributes only the reading, $2\pi/\kappa$.*

Proof. (i) is the positivity selection already used throughout this Part (the static exterior Killing frame is globally timelike outside the ergoregion; for rotation the selection moves to the asymptotic frame, Theorem 3(ii)), read through the polar decomposition of T-P5-B1: a skew generator splits uniquely into rotor times speed, and the rotor is the clock. (ii) is T-P5-B1 applied verbatim: on the physical exterior modes the generator’s spectrum is one-sided, so the evolution continues to the bounded half-plane and the contraction semigroup exists; a two-sided spectrum would admit no such shore (the Paley–Wiener obstruction). (iii) composes two quasi-free-grade results: the uniqueness-and-thermality theorems on bifurcate Killing horizons (Kay and Wald, 1991), which fix the regular exterior state as KMS at $2\pi/\kappa$, and the mode-level display of §26.4, whose reduced two-mode-squeezed state satisfies the KMS ratio $\langle aa^\dagger \rangle / \langle a^\dagger a \rangle = e^{\beta\omega}$ exactly. Numerically: the polar split, the KMS ratio, and the Planck occupation are verified to machine precision (10^{-15} relative) at $\omega/\kappa = \frac{1}{2}, 1, 2$, and the one-sidedness obstruction is displayed explicitly, in the program’s verification

suite (www.arisaka-gg.org). What is *not* claimed: the gravitating, back-reacting extension beyond the quasi-free class — imported where used, exactly as in Theorems 3 and 4. ■

What the theorem retires. Three habits of the horizon literature become clauses rather than assumptions. The “prescription” of matching positive-frequency modes by continuation is clause (ii). The “trick” of demanding regularity in imaginary time is the existence half of clause (iii) — the section exists because the spectrum is one-sided; the geometry’s regularity datum only sets the period. And the choice-of-vacuum language (“the thermal character of the Hartle–Hawking state”) inverts: thermality is not a property a vacuum happens to have but the KMS face every regular quasi-free exterior state must show, because the clock it faces is derived. The custody division of §27 is unchanged — T-BH-5 adds no sixth instance; it names the clock that all five instances leave untouched. And the fiber supplies no rival: the phase circle’s winding spectrum is two-sided, so it damps no direction and admits no half-plane — among the bulk’s directions, time alone has a shore, which is why the horizon’s thermality is the *clock’s* face and never the fiber’s.

Physics intuition. A harbor thermometer does not manufacture temperature; it reads the water it is dipped in. The older derivations bought the thermometer in parts — a dial here, a damping law there, a calibration at the end — and the parts always fitted, without anyone saying why. T-P5-B1 is the statement that the parts are one instrument — rotor, shore, and holomorphic tube, three notations of the single clock the causal structure hands every static observer — and the thermal face is what this paper proves on that instrument at the horizon (T-BH-5, composing the quartet with [Kay and Wald, 1991](#)). The black hole contributes exactly one number to the reading — its surface gravity. Everything else about Hawking thermality was already in the clock.

29 The real-bulk register at the horizon

The cohort maintains a running register of results in which a quantity the boundary literature treats as intrinsically complex is exhibited as real bulk data read through the derived dictionary — the *real-bulk series*, whose home is the quantum-mechanics companion ([Arisaka, 2026, A8-QM](#)) and whose full-length member treatment is GG-P4-Real ([Arisaka, 2026, P4-Real](#)). This paper holds two entries of that series, claimed at §27 and until now never displayed. We display them, together with the further entry that this paper’s entropy chain has earned — proposed to the series’ home papers and, on their ruling of July 2026, registered as Statement RB-13 of the series home ([Arisaka, 2026, A8-QM](#); registrational: no new theorem, no new custody instance, and no series-entry number — the proven series stands at R15).

Table 2: The real-bulk series entries held by this module. Each exhibits a celebrated “quantum” datum as real, classical, or kinematic bulk structure, with the derived complex window untouched. Two entries, and both are model-level rather than general. The real-bulk series claims that a datum usually called quantum is in fact real, classical or kinematic bulk structure; what is verified here is that claim inside a solvable model, at the numerical tolerance the last column prints. That is a demonstration and not a theorem about generic black holes, and the grade column says so rather than leaving it to be inferred.

Entry	Statement	Where proven	Grade
R14	The rotating anchor: the ergoregion defeats the rotating-frame positivity selection for a real reason (an indefinite quadratic form); the asymptotic frame’s polar phase selects J uniquely; superradiance is exactly symplectic — real — Gaussian pumping	Theorem 3 (T-BH-3)	Model level, verified (10^{-9} , 4×10^{-14})
R15	The collapse channel: the Planck factor is the branch-cut asymmetry of the exponential redshift — a real, classical, kinematic datum — and the late-time channel factorizes universally	Theorem 4 (T-BH-4)	Model level, verified (10^{-3})

The registered entry: horizon entropy as real-section geometry (Statement RB-13).

The following is registered in the series’ home paper as Statement RB-13 (Arisaka, 2026, A8-QM; registrational, carrying no series-entry number). *The horizon’s entropy data are real-section geometry.* Every geometric object in the exact-1/4 chain is real: the conical deficit of Step 3 is a curvature datum of the real Euclidean section (§31.5); the area A is a real functional of the real metric; the entanglement the entropy measures is of the three real bulk factors, traced at the horizon (Step 4); and the one place the imaginary unit appears — the analytic continuation and the KMS face — is the boundary’s derived clock (Theorem 5), whose shore belongs to time and never to the fiber. In one line: $A/4G_N$ is real geometry read through a derived complex window — the dictionary owns the sectors; the geometry owns the i . No new custody instance is claimed (the count stands at five, §27); the entry names the real-bulk standing of the entropy chain itself.

Physics intuition. A ledger written in black ink, audited through tinted glass. Every quantity on the page — the deficit angle, the area, the traced fiber count — is black-ink real; the tint is the boundary’s one derived window, the clock’s i , and it colors the reading without ever touching the page. The real-bulk series is the cohort’s running proof that the tint comes off; the horizon, where the ledger’s most famous number lives, is its sharpest page.

The Part in one breath, and the door to the next. A horizon, this Part has argued, is not a wall but a port: the one locus where the constitution’s central operation — the partial trace that elsewhere an experimenter chooses — is performed by the causal structure itself (L-BH-Port). Everything thermal then followed from that single fact, with no further input: the Hawking spectrum as the channel’s output; the temperature as the clock face of a proven KMS state, the old continuation tricks retired by theorem (T-BH-5); the rotating and collapsing cases closed at the model level, with their custody clauses intact (T-BH-3, T-BH-4); and the whole chain standing on

real-section geometry read through one derived complex window (Statement RB-13). The physics intuition to carry forward is the simplest one in the paper: *a black hole is an open system, and geometry is its experimenter*. What the port has not yet told us is the size of the account it keeps. The radiation is thermal, and thermal states carry entropy; the horizon’s entropy has a famous, exact, and — until now — underived coefficient. Counting what the port traces, and proving that the count is $A/4G_N$ with the $1/4$ exact, is the business of the next Part.

30 Part IV in plain English: why the thermality is forced rather than derived, and what that does not buy

A horizon is a one-way surface: what crosses it does not come back. That is a statement about causal structure, and it is as true in the six-dimensional bulk as it is in four dimensions.

The framework’s central operation is a projection from the bulk to a four-dimensional boundary, and every projection discards something. What this Part observes is that at a horizon the discarding is not a modeling choice but a fact about the geometry: the region behind the surface is causally inaccessible, so any description written on this side must trace over it. A trace over part of an entangled state gives a mixed state. In this reading the thermality of Hawking radiation is therefore not a result to be derived but a consequence of the one operation the geometry leaves no choice about.

That is a strong claim, and it is worth naming what it does not buy. It says nothing about the value of the temperature until the calculation is done, so the Part does the calculation both ways — the standard Bogoliubov route and the open-system one — and shows they agree. It also cannot be tested by any instrument now or soon, which is why the Part’s two theorems are stated at model level, with the tolerance of their numerical verification printed rather than described.

What the port has not said is how large the account it keeps is. The horizon has an entropy, that entropy has a coefficient, and the coefficient has been exact and unexplained for fifty years. That is the next Part.

Part V

The Exact $1/4$

The Bekenstein–Hawking Coefficient from Induced Gravity, with Its Honest Accounting

This is the Part the paper exists to deliver. Part IV established that the horizon is a forced OGGE port and that what leaves it is thermal; the account the port keeps — its entropy — was written down by Bekenstein and Hawking half a century ago, with the most famous coefficient in the subject standing in front: the $1/4$. Part I told the story of how that number resisted derivation for the black holes that actually exist, and Part III stated the bar precisely: derive the coefficient, from a microscopic account, for the generic hole, with nothing tuned. This Part clears that bar. The instrument is the constitution’s second axiom: gravity is induced — there is no bare four-dimensional Newton constant — so the same bulk fluctuations that make gravity stiff make the horizon entangled, and the two quantities are locked at the source rather than matched at the end. The plan runs in four steps and touches no dial: the induced Newton constant is computed from the canonical scale profile; one heat kernel is shown to price two quantities at once; the conical replica evaluates their ratio in closed form, $1/4 = 4\pi/16\pi$; and the microstates are named — the hidden-fiber holonomies the port traces. The Part closes with the honest accounting: what is rigorous, what is conditional, and the one computation the module still owes itself.

31 The Bekenstein–Hawking Entropy and the Exact $1/4$ from Induced Gravity

31.1 Entropy as entanglement

The Bekenstein–Hawking entropy is

$$S_{\text{BH}} = \frac{Ac^3}{4G_N\hbar} k_B, \quad (20)$$

where A is the horizon area. In GDOS, this is the *entanglement entropy* of the 4D exterior with the bulk interior:

$$S_{\text{BH}} = S_{\text{vN}}(\rho_{\partial}^{\text{out}}) = -\text{Tr}(\rho_{\partial}^{\text{out}} \ln \rho_{\partial}^{\text{out}}) k_B. \quad (21)$$

The interior here includes both the 4D interior $r < r_s$ and the non-spatial fiber — the phase circle S_{θ}^1 and the inaccessible radial tail of the scale direction I_r . The rest of this Part computes this entanglement entropy from first principles and shows that it equals $A/4G_N$ *exactly*, coefficient included, by the induced-gravity mechanism.

31.2 Strategy: gravity is induced, so the horizon entropy *is* the entanglement

Equation (21) identifies S_{BH} as an entanglement entropy. But the entanglement entropy of quantum fields across a surface is famously ultraviolet-divergent and, in an ordinary field theory, scheme-dependent: it is a *correction* to a pre-existing $A/4G_N$, not the quantity itself (Bombelli et al., 1986; Susskind and Uglum, 1994). What turns it into the exact Bekenstein–Hawking value — coefficient and all — is a structural feature that GG-Theory supplies and ordinary field theory does not:

there is no bare four-dimensional Newton constant.

By GGEP (AX-2) gravity is the geometry of the bulk, not a fundamental boundary field carrying its own coupling; the four-dimensional gravitational action that a boundary observer reads is generated *entirely* by integrating out the bulk fluctuation modes. This is Sakharov’s induced-gravity hypothesis (Sakharov, 1968), here promoted from hypothesis to a consequence of the constitution. Its decisive corollary, established below, is that the divergence of the entanglement entropy is not a correction to $1/4G_N$ — it *is* $1/4G_N$ (Frolov and Fursaev, 1997; Jacobson, 1994). The argument is three analytic steps. First, the induced Newton constant is generated by the canonical scale profile, finite and ultraviolet-localized (§31.3). Second, the *same* bulk modes that induce $1/G_N$ produce the horizon entanglement entropy, through the *same* heat-kernel coefficient, so their ratio is fixed once and for all (§31.4). Third, the conical replica method evaluates that ratio in closed form: $S_{\text{BH}} = A/4G_N$, with the coefficient $1/4 = 4\pi/16\pi$ emerging *regulator-independently* (§31.5).

31.3 Step 1 — the induced Newton constant from the canonical scale profile

Write the six-dimensional gravitational action that GGEP attaches to the bulk,

$$I_6 = -\frac{1}{16\pi G_6} \int_{\mathcal{M}_6} d^6 X \sqrt{g_6} R_6, \quad (22)$$

and reduce it on the non-spatial fiber using the canonical metric (10), $g_{\mu\nu}^{(6)}(X) = e^{-2r} g_{\mu\nu}(x)$. The measure and curvature reduce as

$$\sqrt{g_6} = e^{-4r} \sqrt{g_4}, \quad R_6 \supset e^{+2r} R_4[g], \quad (23)$$

the factor e^{+2r} arising because the four-dimensional Ricci scalar is built with the scale-profile inverse metric $g_{(6)}^{\mu\nu} = e^{+2r} g^{\mu\nu}$. Hence the Einstein–Hilbert term of the boundary theory is

$$I_6 \supset -\frac{1}{16\pi G_6} \left(\int_0^{2\pi} d\theta \int_0^K dr e^{-2r} \right) \int d^4 x \sqrt{g_4} R_4, \quad (24)$$

with the two fiber integrals *finite and parameter-free* (Theorem 2: no free k , R_θ , or r_c):

$$\int_0^{2\pi} d\theta = 2\pi, \quad \int_0^K e^{-2r} dr = \frac{1}{2}(1 - e^{-2K}). \quad (25)$$

Reading off the coefficient of $\frac{1}{16\pi} \int \sqrt{g_4} R_4$,

$$\boxed{\frac{1}{G_N} = \frac{\pi(1 - e^{-2K})}{G_6}} \quad (\text{fundamental scale restoring units}). \quad (26)$$

Three features matter for what follows. (i) G_N is *finite*: the radial integral converges because the scale profile e^{-2r} is integrable on $[0, K]$. (ii) It is *ultraviolet-localized*: the integrand e^{-2r} is peaked at the UV boundary $r = 0$, so Newton's constant is generated near E_{GUT} — the geometric statement that gravity is a UV phenomenon. (iii) It is *induced*, not posited: G_N is a property of the scale profile, exactly as GGEP demands. This is the *same* boundary Newton constant obtained from the scale-overlap reduction of the boundary-gravity derivation in GG-A5-GDOS (Arisaka, 2026, A5-GDOS): both read $1/G_N$ off the identical canonical scale profile of GG-A2-GG6, so the two papers' Newton constants coincide by construction rather than by assumption. Moreover, the $A/4$ coefficient established below depends only on the heat-kernel *ratio* of Step 2, not on the absolute normalization of G_N , so the entropy law is robust even where that normalization is later refined. At one loop the same structure is reproduced quantum-mechanically — the bulk modes' determinant generates (24) with a coefficient set by their spectrum — which is the form we need next.

31.4 Step 2 — one heat kernel, two quantities

Both the induced Newton constant and the horizon entanglement entropy are controlled by a *single* object: the heat-kernel expansion of the bulk-mode fluctuation operator Δ on the four-dimensional boundary,

$$\text{Tr } e^{-s\Delta} = \frac{1}{(4\pi s)^2} \int d^4x \sqrt{g_4} [a_0 + a_1 s + a_2 s^2 + \cdots], \quad a_1 = \left(\frac{1}{6} - \xi\right) R, \quad (27)$$

with ξ the curvature coupling ($\xi = 0$ for a minimal scalar). The coefficient a_1 — the coefficient of R — is the common ancestor of both quantities.

The induced $1/G_N$. The one-loop effective action is $W = \frac{1}{2} \ln \det \Delta = -\frac{1}{2} \int_{\epsilon^2}^{\infty} \frac{ds}{s} \text{Tr } e^{-s\Delta}$, regulated by a proper-time cutoff $s \geq \epsilon^2$ (ϵ the physical short-distance length). The a_1 term contributes the Einstein–Hilbert piece

$$W \supset -\frac{1}{2} \frac{1}{(4\pi)^2} \left(\frac{1}{6} - \xi\right) \int d^4x \sqrt{g_4} R \int_{\epsilon^2}^{\infty} \frac{ds}{s^2} = -\frac{1}{16\pi} \left[\frac{1}{12\pi\epsilon^2}\right]_{\xi=0} \int d^4x \sqrt{g_4} R, \quad (28)$$

so that, per minimal scalar,

$$\frac{1}{G_{\text{ind}}} = \frac{1}{12\pi \epsilon^2}, \quad \implies \quad \frac{1}{4G_{\text{ind}}} = \frac{1}{48\pi \epsilon^2}. \quad (29)$$

The entanglement entropy. The leading (area) term of the entanglement entropy across the horizon, computed from the same heat kernel by the conical method (next subsection), is

$$\frac{dS_{\text{ent}}}{dA} = \frac{1}{48\pi\epsilon^2} \quad (\text{per minimal scalar}). \quad (30)$$

The match. Comparing (29) and (30),

$$\boxed{\frac{dS_{\text{ent}}}{dA} = \frac{1}{4G_{\text{ind}}}} \quad \text{identically.} \quad (31)$$

The specific number $1/48\pi\epsilon^2$ is scheme-dependent — it shifts with the regulator and with ξ — but the *equality* (31) is not, because both sides are the *same* heat-kernel coefficient a_1 of the *same* modes (Callan and Wilczek, 1994; Demers et al., 1995; Susskind and Uglum, 1994). This is the Susskind–Uglum identity, and in GG-Theory it is forced, not assumed: by GGEP there is no bare G_N to which the entanglement divergence could be a mere correction, so (31) reads as an equality of the entropy itself with $A/4G_N$.

31.5 Step 3 — the conical replica and the exact 1/4

The value of the ratio in (31) is fixed in closed form, independent of ϵ and of the field content, by the replica (conical-deficit) construction.


The arena, licensed. One word on where the computation takes place, because the replica method’s Euclidean arena is the step referees most often challenge as a formal device. In the cohort it is not a device. Time carries four endowments (Arisaka, 2026, P5-CPT): order from the causal cone, repetition from the equivalence of instants, the arrow, and the *analytic shore* — the half-plane of lawful imaginary time that the clock’s one-sided spectrum opens (Theorem 5(ii)). The Euclidean section on which the replica cover is built *is* that shore: for the quasi-free exterior algebra the continuation to imaginary Killing time exists as a theorem, and the τ -circle of period $2\pi/\kappa$ whose pinching produces the conical deficit below is the shore’s thermal face (Theorem 5(iii)), not a regularity trick. One import remains, and we declare it: evaluating the *gravitational* action semiclassically on the cover — $\ln Z_n = -I_E[\mathcal{M}_n]$ below — is the saddle step of the replica literature (Solodukhin, 2011), carried at that grade. The named-statement deck registers this license as the text key L-BH-Shore. The arena’s status is therefore Derived-conditional — derived through the clock’s shore, conditional on the quasi-free class and the declared saddle import — and the coefficient extraction that follows is exact arithmetic on both.

Prepare the Euclidean n -fold cover $\mathcal{M}_n$ with a conical deficit angle $2\pi(1-n)$ on the bifurcation surface Σ (the horizon, of area A). The curvature acquires a distributional piece concentrated on Σ (Solodukhin, 2011):

$$\int_{\mathcal{M}_n} \sqrt{g} R = n \int_{\text{smooth}} \sqrt{g} R + 4\pi(1-n) A. \quad (32)$$

Figure 9 displays the two normalizations whose ratio is the theorem.

Euclidean time τ (replica cover $\mathcal{M}_n$)



Σ : horizon, area A

$$\int_{\mathcal{M}_n} \sqrt{g} R \supset 4\pi(1-n) A$$

$$\Rightarrow S_{\text{BH}} = \frac{A}{4G_N}, \quad \frac{1}{4} = \frac{4\pi}{16\pi}$$

Figure 9: The conical-deficit origin of the exact 1/4. On the replica cover $\mathcal{M}_n$ the Euclidean-time circle is pinched at the horizon Σ with deficit angle $2\pi(1-n)$. The curvature concentrated on Σ contributes $4\pi(1-n)A$ to $\int \sqrt{g} R$; fed through the induced Einstein–Hilbert action and the replica formula, this gives $S_{\text{BH}} = A/4G_N$, the coefficient $1/4 = 4\pi/16\pi$ being the ratio of the conical-curvature normalization to the Einstein–Hilbert normalization (Theorem 1).

Insert this into the induced gravitational action $I_E[\mathcal{M}_n] = -\frac{1}{16\pi G_N} \int_{\mathcal{M}_n} \sqrt{g} R$:

$$I_E[\mathcal{M}_n] = n I_E[\mathcal{M}_1] - \frac{(1-n)A}{4G_N}, \quad \ln Z_n = -I_E[\mathcal{M}_n]. \quad (33)$$

The replica entropy is $S = -\text{Tr } \rho \ln \rho = (\ln Z_1 - \partial_n \ln Z_n)|_{n=1}$, since $\text{Tr } \rho^n = Z_n/Z_1^n$. Writing $R_1 \equiv \int_{\mathcal{M}_1} \sqrt{g} R$,

$$\ln Z_n = \frac{n R_1}{16\pi G_N} + \frac{(1-n)A}{4G_N}, \quad \partial_n \ln Z_n|_{n=1} = \frac{R_1}{16\pi G_N} - \frac{A}{4G_N},$$

$$\ln Z_1 = \frac{R_1}{16\pi G_N}, \quad S = \ln Z_1 - \partial_n \ln Z_n|_{n=1} = \boxed{\frac{A}{4G_N}}. \quad (34)$$

The bulk terms R_1 cancel identically; the entropy is purely the boundary (area) term. The coefficient is transparent:

$$\frac{1}{4} = \frac{4\pi}{16\pi} = \frac{(\text{conical-deficit curvature coefficient})}{(\text{Einstein–Hilbert normalization})}, \quad (35)$$

the ratio of the $4\pi(1-n)A$ in (32) to the $16\pi G_N$ in the action. Neither number knows anything about black holes: the 4π is fixed by the geometry of spheres (Gauss–Bonnet), the 16π by demanding Newton’s law in the weak-field limit — both were set by consistency requirements decades before anyone computed a horizon entropy. No free parameter enters, no regularization choice survives, and — crucially — supersymmetry is never used. Restoring k_B and $\hbar$,

$$\boxed{S_{\text{BH}} = \frac{k_B c^3 A}{4 G_N \hbar}} \quad (36)$$

which is Eq. (20), now *derived* rather than posited.

Theorem 1 (Exact Bekenstein–Hawking coefficient from induced gravity (T-BH-1)). *Let the four-dimensional Newton constant be induced (GGEP), the bulk metric be canonical (GG-A2-GG6, Theorem 2), and observation be the OGGE trace. Then the horizon entropy is $S_{\text{BH}} = A/4G_N$, with the coefficient $1/4 = 4\pi/16\pi$ fixed regulator-independently and without supersymmetry or extremality, for the generic (Schwarzschild, Kerr, Reissner–Nordström) black hole.*

Grade. Theorem 1 is held at *Derived-conditional*: derived regulator-independently by Steps 1–3, conditional on exactly its three stated premises — full inducedness (GGEP), the canonical bulk metric (GG-A2-GG6, Theorem 2), and observation as the OGGE trace — with the Euclidean arena licensed at the same grade (§31.5) and the numerical species closure registered as the module’s internal kill row (F-BH-9, Table 9; Appendix R). Nothing in the coefficient $\frac{1}{4}$ waits on that closure.

31.6 Step 4 — the microstates are the hidden fiber

The replica computation tells us *that* $S_{\text{BH}} = A/4G_N$; Step 1 and Step 2 tell us *which* modes are being counted — those whose one-loop determinant induces G_N . By OGGE (AX-3) these are exactly the bulk-fiber modes that a horizon traces out: the $U(1)_{Q_i}$ Wilson-line holonomies $\Theta_H = \oint A_\theta d\theta$ on the phase circle S_θ^1 — protection by holonomy rather than by symmetry, in GG-W2-Holonomy’s sense (Arisaka, 2026, W2-Holonomy) —, together with the Kaluza–Klein tower of the scale direction I_r , evaluated on the horizon two-sphere. In the hidden-fiber language of GG-A8-QM (Arisaka, 2026, A8-QM), the Bekenstein–Hawking entropy is

the entropy of the non-spatial fiber, traced out over the horizon.

The same compact phase circle S_θ^1 that carries quantum coherence in the bulk (where its holonomy is the Aharonov–Bohm/Berry phase) becomes, at a horizon, the carrier of black-hole microstates; the black-hole counting and the origin of quantum mechanics are one mechanism evaluated on two different loci. This identifies the microstates concretely and finitely — a horizon of area A supports $\exp(A/4\ell_P^2)$ distinguishable fiber configurations — without any appeal to branes, extremal limits, or replica wormholes.

31.7 Generality: neither supersymmetry nor extremality is used

It is worth stating plainly what the derivation did *not* assume. It used only (i) that G_N is induced (GGEP), (ii) the canonical scale profile (Theorem 2), and (iii) that observation is the OGGE trace. It never used supersymmetry, a BPS bound, extremality, or a near-horizon AdS_2 throat. Consequently (34) holds for the *generic* black hole — Schwarzschild, Kerr, Reissner–Nordström alike — the astrophysical case. This is the decisive point of contrast with the string-theoretic microstate count (Strominger and Vafa, 1996), which is exact but available only for supersymmetric, extremal or near-extremal black holes. GG-Theory reproduces the exact coefficient precisely where that method does not reach.

31.8 Kerr, Reissner–Nordström, and the universality of the 1/4

Theorem 1 was stated for the generic black hole, but the conical derivation (§31.5) was carried out on the Schwarzschild horizon. We now show the result is unchanged for rotating and charged black holes — the case that matters astrophysically, and the case the explicit string count does not reach.

The coefficient is metric-independent. The replica derivation used only two facts: that the horizon is a stationary Killing bifurcation surface Σ of area A , around which the Euclidean-time circle caps off smoothly with period $\beta = 2\pi/\kappa$; and the induced Einstein–Hilbert action. The conical-deficit identity $\int_{\mathcal{M}_n} \sqrt{g} R \supset 4\pi(1-n)A$ holds for *any* such Σ . Neither input distinguishes Schwarzschild from Kerr or Reissner–Nordström, so

$$\boxed{S = \frac{A}{4G_N} \quad \text{for every stationary black hole,}} \quad (37)$$

with the coefficient $1/4 = 4\pi/16\pi$ identical in every case and only the horizon area $A(M, J, Q)$ changing. The microstates are the same hidden-fiber holonomies traced over Σ (§31.6).

The three families. In geometrized units ($G_N = c = 1$, restored at the end), with $a = J/M$,

$$\begin{aligned} \text{Reissner–Nordström: } r_+ &= M + \sqrt{M^2 - Q^2}, & A &= 4\pi r_+^2, \\ \text{Kerr: } r_+ &= M + \sqrt{M^2 - a^2}, & A &= 4\pi(r_+^2 + a^2), \\ \text{Kerr–Newman: } r_+ &= M + \sqrt{M^2 - a^2 - Q^2}, & A &= 4\pi(r_+^2 + a^2), \end{aligned} \quad (38)$$

and in each case $S = A/4G_N$ (Kerr, 1963; Newman et al., 1965; Reissner and Nordström, 1916).

First-law consistency. With $T = \hbar\kappa/2\pi ck_B$ and $S = A/4G_N$, the GG-Theory entropy automatically satisfies the full first law of the (M, J, Q) family (Bardeen et al., 1973),

$$dM = T dS + \Omega_H dJ + \Phi_H dQ, \quad \Omega_H = \frac{a}{r_+^2 + a^2}, \quad \Phi_H = \frac{Qr_+}{r_+^2 + a^2}, \quad (39)$$

because $A/4G_N$ is precisely the Bekenstein–Hawking–Wald entropy the first law requires (Wald, 1993). GG-Theory thus reproduces the entire stationary first law, not only the static coefficient.

Table 3: The exact 1/4 across the stationary black-hole family. In every case $S = A/4G_N$ with the identical coefficient $1/4 = 4\pi/16\pi$; supersymmetry and extremality are never used. The point of running the whole family across the table is that the coefficient does not move. Charge and rotation change the horizon radius, the area formula and the work term, and leave the 1/4 exactly where it was. No row uses supersymmetry, none is extremal, and none was chosen because it worked — which is the difference between this and a microstate count in a controlled sector.

Black hole	Horizon r_+ ($G_N = c = 1$)	Area A	First-law work
Schwarzschild	$2M$	$16\pi M^2$	—
Reissner–Nordström	$M + \sqrt{M^2 - Q^2}$	$4\pi r_+^2$	$\Phi_H dQ$
Kerr	$M + \sqrt{M^2 - a^2}$	$4\pi(r_+^2 + a^2)$	$\Omega_H dJ$
Kerr–Newman	$M + \sqrt{M^2 - a^2 - Q^2}$	$4\pi(r_+^2 + a^2)$	$\Omega_H dJ + \Phi_H dQ$

The astrophysical point. Every observed black hole rotates; the astrophysical objects are near-Kerr and non-extremal. The induced-gravity derivation covers exactly this generic case, with no supersymmetry and no extremal limit — the regime in which the explicit string-theoretic microstate count is not available. This is the concrete content of the generality claimed in §31.7.

31.9 The subleading logarithmic correction

Beyond the area law, every microscopic theory of black-hole entropy predicts a *universal* subleading term,

$$S = \frac{A}{4G_N} + c_{\log} \ln \frac{A}{\ell_P^2} + \cdots, \quad (40)$$

and the coefficient $c_{\log}$ is a sharp, parameter-free discriminator between frameworks. In the induced-gravity derivation it has a clean origin: where the leading 1/4 came from the a_1 Seeley–DeWitt coefficient (the R term, §31.4), the logarithm comes from the *next* coefficient a_2 (the dimension-four, conformal-anomaly term). On the replica cover the a_2 density integrated over the conical singularity carries an n -dependence whose ∂_n at $n = 1$ is logarithmic in the cutoff, hence in A/ℓ_P^2 ; the coefficient is fixed by the conformal-anomaly charges of the massless fields crossing the horizon (Carlip, 2000; Sen, 2013; Solodukhin, 2011).

The GG-Theory prediction. Because $c_{\log}$ is set by the same boundary spectrum that fixes the induced G_N , GG-Theory predicts a definite value, $c_{\log} = -\frac{1}{2} b_2[\text{GG-Theory massless content}]$, with b_2 the summed a_2 -anomaly coefficient of the Standard-Model fields and the light bulk modes. For its *exact number* this shares the open status of the leading species sum (Appendix R); its *structure* is robust, and it is an independent falsifiable check. Two consequences follow. First, internal consistency: the GG-Theory value must coincide with the macroscopic a_2/Sen value for its spectrum. Second, a discriminator: it generically *differs* from the loop-quantum-gravity prediction $c_{\log} = -\frac{3}{2}$, which arises from a horizon Chern–Simons state count rather than from a bulk anomaly. A microscopic count that reproduced $A/4G_N$ but missed the GG-Theory $c_{\log}$ would falsify the induced-gravity reading.

31.10 What is rigorous here, and what remains

In the spirit of MF (POS-7) we separate the secure from the outstanding.

Established. The area law and the exact coefficient $\frac{1}{4}$ follow from (32)–(35), regulator-independently, for generic non-supersymmetric black holes, given only that G_N is induced. The microstates are identified (Step 4). The Euclidean arena of Step 3 is itself licensed constitutionally — the analytic shore of the four-endowment time, consumed through Theorem 5 — at Derived-conditional grade, with the semiclassical gravitational saddle the one declared import (§31.5). This is a genuine derivation of the coefficient that made the Strominger–Vafa result a landmark, in a setting that result cannot reach.

Outstanding. Three items remain, and we flag them honestly. (a) **Species closure:** $S_{\text{BH}} = A/4G_{\text{ind}}$ holds for *whatever* G_{ind} the modes produce; obtaining a pure number requires summing the actual GG-Theory spectrum (the Standard Model content of GG-A9-SM plus the bulk Kaluza–Klein tower) and checking that $1/G_{\text{ind}}$ equals the G_N of Step 1 that appears in the four-dimensional Einstein equations. We show in Appendix R that this closure is *structural* — the Kaluza–Klein tower sum reproduces the scale reduction — while the precise leading numerical coefficient remains the principal open computation. (b) **Non-minimal and gauge fields:** the matching (31) is clean for minimally coupled scalars but carries known subtleties (contact terms) for gauge fields and the graviton (Kabat, 1995); the full spectrum requires care. (c) **Subleading** $\log A$: the a_2 heat-kernel coefficient yields a universal logarithmic correction $S_{\text{BH}} \supset c_{\log} \ln(A/\ell_P^2)$, an independent check the framework must pass (developed in §31.9). None of these touches the coefficient $\frac{1}{4}$ itself, which is settled by Step 3.

The Part in one breath, and the door to the next. Gravity’s stiffness and the horizon’s memory are the same purchase, priced once. Because no bare Newton constant exists, the entanglement across the horizon does not approximate $A/4G_N$ — it *is* $A/4G_N$; and the coefficient is the ratio of two normalizations fixed by non-negotiable physics decades before anyone computed an entropy. The physics intuition to carry forward: *the $1/4$ is not a count that could drift; it is a ratio that cannot.* What the coefficient does not yet tell us is how the account settles over a black hole’s lifetime — when the information comes back, why no firewall guards the door, and what the second law charges for every exchange. That ledger, set in motion, is the business of the next Part.

32 Part V in plain English: why the coefficient is a ratio rather than a count, and why that is checkable

This is the paper’s central technical result, and the shape of the argument is unusual enough to be worth stating before the argument itself.

Everyone agrees the horizon entropy is proportional to area. The coefficient is one quarter, exactly, in units of Newton’s constant. The usual way to explain it is to count something — microstates, in some sector where counting can be controlled — and check that the count comes out right. That works where it works, and it does not reach the black hole at the center of a galaxy.

This Part does not count. It observes that in this framework there is no fundamental Newton

constant at all: gravity is induced, which means the constant is generated by the same fields whose entanglement across the horizon one would be trying to count. The two quantities are therefore not independent, and both fall out of a single heat-kernel calculation. The entanglement entropy does not approximate a quarter of the area over Newton's constant. It is that, because the quarter is a ratio between two normalizations that other physics had fixed decades before anyone asked about entropy.

The consequences are what make the claim attackable. No supersymmetry and no extremality enter, so the result covers Kerr, Reissner–Nordström and real astrophysical holes rather than a protected sector. And a ratio cannot drift the way a count can — which is exactly why a measurement or a calculation returning a coefficient other than one quarter would end this reading rather than adjust it.

What is not settled is in the Part's own last section rather than left for a reader to notice. The species sum that closes the argument is a computation nobody has performed, and it is registered as a falsifier rather than filed as a detail.

Part VI

Information: Page Curve, Firewalls, and the Interior

The Ledger, the Scheduled Expiry of Freshness, and the Direction of the Second Law

Part V fixed the size of the horizon’s account; this Part watches the account settle in time. Everything here is bookkeeping on one object — the free-energy ledger that the open-system constitution supplies before any black hole is mentioned — and the Part’s results are the entries that ledger forces: the Page curve as a branch switch whose clock is an audited counting premise; the firewall dissolved as an artifact of describing one system twice; the interior read as bulk rather than as an ending; the four laws upgraded from analogy to GDOS statements; and the generalized second law given both its direction and its price. The reader who wants the mechanism in one sentence can have it now: information is never lost, only displaced — and its return is not a mechanism switching on but an assumption expiring on schedule. The sections that follow earn that sentence, starting with the ledger itself.

33 The GG-Theory free-energy ledger and generalized entropy

The Bekenstein–Hawking entropy, derived in Part V as an entanglement entropy, is a special case of a structural object we now record in full, because that is what makes the black-hole result an instance of a universal law rather than a gravitational accident. In the modern island and quantum-extremal-surface (QES) literature, the entropy of Hawking radiation is obtained from the generalized-entropy functional

$$S_{\text{gen}}[X] = \frac{A[X]}{4G_N} + S_{\text{bulk}}[X]. \quad (41)$$

Within GG-Theory this is immediately recognized as the black-hole instance of the universal scalar ledger of GDOS ([Arisaka, 2026](#), A5-GDOS),

$$\text{GFF} = F_{\text{GEP}} + F_{\text{CEP}}, \quad (42)$$

the sum of a geometric extremal piece (GEP) and a curvature–entropy piece (CEP). For a black hole the natural identification is

$$F_{\text{GEP}}^{\text{BH}} = \frac{A}{4G_N}, \quad F_{\text{CEP}}^{\text{BH}} = S_{\text{bulk}}, \quad F_{\text{GFF}}^{\text{BH}} = \frac{A}{4G_N} + S_{\text{bulk}}. \quad (43)$$

The significance is interpretive, not notational. In the standard horizon story, generalized entropy looks like the right formula discovered *within* black-hole semiclassicals. In GG-Theory it is the black-hole reading of a ledger that the open-system constitution *already required*, before any black hole is

introduced.

33.1 Why the area term belongs to GEP

In ordinary thermodynamics the free-energy ledger carries a cost for occupying macroscopically constrained configurations. In gravity, horizon area is exactly such a constraint: for a fixed exterior observer, changing the area changes not only mass and temperature but the causal partition between accessible and inaccessible sectors. The area term therefore measures more than surface size — it is the geometric price of the horizon itself, and so belongs to the geometric piece GEP of the ledger.

33.2 Why the bulk-entanglement term belongs to CEP

The second term, S_{bulk} , measures the unresolved internal state volume compatible with the chosen branch; it is where entanglement, coarse-graining, and state counting enter. The exterior observer does not see microscopic bulk data directly, only a projected channel whose entropy must include the volume of discarded interior possibilities — precisely the role of the curvature–entropy piece CEP. The conceptual payoff is that GG-Theory removes the air of special pleading from generalized entropy: its two terms are not unrelated objects glued together because black holes happened to demand them, but the two natural halves of one open-system ledger.

33.3 Operational entropy as a branch minimum

Once the ledger is written this way, the physically accessible radiation entropy is not the naive thermal entropy of the emitted quanta. It is the branch-minimized ledger over admissible coarse-grained saddles,

$$S_R(t) = \min_{X \in \mathfrak{B}_{\text{adm}}} F_{\text{GFF}}^{\text{BH}}[X; t]. \quad (44)$$

This is deliberately close to the island prescription, but its reading is different: in island language one extremizes and minimizes over candidate quantum extremal surfaces; in GG-Theory the operational boundary system explores admissible coarse-grained branches and the realized entropy is the minimum of the universal ledger over them. The black-hole problem is then only one especially vivid instance of a rule already needed for open systems, non-equilibrium thermodynamics, and measurement.

33.4 The factorization problem across the horizon

A hidden assumption in many textbook formulations of the information paradox is that the total Hilbert space factorizes cleanly, $\mathcal{H}_{\text{tot}} \stackrel{?}{=} \mathcal{H}_{\text{in}} \otimes \mathcal{H}_{\text{out}}$. In a non-gravitational theory this is often harmless; in gravity it is dangerous. Physical states obey constraints $\mathcal{C}_\alpha |\Psi_{\text{phys}}\rangle = 0$, and states related by gauge redundancy are not distinct, so the physical Hilbert space is a constrained quotient $\mathcal{H}_{\text{phys}} = \ker \mathcal{C}_\alpha / (\text{gauge redundancy})$ rather than an unconstrained tensor product.

In GG-Theory this is the DGI content of the factorization puzzle: the horizon split cannot be treated as a cut through independent microscopic degrees of freedom, because diffeomorphism constraints and edge data already make the subsystem notion subtler than the naive product. The




Figure 10: Information “loss” as geometric displacement. The six-dimensional bulk state is pure and its information is never destroyed; the OGGEP projection π traces the non-spatial fiber, so a mode displaced along (I_r, S_θ^1) becomes inaccessible to the four-dimensional observer and the boundary state looks mixed. What the 4D observer reads as lost information is, in the bulk, merely *displaced* — distance along a non-spatial direction, not destruction.

island story is indirect evidence for the same point — when the radiation’s entanglement wedge acquires an island, one is already conceding that its operational subsystem is not the naive geometric exterior.

Islands are then not a primitive ontology but the late-time support of the minimizing operational branch once DGI and OGGEP projection are respected. Finally, the data-processing inequality makes the boundary description operationally honest: for the CPTP projection Π from bulk to boundary, relative entropy cannot increase, so the boundary may *lose* information without the parent theory *destroying* it. Hawking’s original argument acquires its paradoxical force only by treating the semiclassical exterior channel as the full ontology — the move that DGI and OGGEP together reject.

34 The explicit OGGEP channel for a Schwarzschild black hole

Figure 10 plots the distances just computed.

34.1 Setup

Consider a Schwarzschild black hole of mass M in flat 4D spacetime. The bulk geometry is approximately

$$ds_6^2 = e^{-2r} \left[-f(r_4)dt^2 + dr_4^2/f(r_4) + r_4^2 d\Omega^2 \right] + dr^2 + d\theta^2, \quad (45)$$

where r_4 is the 4D radial coordinate with horizon at $r_4 = r_s = 2G_N M/c^2$ and $f(r_4) = 1 - r_s/r_4$. The two non-spatial coordinates are the scale direction r and the phase angle θ ; the scale profile e^{-2r} and the unit circle period are fixed by GG-A2-GG6 (Theorem 2), with no free scale rate or circle radius.

34.2 The OGGEP channel

The OGGEP channel takes the 6D bulk state and produces the 4D exterior state by partial trace:

$$\mathcal{D}_{\text{BH}} : \rho_{\text{bulk}} \mapsto \rho_\partial^{\text{ext}} = \text{Tr}_{\text{in}}(\rho_{\text{bulk}}), \quad (46)$$

where Tr_{in} traces over (i) the compact dimensions (r, θ) and (ii) the 4D interior $r_4 < r_s$.

34.3 Kraus representation

Any CPTP channel admits a Kraus representation $\mathcal{D}(\rho) = \sum_j K_j \rho K_j^\dagger$ with $\sum_j K_j^\dagger K_j = \mathbf{1}$. For $\mathcal{D}_{\text{BH}}$, the Kraus operators are labeled by (a) the Hawking-radiation outgoing mode frequency ω , (b) angular-momentum partial wave ℓ, m , and (c) occupation number n :

$$K_{\omega\ell mn}^T = \sqrt{p_n(\omega) \Gamma_\ell(\omega)} |\omega\ell m; n\rangle \langle \phi_{\omega\ell m}^{\text{in}}|, \quad K_{\omega\ell mn}^R = \sqrt{p_n(\omega) (1 - \Gamma_\ell(\omega))} |\chi_{\omega\ell mn}\rangle \langle \phi_{\omega\ell m}^{\text{in}}|, \quad (47)$$

where $p_n(\omega) = (1 - e^{-\beta\hbar\omega})e^{-n\beta\hbar\omega}$ is the thermal probability at temperature $T_{\text{H}} = 1/k_B\beta$, $\Gamma_\ell(\omega)$ is the graybody factor acting as a beamsplitter of the emitted flux, and the $|\chi\rangle$ are the reflected, horizon-returning states; the input bra is the single infalling state of the mode, so completeness closes exactly: $\sum_n (K^{T\dagger} K^T + K^{R\dagger} K^R) = \sum_n p_n |\phi_{\omega\ell m}^{\text{in}}\rangle \langle \phi_{\omega\ell m}^{\text{in}}| = \mathbf{1}$ on the per-mode input span.

The observable output for an initial pure infalling state traces the reflected branch; per mode it is

$$\rho_{\partial}^{\text{out}}|_{\omega\ell m} = \sum_n \tilde{p}_n(\omega) |\omega\ell m; n\rangle \langle \omega\ell m; n|, \quad (48)$$

with $\tilde{p}_n(\omega)$ the thermal distribution at the *filtered* occupation $\bar{n}_{\text{obs}} = \Gamma_\ell(\omega) \bar{n}(\omega)$: at the per-quantum (fine-grained) level the barrier is a Gaussian beamsplitter, and a thermal state through a beamsplitter is exactly thermal at the transmitted occupation — the same Gaussian fact the exchange ledger of Appendix O runs on. The two-branch set above is the mode-level coarse-graining of that channel, matching it in mean flux — the graybody-filtered Planck flux of Eq. (19). One assumption in this construction now carries its cohort name: labeling the Kraus operators mode by mode, and composing them over the emission history, presumes that each engaged mode pair presents fresh — uncorrelated with the boundary record — at engagement. That is the Port-Freshness Clause of the open-system law (GG-A5-GDOS, Definition 5; Arisaka, 2026, A5-GDOS), consumed here with a validity window rather than silently: the window, and its scheduled edge at the Page time, are the content of Lemma 2.1.

34.4 The entropy of the output state

The von Neumann entropy of $\rho_{\partial}^{\text{out}}$ is

$$S_{\text{vN}}(\rho_{\partial}^{\text{out}}) = - \sum_{\omega\ell mn} \tilde{p}_n(\omega) \ln \tilde{p}_n(\omega) \approx S_{\text{BH}} = \frac{A}{4G_N \hbar} k_B, \quad (49)$$

at leading order, with the approximation holding when the thermal radiation dominates the phase space. The coefficient of $1/4$ is reproduced, and the area-scaling is automatic.

34.5 Compositionality: post-evaporation state

After complete evaporation, the black hole is gone and the radiation carries the full bulk state projected to 4D. In GDOS, this is:

$$\rho_{\partial R}^{\text{final}} = \mathcal{D}_{\text{total}}(\rho_{\text{initial}}), \quad (50)$$

where $\mathcal{D}_{\text{total}}$ is the composition of all the Kraus operations over the evaporation history. Crucially, for a bulk-pure initial state, the final 4D state is a pure state in the radiation — provided the radiation sector includes *all* degrees of freedom, including very-late subtle correlations. This is the manifest unitarity of the total evolution. Those very-late subtle correlations are now named objects: they are the stale-port data of Lemma 2.1 — the correlations that *must* appear once the freshness budget of the pure total state is spent.

35 The Page curve in GDOS

35.1 Setup

Consider a black hole formed from a pure quantum state $|\psi_0\rangle$. At time t , the system consists of two registers: the radiation $R(t)$ — Hawking quanta that have already been emitted and propagated outward — and the black hole $B(t)$, the remainder, with reduced mass $M(t) < M(0)$. Bulk unitarity says that the total state $|\psi(t)\rangle_{R \otimes B}$ is still pure, evolving via the unitary 6D evolution.

The Page curve is the time evolution of the entanglement entropy $S_R(t) = S_{\text{vN}}(\rho_R(t))$ of the radiation.

35.2 GDOS derivation

In GDOS, the radiation is produced by the OGGEF channel at the horizon. The bulk state initially is approximately:

$$|\Psi_0\rangle_{\text{bulk}} = |\psi_0\rangle_{M_4} \otimes |\Omega\rangle_{\text{compact}}, \quad (51)$$

where $|\psi_0\rangle_{M_4}$ is the 4D matter state that will collapse, and $|\Omega\rangle$ is the bulk-compact ground state.

As evaporation proceeds, bulk unitary evolution entangles $|\psi_0\rangle_{M_4}$ progressively with the compact bulk degrees of freedom, and the OGGEF channel projects the outside- horizon subsector.

Early times ($t \ll t_P$). Most of the bulk state is still “inside” the black hole (interior + compact). The radiation $R(t)$ has only weak correlations with the black hole $B(t)$. Its entanglement entropy grows as

$$S_R(t) \approx S_{\text{BH}}^{\text{emitted}}(t) \quad (\text{linear in radiated area}), \quad (52)$$

consistent with the Hawking result.

Page time ($t = t_P$). Half the initial area has been radiated. The bulk bipartition has equal-dimensional sides. The entanglement entropy is maximal.

Late times ($t > t_P$). The radiation is now the larger subsystem. Its entanglement entropy equals

that of the black hole:

$$S_R(t) = S_{\text{BH}}^{\text{remaining}}(t) \quad (\text{decreasing}), \quad (53)$$

by the bipartite pure-state property $S_{\text{vN}}(\rho_R) = S_{\text{vN}}(\rho_B)$. The radiation entropy decreases as the black hole shrinks.

End of evaporation. Final state is pure; radiation entropy is zero.

This is the Page curve’s shape, read directly off the GDOS structure without any exotic mechanism. The key point: once bulk evolution is unitary and the boundary sees a CPTP partial trace, the radiation entropy must rise and then fall. What the argument does not supply is the curve’s quantitative form: the turnover time is fixed below by counting under an imported typicality premise (Lemma 2.1), and the replica or island computation that would produce the curve itself is not performed here. No replica wormholes, no quantum extremal surfaces (as independent principles), no islands formula as a separate postulate are needed for the shape.

35.3 Connection to the islands formula

The islands formula (Eq. 9) gives the entropy of the radiation as

$$S(R) = \min \text{ext} \left\{ \frac{A(\partial I)}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I) \right\}. \quad (54)$$

At sub-Page times, no island is picked (the trivial $I = \emptyset$ minimizes), and $S(R) = S_{\text{semi-cl}}(R)$ (Hawking’s result). At super-Page times, a non-trivial island $I \subset$ black hole interior picks up, and $S(R)$ starts to decrease.

In GDOS, the island I is exactly the *bulk region that the OGGEF functor traces over*. At sub-Page times, this region is mostly the interior; its entanglement contribution is already captured by the black-hole entropy term. At super-Page times, the bulk-traced region effectively extends into the black-hole interior “from the radiation side,” giving the island its bulk extent.

The islands formula is therefore a formal consequence of the GDOS structure. The QES (quantum extremal surface) is the boundary ∂I of the traced-over region, and the extremization prescription is the saddle point of the OGGEF path integral.

GDOS and the islands formula agree on the Page curve. The advantage of GDOS is that the derivation is *physically concrete*: the 6D bulk provides a realization, the OGGEF partial trace is a specific operation, and the resulting Page curve is geometric rather than arising from replica wormholes.

36 Hawking–Page and the Page curve as branch-switch events

36.1 Hawking–Page as a gravitational branch selection

The Hawking–Page transition is usually presented as a result of Euclidean gravity, but its importance here is methodological: it is one of the cleanest cases in standard gravity where the realized history is selected by a competition among scalar thermodynamic branches rather than by local equations of motion alone. Writing the coarse-grained ledger of the two competing saddles




Figure 11: The radiation entanglement entropy. The naive thermal calculation (red dashed) rises monotonically to $S_{\text{BH}}(M_0)$, implying information loss. Unitarity requires the Page curve (blue): it follows the thermal branch until the *Page time* t_{Page} , when the area-limited branch becomes dominant in the GFF ledger, and then falls to zero as the black hole evaporates. In GG-Theory the turnover is the branch switch of §36, forced in shape by bulk unitarity plus the OGGE partial trace and dated by the counting of Lemma 2.1; no islands postulate is added, and the quantitative curve is not computed here.

as $\mathcal{F}_{\text{th-AdS}}(T)$ and $\mathcal{F}_{\text{BH}}(T)$, the equilibrium phase is the one of lower free energy; in GG-Theory language the runtime admits multiple thermodynamically admissible branches and the realized branch minimizes the scalar ledger. The Hawking–Page temperature is thus a gravitational instance of branch selection, showing that the GG-Theory branch grammar is not imported from chemistry or from measurement — it is already latent in standard black-hole thermodynamics (the explicit free-energy comparison is recorded in the Hawking–Page appendix). In the two-regime vocabulary of the runtime (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A11-Origin) this transition now has its canonical name: it is a **GRIP-Out** — temperature is a control parameter of the ledger, and as it crosses T_{HP} the free-energy landscape *itself* reorganizes, the dominant saddle structure changing identity. It is, in fact, the one genuine GRIP-Out of black-hole physics (§42.3); everything else in the black hole’s life — collapse, merger, evaporation — happens on a fixed landscape.

36.2 The Page curve as a GFF branch minimum

Consider two competing operational branches for the radiation entropy. The early-time Hawking branch grows as the black hole loses entropy,

$$S_R^{(1)}(t) \approx \eta[S_{\text{BH}}(M_0) - S_{\text{BH}}(M(t))], \quad (55)$$

while the late-time bridge/island-like branch is capped by the remaining geometric ledger,

$$S_R^{(2)}(t) \approx S_{\text{BH}}(M(t)). \quad (56)$$

The operational entropy is the branch minimum $S_R(t) = \min\{S_R^{(1)}, S_R^{(2)}\}$ of (44) (Figure 11). For $\eta \simeq 1$ the crossing occurs when $S_{\text{BH}}(M_{\text{Page}}) = \frac{1}{2}S_{\text{BH}}(M_0)$, i.e.

$$M_{\text{Page}} = \frac{M_0}{\sqrt{2}}, \quad t_{\text{Page}} \approx (1 - 2^{-3/2}) t_{\text{evap}}. \quad (57)$$

The Page turnover is therefore not a miraculous late correction tacked onto Hawking’s early-time story; it is the moment when a different operational branch becomes dominant in the entropy ledger; equivalently, in the boundary-side reading of (Arisaka, 2026, A5-GDOS), the Page curve is a theorem of OGGEp partial-tracing plus bulk unitarity (the bipartite-pure-state argument), of which the branch-switch is the free-energy-ledger view. In regime language the classification is sharper still, and it strengthens the no-firewall conclusion of §37: the Page time is a *milestone of GRIP-In*, not a transition — like matter–radiation equality in cosmology, it is a bookkeeping crossover on a landscape that does not move. Nothing physical reorganizes at t_{Page} ; there is therefore nothing at the Page time for a firewall to ride on. The taxonomy itself rules the drama out.

Theorem 2 (GG-Theory Page-turnover as a branch switch (T-BH-2)). *Within the minimal black-hole module of GG-Theory/GDOS, the Page turnover is the branch-switch time at which the GFF-minimizing operational saddle ceases to be the Hawking-like thermal branch and becomes an area-limited bridge branch. Page-curve recovery is in this language a GSSB event in the entropy ledger, not a late violation of early-time semiclassical validity.*

36.3 The mechanism: port freshness and its scheduled expiry

Theorem 2 *names* the turnover; this subsection supplies its mechanism, in the vocabulary the cohort maintains for exactly this purpose. The open-system law runs on one named working assumption, stated once for the whole program as the *Port-Freshness Clause* (L-GDOS-Fresh; GG-A5-GDOS, Definition 5): every engaged environment port presents a fresh, memory-free state to the operational dictionary — the boundary system and each engaged port are uncorrelated at engagement. That clause is GG-A5-GDOS’s single home of the factorizable initial state, the Markov property, and what GG-A8-QM calls the memoryless bath; GG-A5-GDOS grades it honestly as the correlational half of molecular chaos, the cohort’s declared past-hypothesis-equivalent (Arisaka, 2026, A5-GDOS). The channel constructions of this paper consume it — the mode-by-mode Kraus product and its composition over the emission history (§34) are exactly factorized-engagement statements — and until now they consumed it silently. At a black-hole horizon the clause has a sharply bounded validity window, and the window’s edge is not set by dynamics at the port: it is set by bookkeeping, and it *is* the Page time.

Lemma 2.1 (The stale port: freshness expires at the Page time (L-BH-Fresh)). *For a black hole formed from a pure state, evolving unitarily in the bulk, with boundary record $R(t)$ the emitted radiation, and granting the typicality premise of Appendix K (an imported dynamical assumption, flagged there): (i) (fresh regime) for $t \ll t_{\text{Page}}$ each newly engaged horizon mode is uncorrelated with $R(t)$ up to corrections exponentially small in the unread entropy, so the Port-Freshness Clause holds at the horizon and the mode-by-mode channel composition of §34 is exact at that grade; this is*

the regime in which the GFF-minimizing branch is the Hawking-like thermal branch of Eq. (55); (ii) (finite budget) the clause’s correlational content has a finite budget at a horizon: a pure total state can present fresh port modes only while the unread sector holds more than half the total information, so freshness expires at t_{Page} by counting alone — no new interaction, no new scale, no change in the port dynamics; (iii) (stale regime) for $t > t_{\text{Page}}$ every newly engaged mode is already correlated with the record — in the coarse-grained model the mutual information per emitted unit approaches its maximum, $2 \ln 2$ — and these forced correlations are precisely the late-time data that cap the radiation entropy at the area-limited bridge branch of Eq. (56). The branch switch of Theorem 2 is therefore the scheduled expiry of L-GDOS-Fresh at the horizon: the information return is not a mechanism switching on but an assumption switching off. Grade: Derived-conditional — derived from bulk unitarity and the bipartite pure-state property, conditional on the typicality premise, and consuming GG-A5-GDOS’s clause as the name of what expires.

Proof (coarse-grained model). Work in the N -unit model of Appendix K: after m units are emitted the record is R (dimension 2^m), the next engaged mode is one unit q , and the unread remainder has dimension 2^{N-m-1} . Freshness at engagement is the statement $I(q:R) = S_q + S_R - S_{Rq} \approx 0$. (i) For $m < N/2$ Page’s averages give $\langle S_X \rangle = \ln d_X - d_X/2d_{\bar{X}} + \dots$ for every subsystem X smaller than its complement, so the three leading logarithms cancel in $I(q:R)$ and the survivor is the correction term, of order $d_R d_q / d_{\bar{X}} \sim 4^{m-N/2}$: exponentially small in $N - 2m$. (ii) is the pure-state constraint $S_R = S_{\bar{R}}$: once $d_R > d_{\bar{R}}$ the record’s entropy is capped by the shrinking complement, $S_R \approx (N - m) \ln 2$, so each further emission *lowers* S_R by $\ln 2$ while adding a unit — arithmetic possible only if the added unit arrives correlated with the record. (iii) then follows by the same counting: $I(q:R) \rightarrow 2 \ln 2$, the maximum for one unit. Numerically ($N = 12$, Haar-random states, five realizations, verification suite at www.arisaka-gg.org): $I(q:R) = 0.002 \ln 2$ at $m = 1$, growing by the predicted factor $\simeq 4$ per step (4.05, 4.13, 4.12, 4.03 measured), crossing at the Page point $m = 6$, and saturating at $1.997 \ln 2$ of the maximum $2 \ln 2$ at $m = 10$; the record’s entropy traces the Page tent throughout. What is *not* claimed: the gravitating micro-dynamics by which the correlations are delivered (scrambling; Hayden and Preskill, 2007), any modification of the clause’s status or grade at its GG-A5-GDOS home, and any failure of the clause elsewhere in the cohort — the black hole is the one system whose own schedule drives the clause from holding exactly to failing maximally within a single history, which is what makes it the cohort’s sharpest stress test of L-GDOS-Fresh. ■

Physics intuition. A dealer dealing from a finite, well-shuffled deck: early hands are fresh — nothing on the table tells you what comes next. Past the midpoint of the deck, every further card is increasingly readable from the table, and the last card is fully determined by it. The dealer never changes the rules; the deck simply runs out of newness. A black hole formed from a pure state deals from a finite deck, the Page time is the midpoint, and the celebrated “information return” is the table beginning to read the deck. Firewall arguments, in this language, mistook the growing readability of the cards for a change in how the dealer deals; Theorem 2 and the lemma say the dealing — the local horizon physics — never changes at all.

36.4 Hayden–Preskill, and why no firewalls or remnants are forced

The Hayden–Preskill scenario fits naturally: once the Page turnover has occurred, and after a scrambling delay, newly injected information can in principle be reconstructed from the outgoing channel, because the boundary-accessible map is no longer trapped on the purely thermal branch. There is no contradiction with early-time thermality, because the branch structure is itself time-dependent. The key refinement is that GG-Theory distinguishes three layers — bulk purification, boundary accessibility, and detector-level reconstruction — and confusion among them inflates paradoxes in the literature. Once the ledger is branch-structured and the subsystem notion is operational rather than naive, several older cures lose their necessity: one need not postulate a generic firewall to preserve entropy bookkeeping, nor hide the state count in an implausibly degenerate remnant. Honesty requires the narrow form of the claim: GG-Theory does not prove firewalls impossible in every UV completion; it shows that entropy bookkeeping by itself does not force them once the global-versus-operational distinction is made correctly.

37 No firewalls in GDOS

37.1 The monogamy puzzle dissolves

Recall the AMPS argument: the late-time Hawking quantum B must be entangled with both the earlier radiation R (for unitarity) and its interior partner A (for no-drama at the horizon). Entanglement monogamy rules this out.

In GDOS, the resolution is:

There is no independent “interior partner A .” The apparent partner is part of the 6D bulk, which the OGGEF functor traces over. B is entangled with a *bulk state*, and the bulk state has been partial-traced to produce both the “interior” (as seen by an infalling observer) and the “exterior” (as seen by the outside observer).

Formally: AMPS assumes that the 4D Hilbert space factorizes as $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_R$ with $\mathcal{H}_A, \mathcal{H}_B, \mathcal{H}_R$ independent. In GDOS, the 4D description arises from tracing over the bulk, and the three subsystems are not really independent at the 6D level.

More concretely: the apparent partner A is part of a bulk configuration that is also entangled with R through the bulk dynamics. The entanglement is not monogamous because it is redistributed over the bulk, which has much more room than the 4D-Hilbert-space decomposition admits.

37.2 Why no firewall is forced on the infalling observer

No firewall is forced on an observer falling through the horizon, for three reasons that are really one. In the bulk 6D picture, the observer simply transitions from accessing scale-profile- $e^{A(r_*)}$ degrees of freedom (the exterior) to accessing scale-profile- $e^{A(r'_*)}$ degrees of freedom (the interior, at smaller scale profile). The horizon is a regular 6D surface in the canonical bulk, and nothing in the entropy bookkeeping makes high-energy quanta accumulate there. And the apparent high energies at the horizon — those of the Unruh frame of an accelerated exterior observer — are frame-dependent.




Figure 12: Schematic conformal (Penrose) diagram of gravitational collapse and evaporation. Time runs upward; each 45° line is null. The event horizon (blue) is the one-way OGGEF port: the exterior observer on $\mathcal{I}^+$ receives the Hawking radiation (orange) but cannot access the interior, which — together with the non-spatial fiber — is traced out. In GG-Theory the classical singularity is a coordinate artifact of the small-scale-profile region and is regular in the six-dimensional bulk, and the horizon carries no firewall because it is a regular 6D surface.

This is the standard general-relativity no-drama picture, left standing at the 6D level: the firewall argument needed an independent interior partner, and the six-dimensional description has none. It does not follow that every kind of horizon-scale drama is excluded; it follows that none is implied by the paradox, which is what the firewall argument claimed. The paradox was an artifact of insisting on a purely 4D description of a fundamentally 6D system.

38 The interior of the black hole in GDOS

Figure 12 shows the causal structure on which this Part reads the information flow.

38.1 What's inside?

In 4D, the interior of a Schwarzschild black hole ($r < r_s$) contains a curvature singularity at $r = 0$. In classical GR, this singularity is a region of infinite tidal forces where classical theory breaks down.

In GDOS, the 4D interior is part of the 6D bulk. The apparent 4D singularity corresponds to the region where the scale profile is very small, $e^{-r} \ll 1$. In 6D, this is not a singular region; it is just a region where the 4D metric is strongly scale-running, and the effective low-energy physics

there is different from the exterior.

What the infalling 4D observer sees is a journey through the projection, not a march toward a wall. Near the horizon the 4D metric looks Schwarzschild, and the observer still sees the exterior. Just inside, the metric becomes unusual: what was “spatial” becomes “temporal” in Schwarzschild coordinates. Deep inside, the scale running is extreme, and the 4D metric becomes effectively a null surface. And at $r = 0$ the 4D metric is singular — but the 6D metric is smooth: the “singularity” is an artifact of the 4D projection.

38.2 Information doesn’t “fall” into the singularity

In the classical picture, information that crosses the horizon is destined to end up at the singularity. In GDOS, this is not quite right: the 4D information that “crosses” the horizon simply enters the region of smaller scale profile, where it becomes inaccessible to the 4D exterior observer but continues to exist in the 6D bulk.

Over time, the bulk dynamics redistribute this information across the compact directions and the radial throat. Eventually, this redistribution leads to it being carried off by Hawking radiation — but the redistribution is slow, and for astrophysical black holes, very long timescales are involved (Page time $\sim (G_N M^2 / \hbar)$).

The singularity, insofar as it exists, is a coordinate singularity in the 4D projection, not a physical endpoint of information.

39 Black-hole complementarity revisited

39.1 Susskind-Thorlacius-Uglum 1993

The original black-hole complementarity proposal (Susskind et al., 1993) was: an exterior observer and an infalling observer both have valid, but mutually inconsistent, descriptions of the black hole. The exterior observer sees information preserved in the Hawking radiation; the infalling observer sees information entering the interior. These two descriptions are complementary in the Bohr sense: one cannot be tested in a single experiment against the other.

The complementarity proposal faced the AMPS challenge in 2012 because it appeared to require the same information to be present in two different places (radiation AND interior) without a single observer being able to verify both — hence the monogamy violation.

39.2 GDOS complementarity

In GDOS, complementarity has a clean geometric origin: the exterior and infalling observers access *different projections of the 6D bulk state*. The exterior observer sees $\rho_\partial^{\text{out}}$, the partial trace over the interior 4D region and the compact dimensions. The infalling observer sees a different partial trace — one that retains the 4D interior but still traces over the compact dimensions, which remain inaccessible from their reference frame.

The two observers see different density matrices, but both are valid OGGEF projections of the same unique 6D bulk state. There is no contradiction, and no firewall is required to enforce

consistency, because the complementarity is just two different partial traces of the same underlying state.

40 The four laws of black-hole mechanics as GDOS statements

The Bardeen–Carter–Hawking laws (Bardeen et al., 1973) were originally a formal *analogy* between black-hole mechanics and thermodynamics. In GG-Theory the analogy is an identity: each law is a statement about the OGGEp port and the GDOS open-system dynamics.

Zeroth law (constant κ). The surface gravity κ is constant over the horizon of a stationary black hole. In GDOS language κ is the single scale/Unruh–Tolman rate of the OGGEp port (§26.4); its constancy says the port is characterized by one temperature $T = \hbar\kappa/2\pi ck_B$ — the boundary CPTP channel has a single thermal scale, the gravitational counterpart of a body in equilibrium having one temperature. In regime language: the zeroth law is basin-floor equilibrium — the certificate that the black hole is a genuine Geometric Attractor in residence, not a passing configuration.

First law (energy balance). $dM = \frac{\kappa}{8\pi G_N} dA + \Omega_H dJ + \Phi_H dQ$. With the induced-gravity entropy $S = A/4G_N$ and $T = \hbar\kappa/2\pi$, this is exactly $dM = T dS + \Omega_H dJ + \Phi_H dQ$ (§31.8); the area term is the geometric piece F_{GEP} of the GFF ledger (§33), and the first law is the thermodynamic energy balance of the OGGEp-projected open system — the bookkeeping of the descent, debiting the ledger as the exit channel carries mass–energy away.

Second law (generalized entropy increase). Hawking’s area theorem $\delta A \geq 0$, sharpened by Bekenstein to the generalized second law $\delta S_{\text{gen}} \geq 0$ (Bekenstein, 1973). This is the sharpest GG-Theory reading of the four: it is the *monotonicity of the GFF ledger under the GDOS CPTP/Markov semigroup* — the H -theorem of the open-system dynamics, equivalent to the data-processing inequality for the boundary channel. The irreversibility is the forward renormalization-group arrow of the non-compact scale fiber (the GDOS generator is the dilatation generator along I_r); the generalized second law and the arrow of measurement are one and the same arrow. The second law is thus not an analogy to thermodynamics but the open-system H -theorem itself — and, in the runtime vocabulary, it *is* the Lyapunov monotonicity of GRIP-In (the descent theorem of (Arisaka, 2026, A11-Origin), anchored in the GFF ledger structure theorems of (Arisaka, 2026, A6-GRIP)) evaluated on the horizon stratum. The *direction* of this identity is no longer carried by citation: it is derived link by link, with its ceiling declared, in §41 (Lemma 2.2, L-BH-Arrow) — the arrow consumed from its constitutional owner, the monotonicity supplied by the channel identity, and the loss priced in area by the clock’s thermal face.

Third law (unattainability of extremality). $\kappa \rightarrow 0$ (the extremal limit) cannot be reached in a finite number of operations. In GDOS language extremality is the boundary of the admissible OGGEp-port family — a port of zero temperature, whose CPTP channel would have a pure output — and the third law is the statement that the open-system dynamics cannot drive the port to a pure, zero-entropy-production output in finite steps: the gravitational counterpart of the unattainability

of absolute zero. In regime language the third law acquires a pointed reading: an extremal hole would be a zero-temperature port — an attractor whose Hawking exit channel has been closed *by hand* — and the law says no finite procedure can manufacture such an artificial ground state. Nature’s GGAs are earned by descent, not engineered by extremization.

Table 4: The four laws as mechanics, thermodynamics, and GDOS/OGGEP statements. The first three columns are standard and are reprinted so that the fourth can be read against them. That fourth column is the claim: each law of black-hole mechanics is a statement about an open system at a port, and the analogy Bardeen, Carter and Hawking could only note in 1973 becomes an identification. The first-law row carries the weight, because it is where the area term becomes a free-energy term rather than merely resembling one.

Law	BH mechanics	Thermodynamics	GG-Theory/GDOS reading
0th	κ constant on horizon	T uniform in equilibrium	OGGEP port has one scale/Unruh rate κ ; one channel temperature
1st	$dM = \frac{\kappa}{8\pi G_N} dA + \Omega dJ + \Phi dQ$	$dE = T dS + \text{work}$	Energy balance of the projected open system; $A/4G_N = F_{\text{GEP}}$
2nd	$\delta A \geq 0$; $\delta S_{\text{gen}} \geq 0$	$\delta S \geq 0$	Monotonicity of the GFF ledger under the GDOS CPTP semigroup (H -theorem; RG arrow)
3rd	$\kappa \rightarrow 0$ unattainable	$T \rightarrow 0$ unattainable	Port cannot be driven to a pure (zero- T) output in finite steps

41 The direction of the second law: the arrow consumed, the price computed

The second-law paragraph of §40 makes the strongest single claim of the four-laws reading: that the generalized second law and the arrow of measurement are one and the same arrow. A claim of that size should not rest on a citation. This section derives its *direction* — why S_{gen} increases toward the future rather than the past, and at what rate — from named inputs, each consumed from its constitutional owner, with the one genuinely hard remaining step declared open rather than absorbed. The result is quantitative: in the solvable class that models the horizon exchange, the generalized-entropy production between successive cuts of the record is not merely non-negative — it is *exactly* the relative entropy spent toward the clock’s thermal face, with closed-form values and a sharp equality point.

41.1 Three links, three owners

The claim “ $\Delta S_{\text{gen}} \geq 0$ toward the future” decomposes into three logically independent links, and the cohort’s custody discipline assigns each to a different owner. *The paper that uses a link is not the paper that proves it*; this module consumes all three and adds only the horizon translation. The cohort’s survey of how a conservation law comes to stand without a symmetry beneath it is GG-W1-Symmetry (Arisaka, 2026, W1-Symmetry).

Table 5: The three links of the second law’s direction, their owners, and their status in this paper. The custody chain for the arrow is the cohort’s standing discipline: AX-3 owns it, the CPT companion derives it at law level over the exactly T-symmetric bulk, GG-A5-GDOS implements it dynamically, and the runtime papers — this one now among them — consume it and never re-derive it.

Link	Content	Owner	Status here
D1	Which time direction the boundary channel family composes in: the record is ordered, and the ordering is not the horizon’s choice	AX-3 (non-invertibility; Arisaka, 2026, A1-Constitution); law-level derivation in the CPT companion (Arisaka, 2026, P5-CPT); dynamical implementation GG-A5-GDOS (Arisaka, 2026, A5-GDOS)	Consumed, never re-derived
D2	No channel increases distinguishability: relative entropy contracts under every CPTP map	Identity of quantum channels (the data-processing inequality, Appendix O); carries <i>no</i> arrow content of its own	Identity; verified numerically
D3	The horizon translation: contraction of relative entropy toward the clock’s KMS face <i>is</i> generalized-entropy increase, with the loss priced in area	This paper: T-BH-5 (the face) + the induced first law (§31.8) + the exchange-channel identity below	Derived here (quasi-free exchange class; linearized first law)

Two remarks before the lemma, because they are where confusion most often enters. *First*, D2 is directionless: the data-processing inequality constrains any channel, run in any direction; by itself it cannot distinguish future from past. The direction enters once and only once, at D1 — the constitutional statement that the boundary description composes as channels toward the future and as nothing lawful toward the past. Papers that “derive the arrow from the second law,” or the reverse, are circling between D1 and D2 without naming them; the cohort’s custody discipline exists to prevent exactly that circle. *Second*, D3 is where all the physics lives: it converts an information inequality into a geometric one, and the conversion rate — area per unit of spent distinguishability — is fixed by the same two theorems this paper has already paid for: the KMS face of the clock (Theorem 5) and the induced first law (§31.8).

41.2 The lemma

Lemma 2.2 (The direction of the generalized second law (L-BH-Arrow)). *Let the boundary record at the horizon evolve by the CPTP family of the OGGE channel (§34), ordered by the constitutional arrow (link D1, consumed). Then: (i) (monotone) for every pair of record times $t_1 \leq t_2$ in the constitutional ordering and every reference state σ carried by the same channel, $S(\rho(t_2) \parallel \sigma(t_2)) \leq S(\rho(t_1) \parallel \sigma(t_1))$ — the data-processing inequality, an identity of CPTP maps; this holds for the total map between any two record times and therefore survives the expiry of port freshness (Lemma 2.1): staleness breaks the Markov semigroup form, never the direction; (ii) (the price list) in the Gaussian exchange class that models the horizon channel at frequency ω — transmissivity $\Gamma \in (0, 1]$, the clock’s KMS state σ_{KMS} of occupation $\bar{n}_H = (e^{\hbar\omega/k_B T_H} - 1)^{-1}$*

as fixed point, thermal record states, and the linearized induced first law $\delta(A/4G_N) = \delta E_H/T_H$ at fixed (J, Q) — the generalized-entropy production between successive record cuts is exactly the spent relative entropy toward the thermal face,

$$\Delta S_{\text{gen}} = S(\rho_{\text{in}} \parallel \sigma_{\text{KMS}}) - S(\rho_{\text{out}} \parallel \sigma_{\text{KMS}}) \geq 0, \quad (58)$$

with equality if and only if the record state is the fixed point ($\bar{n}_b = \bar{n}_H$, i.e. $T_b = T_H$); (iii) (closed forms) in that class, per mode of frequency ω : pure emission ($\bar{n}_b = 0$, $\Gamma = 1$) produces $\Delta S_{\text{gen}} = \ln(1 + \bar{n}_H) = -\ln(1 - e^{-\hbar\omega/k_B T_H}) > 0$; graybody emission produces $\Delta S_{\text{gen}} = g(\Gamma \bar{n}_H) - \Gamma \hbar\omega \bar{n}_H/k_B T_H \geq 0$; and exchange with a bath at T_b produces $\Delta S_{\text{gen}} = g(\bar{n}_{\text{rec}}) - g(\bar{n}_b) + \Gamma(\bar{n}_b - \bar{n}_H) \hbar\omega/k_B T_H \geq 0$ with $\bar{n}_{\text{rec}} = (1 - \Gamma)\bar{n}_b + \Gamma\bar{n}_H$ and $g(n) = (n+1)\ln(n+1) - n\ln n$ — positive for net emission and for net absorption alike. Grade: (i) identity on a consumed direction; (ii)–(iii) Derived (exchange class), Derived-conditional as a statement about the astrophysical horizon (quasi-free grade, linearized first law). The general interacting statement is Open (§41.3).

Proof. Step 1 (the direction, consumed). By AX-3 the boundary description is the output of a one-way dictionary; the law-level location of the arrow — the dictionary’s one-way reading, derived over an exactly T-symmetric bulk — is Part IX of the CPT companion (Arisaka, 2026, P5-CPT), and its dynamical implementation as the forward-composing CPTP family is GG-A5-GDOS’s (Arisaka, 2026, A5-GDOS). Nothing in this proof re-derives any of it; the record times $t_1 \leq t_2$ of clause (i) are ordered by that consumed arrow, full stop.

Step 2 (the monotone). Clause (i) is the data-processing inequality of Appendix O applied to the total map $\Lambda_{t_2 \leftarrow t_1}$, which is CPTP by T2 of GG-A5-GDOS regardless of whether the Markov factorization holds — hence the robustness against stale ports noted in the clause: the freshness window of Lemma 2.1 bounds where the *semigroup* form is exact, while the direction statement needs only complete positivity of the total map. Verified numerically: 300 Haar-random Stinespring channels on random full-rank state pairs give contraction in every trial (verification suite, www.arisaka-gg.org).

Step 3 (the reference and the price). Take $\sigma = \sigma_{\text{KMS}}$, the thermal face the clock forces on the regular exterior state (Theorem 5(iii)) — the unique reference the horizon itself supplies. For single-mode thermal states the relative entropy is

$$S(\sigma_n \parallel \sigma_m) = -g(n) + n \ln \frac{m+1}{m} + \ln(m+1), \quad g(n) = (n+1)\ln(n+1) - n\ln n, \quad (59)$$

and for $m = \bar{n}_H$ the middle coefficient is $\ln((\bar{n}_H+1)/\bar{n}_H) = \hbar\omega/k_B T_H$: the modular weight of the KMS reference is the Clausius weight of the horizon. This is the entire content of the price list: by the induced first law (§31.8, at fixed J, Q , quasi-static, linearized), energy δE_H crossing the horizon moves the area term by $\delta(A/4G_N) = \delta E_H/T_H$, which is precisely the $n \hbar\omega/k_B T_H$ term of (59) mode by mode.

Step 4 (the exchange identity). Model the horizon channel at frequency ω as the Gaussian exchange map of transmissivity Γ with fixed point σ_{KMS} : a thermal record mode $\bar{n}_b$ goes to

$\bar{n}_{\text{rec}} = (1 - \Gamma)\bar{n}_b + \Gamma\bar{n}_H$, and the energy delivered to the hole is $\delta E_H = \Gamma\hbar\omega(\bar{n}_b - \bar{n}_H)$. Then

$$\begin{aligned}\Delta S_{\text{gen}} &= \underbrace{[g(\bar{n}_{\text{rec}}) - g(\bar{n}_b)]}_{\Delta S_{\text{out}}} + \underbrace{\Gamma(\bar{n}_b - \bar{n}_H) \frac{\hbar\omega}{k_B T_H}}_{\Delta(A/4G_N)} \\ &= S(\sigma_{\bar{n}_b} \parallel \sigma_{\bar{n}_H}) - S(\sigma_{\bar{n}_{\text{rec}}} \parallel \sigma_{\bar{n}_H}),\end{aligned}\tag{60}$$

by direct substitution of (59): the g -terms and the Clausius terms reassemble into the two relative entropies with nothing left over. The right side is non-negative because the exchange map is CPTP with $\sigma_{\bar{n}_H}$ as fixed point — clause (i) applied with the horizon’s own reference — and vanishes iff $\bar{n}_b = \bar{n}_H$. This is the black-hole instance of the entropy-production theorem for quantum dynamical semigroups with a KMS fixed point (Spohn, 1978), and the relative-entropy reading of horizon entropy bounds (Casini, 2008), here composed with the induced first law so that the produced entropy is carried by *area*. Numerically: the identity (60) holds to 3.6×10^{-15} over a grid of $(\omega/T_H, T_b/T_H, \Gamma)$; the minimum of ΔS_{gen} over a 1.2×10^5 -point sweep is $+1.2 \times 10^{-8}$ (attained approaching $\bar{n}_b \rightarrow \bar{n}_H$), and the closed forms of clause (iii) are reproduced to 10^{-12} (verification suite).

Step 5 (closed forms and their reading). Setting $\bar{n}_b = 0$, $\Gamma = 1$ in (60) gives the pure Hawking emission of one mode:

$$\Delta S_{\text{gen}}|_{\text{emit}} = \ln(1 + \bar{n}_H) = -\ln(1 - e^{-\hbar\omega/k_B T_H}) = S(|0\rangle\langle 0| \parallel \sigma_{\text{KMS}}) > 0:\tag{61}$$

the generalized entropy produced by emitting one mode is exactly the relative entropy of the vacuum to the clock’s thermal face — the distinguishability the exterior surrenders by letting the port thermalize that mode. Soft modes ($\hbar\omega \ll k_B T_H$) are the expensive ones, $\Delta S_{\text{gen}} \simeq \ln(k_B T_H / \hbar\omega) \rightarrow \infty$; hard modes are nearly free, $\Delta S_{\text{gen}} \simeq e^{-\hbar\omega/k_B T_H} \rightarrow 0$; summed against the Planck weight and graybody factors this is the familiar statement that Hawking radiation carries more entropy than the area it spends, here obtained per mode and in closed form. The two-temperature form of clause (iii) settles the conditional-GGA regime of §42.2 quantitatively: for $T_b > T_H$ (today’s sky, CMB-fed holes) the flux reverses — the hole grows — yet $\Delta S_{\text{gen}} \geq 0$ unchanged, and the equality point $T_b = T_H$ is exactly the crossover mass $M_\times$ of Eq. (63): *the conditional-GGA crossover is the GSL’s equality point, the one mass at which the horizon exchanges without producing entropy.*

Step 6 (what is not proven). The identity (60) is exact in its class; the class is a model of the horizon channel, not the interacting gravitating field. The ceiling is declared in §41.3. ■

41.3 The declared ceiling

A Wall-class proof of the generalized second law — all semiclassical orders, arbitrary horizon cuts, interacting fields, with the monotonicity of the relative entropy run against the half-sided modular structure of the horizon algebra (Wall, 2012) — is *not* claimed here and is not within reach of the present apparatus: it requires input of ANEC grade and Hadamard control that this module imports rather than owns. What this section establishes is strictly less and strictly named: the *direction* is constitutional and consumed (D1); the *monotone* is a channel identity (D2); and the

price is computed exactly in the exchange class that models the port, with the linearized induced first law as the conversion rate (D3). The promotion path — from the exchange class to arbitrary cuts of the interacting horizon — is registered as an open problem of this module in §66, alongside the species sum. The Lyapunov reading of §40 (the GRIP-In descent theorem evaluated on the horizon stratum) is thereby supported at the same grade as its inputs: direction derived, production non-negative in the modeled class, full interacting monotonicity open.

Physics intuition. A customs post on a one-way bridge. Which way the traffic flows is decreed far upstream — by the constitution, not by the customs officer (link D1). That no cargo gains value in transit is not a local regulation either; it is the law of every checkpoint everywhere (link D2, the data-processing inequality). What the horizon adds — the only thing it adds — is a *tariff schedule* (link D3): every mode that crosses pays in distinguishability, the price list is printed on the clock’s thermal face, and the till is the area. One line of the schedule is worth memorizing: a single emitted mode costs exactly $\ln(1 + \bar{n}_H)$ — expensive for soft quanta, nearly free for hard ones — and the toll is zero only when the two banks stand at the same temperature, which for a black hole in today’s sky happens at exactly one mass, the crossover $M_\times$ where its conditional ground-state status is decided. The second law’s direction was never the horizon’s to choose; the horizon only sets the prices.

42 The black hole as a way-station attractor: where the descent pauses, and where it ends

42.1 Hawking radiation is the exit channel

The runtime vocabulary makes one structural fact about black holes visible that the thermodynamic language leaves implicit. A black hole is a *Geometric Attractor* of the gravitational stratum — a deep, stable basin of the free-energy ledger, which is why it is no-hair simple and why perturbed holes relax back to Kerr. But it is *not* the Ground Geometric Attractor, because it possesses an exit channel: **Hawking radiation is the channel-mediated exit of GRIP-In on the gravitational stratum** — this paper’s own identification, extending to gravity the catalog of exit channels (weak decay for unstable matter, Hubble friction for the cosmological dark sector) whose named instances are GG-A11-Origin’s (Arisaka, 2026, A11-Origin). A stellar-mass black hole is therefore the most extreme *way station* in nature: residence time $\sim 10^{67}$ yr, the slowest single GRIP-In step known. The descent does not end at the black hole; it ends beyond it, at the fully radiated configuration — which then relaxes, with the rest of the universe, toward the de Sitter vacuum, the cosmological stratum’s true GGA (Arisaka, 2026, A10-Cosmos). In one sentence: *the information paradox was the worry that the way station destroys the luggage; the GDOS reading of this Part shows the luggage is forwarded.*

One precision keeps this claim exact. The black hole *maximizes entropy* at fixed energy (the Bekenstein bound) — how can a maximum be a way station? Because the runtime does not descend entropy; it descends *free energy* ($GFF = F_{\text{GEP}} + F_{\text{CEP}}$, §33; the ledger is minimized at fixed conserved

charges, and on the horizon stratum its monotone content is exactly $\delta S_{\text{gen}} \geq 0$ — lower ledger and higher generalized entropy are one direction, not two claims), and the generalized second law of §40 says precisely that the radiated configuration has *higher* generalized entropy and *lower* ledger than the hole that produced it. The GSL is not in tension with the way-station reading; it is the signpost of the continuing descent.

42.2 Today’s black holes are *conditional* ground states

Whether an exit channel is open is not an intrinsic property of the attractor; it is computed *including the boundary conditions* — the boundary-port clause of the open-system law (Arisaka, 2026, A5-GDOS). For a black hole the relevant port is the cosmic bath: a hole of mass M radiates at

$$T_H = \frac{\hbar c^3}{8\pi G_N M k_B} \simeq 6.2 \times 10^{-8} \text{ K} \frac{M_\odot}{M}, \quad (62)$$

and sits in the CMB at $T_{\text{CMB}} = 2.725 \text{ K}$. For $T_H < T_{\text{CMB}}$ the channel runs *backward*: the hole absorbs more than it emits and grows. Since every astrophysical black hole has T_H at least seven orders of magnitude below the CMB, **every black hole in today’s sky is conditionally exitless** — **a conditional** GGA, in precisely the sense in which neutron matter is the conditional ground state inside a neutron star. The crossover is sharp:

$$T_H = T_{\text{CMB}} \iff M_\times \simeq 4.5 \times 10^{22} \text{ kg} \simeq 0.61 M_{\text{Moon}} : \quad (63)$$

primordial holes below $M_\times$ are way-station GAs in active descent today; everything above it is a conditional GGA. And the condition expires on a calculable schedule: as the CMB cools, $T_{\text{CMB}}(t) \propto 1/a(t)$, the exit channel opens for ever-larger holes,

$$M_{\text{open}}(t) \propto a(t), \quad (64)$$

so a solar-mass hole begins its net descent when the universe has expanded by a further factor $\sim 4 \times 10^7$. *Ground-state-ness is boundary-condition-relative, and for black holes the expiration date is a redshift.* No other framework in black-hole science has the vocabulary in which this sentence can even be written; in the runtime grammar it is a one-line corollary of the boundary-port clause — and we key it as such:

Corollary 2.1 (The conditional-GGA crossover (C-BH-1)). *With the boundary port the cosmic bath at $T_{\text{CMB}}(t)$, a black hole is a conditional GGA if and only if $T_H < T_{\text{CMB}}$, with the crossover mass given by Eq. (63) and the opening schedule by Eq. (64); and by Lemma 2.2 the crossover is simultaneously the generalized second law’s equality point — the one mass at which the horizon exchanges with its bath without producing entropy. Grade: Derived (the induced first law plus the exchange-class identity), conditional on the quasi-static reading of the bath coupling.*




Figure 13: The black-hole life cycle in the two-regime grammar. The entire column is GRIP-In on a fixed gravitational landscape: capture through the instability gates, residence at the way-station attractor, and exit through the Hawking channel toward the radiated configuration — the stratum’s true ground state (GGA). Nothing in this column reorganizes the landscape; the one genuine GRIP-Out of black-hole physics is the Hawking–Page transition (§36), where a control parameter rebuilds the saddle structure. Right: today the exit channel of every astrophysical hole runs backward ($T_H < T_{\text{CMB}}$), making it a *conditional* GGA whose ground-state status expires on the cosmological schedule of Eq. (64).

42.3 Sorting the transitions: capture, milestone, and the one true GRIP-Out

The two-regime question — *does the landscape change?* — sorts the black-hole life cycle cleanly, and two common misreadings are worth pre-empting. *Formation is not a GRIP-Out.* Stellar collapse is descent on the *fixed* gravitational landscape: the Chandrasekhar and Tolman–Oppenheimer–Volkoff thresholds are the saddle gates the trajectory crosses, and the capture into the black-hole basin is ordinary GRIP-In — nothing about the landscape itself changes when a star collapses, however violent the infall. Likewise a binary merger is *basin coalescence* on the fixed landscape (two attractors of the joint system merging into a deeper one), and accretion is port-driven traffic into an existing basin. The one genuine GRIP-Out in this paper is the Hawking–Page transition of §36, where temperature — a control parameter of the ledger — reorganizes the saddle structure of the landscape itself. And the Page turnover is neither: it is a *milestone* of GRIP-In (§36), the moment a bookkeeping crossover occurs while the landscape stays put.

Figure 13 places the black hole on the descent.

43 The GDOS reading, in summary

The GDOS reading of black holes can now be told as one continuous story, and it is worth telling in prose before the predictions Part prices it.

A horizon is an OGGE port: the specific locus where the 4D exterior separates from an inaccessible interior. Hawking radiation is that port’s channel output, its thermal character automatic from the CPTP structure rather than re-derived metric by metric. The Bekenstein–Hawking entropy is the entanglement entropy of the projection, $S_{\text{BH}} = S_{\text{vN}}(\rho_{\partial}^{\text{out}})$, with the exact coefficient the business of Part V. And the Page curve’s shape follows from bulk unitarity plus

CPTP projection, with no replica wormholes required; its quantitative form is open. Its turnover has a name and a schedule: the cohort’s freshness clause (L-GDOS-Fresh; GG-A5-GDOS, Definition 5) holds at the horizon exactly until the Page time and fails after it by counting alone, and the branch switch of T-BH-2 is that expiry, named and verified in the coarse-grained model (L-BH-Fresh).

The paradox-adjacent structures dissolve rather than resist. Firewalls are an artifact of 4D-only thinking: at the 6D bulk level, the monogamy violation never happens. The singularity is a coordinate artifact of the 4D projection, smooth in the bulk. And complementarity is two partial traces of the same 6D state — two lawful pages of one ledger, not two competing copies of one page.

The theorems anchor what would otherwise be metaphor. Rotation and collapse are the same custody: the ergoregion defeats the rotating-frame selection and the anchor moves to the asymptotic frame, with superradiance as GDOS pumping (T-BH-3), while the collapse epoch factorizes into a universal thermal channel times a decaying transient (T-BH-4). Thermality is the clock’s KMS face as a theorem: the positive-frequency selection, the complexified-horizon continuation, and the imaginary-time periodicity are consumed from the cohort’s Wick-quartet theorem T-P5-B1 rather than imported as prescriptions, the geometry contributing only κ (T-BH-5). And the generalized second law is derived with a declared ceiling: the constitutional arrow consumed from its owner, the data-processing inequality supplying the monotone, and the clock’s KMS face pricing the loss in area — ΔS_{gen} equals the spent relative entropy exactly in the exchange class, $\ln(1 + \bar{n}_H)$ per emitted mode, with equality only at $T_b = T_H$, and the Wall-class statement registered open (L-BH-Arrow).

The Part in one breath, and the door to the next. This re-reading resolves the information paradox at the level of physical geometry, not formal duality. The physics intuition to carry forward: *information is never lost, only displaced — and its return is an assumption expiring on schedule, not a mechanism switching on.* A resolution this tidy must now put numbers where instruments can reach. The next Part prices the reading — shadows, ringdowns, primordial ladders, analog horizons, and the universe’s own horizon — and shows that the framework’s boldest quantitative claim is a precise, dated, and falsifiable silence.

44 Part VI in plain English: why the information comes back, and why nothing switches on to bring it

The information paradox in one sentence: if the radiation is thermal and the hole evaporates, the information is gone, and it cannot be.

This Part’s answer is that the radiation is thermal only while a particular assumption holds, and that the assumption has an expiry date built into it from the beginning. The framework keeps a ledger with two terms, an area term and a bulk-entanglement term, and the physical entropy is whichever combination of them is smaller. Early on the area term dominates and the radiation looks thermal. Later the other term does, and it stops looking thermal. Nothing switches on. No new physics enters at the halfway point. What changes is which of two terms that were both there all along is winning, and the crossover is the Page curve.

Two separately puzzling things fall out of that. The firewall argument assumed the radiation’s partner had to be an independent four-dimensional degree of freedom; here it is a bulk one, so the monogamy conflict does not arise and nothing forces a firewall on an infalling observer at the horizon. And the islands formula, which gets the right answer in controlled models by machinery whose meaning is still argued about, is recovered as the semiclassical display of the same ledger — the same arithmetic under a different name, which the Part sets out row by row rather than asserting.

The honest part is the direction of the second law. This Part does not derive the arrow. It consumes the arrow from a companion paper, computes what each exchange costs, and registers the general statement as open with a named ceiling. A reader who wants the arrow derived will not find it here, and the Part says which paper to read instead.

Part VII

Quantitative Predictions

Event-Horizon Telescope, Ringdown, Primordial Black Holes, and the Hawking Spectrum

A framework that resolves paradoxes owes the sky some numbers. This Part collects them, and it is honest up front about their shape: for astrophysical black holes the GG-Theory corrections to general relativity are suppressed by $(\ell_{\text{lock}}/r_s)^2$ — parts in 10^{46} for stellar-mass holes — so the framework’s signature is not a bump or an echo but a precise silence: pure Kerr wherever instruments can see, with the deviations budgeted row by row rather than waved away. Five observational windows carry the exposure, each with its own section below; the worked numbers for four representative holes come first, so the scales are concrete; and the Part closes with the full predictions table — every row either consistent or not yet tested, none fitted.

45 Overview of predictions

The GDOS reading of black holes is more than a conceptual re-framing; it makes specific observational predictions that can (in principle, and sometimes in practice) be tested. The key observational windows are five. EHT imaging reads the shadows of the supermassive black holes M87* and Sgr A* at horizon scales. Gravitational-wave ringdown reads the quasi-normal-mode spectrum of post-merger black holes detected by LIGO and, in the coming decade, LISA. The primordial-black-hole mass spectrum carries the module’s one model-level ladder prediction from the early universe. The Hawking-radiation spectrum is observable in principle for very small black holes, if such holes exist. And entanglement-entropy structure is testable indirectly, through the black-hole-like behavior of analog systems — acoustic and optical horizons in the laboratory. Each of these is discussed in detail below.

46 Four black holes, end to end

To make the predictions concrete, we run the explicit arithmetic for four representative objects, from the supermassive to the primordial. In every case GG-Theory reproduces the standard general-relativistic and thermodynamic values — the framework adds no deviation at observable precision — so the worked numbers are exactly the textbook ones; their purpose is to show the entropy, temperature, and lifetime that the induced-gravity result $S = A/4G_N$ assigns, and to make plain how small the GG correction is.

Table 6: Worked black-hole numbers (order of magnitude). r_s is the horizon radius, T_H the Hawking temperature, $S_{BH} = A/4G_N\hbar$ in units of k_B , and t_{evap} the evaporation time. These are the standard values, which GG-Theory reproduces; the GG correction is suppressed by $(\ell_{\text{lock}}/r_s)^2$ and is unobservable in every row (see §47, §48).

Object	Mass	r_s	T_H	S_{BH}/k_B	t_{evap}
M87*	$6.5 \times 10^9 M_\odot$	$1.9 \times 10^{13} \text{ m}$	$1 \times 10^{-17} \text{ K}$	4×10^{96}	10^{97} yr
Sgr A*	$4.3 \times 10^6 M_\odot$	$1.3 \times 10^{10} \text{ m}$	$1.5 \times 10^{-14} \text{ K}$	2×10^{90}	10^{87} yr
Stellar	$10 M_\odot$	$3 \times 10^4 \text{ m}$	$6 \times 10^{-9} \text{ K}$	1×10^{79}	10^{70} yr
PBH (now)	$5 \times 10^{11} \text{ kg}$	$7 \times 10^{-16} \text{ m}$	$2 \times 10^{11} \text{ K}$	7×10^{39}	$\sim 2 \times 10^{10} \text{ yr}$

The supermassive pair (M87*, Sgr A*). These are the two black holes the Event Horizon Telescope has imaged (Event Horizon Telescope Collaboration, 2022; Event Horizon Telescope Collaboration et al., 2019). Their Hawking temperatures (10^{-17} – 10^{-14} K) are far below the cosmic microwave background, so they absorb rather than radiate; their entropies (10^{90} – $10^{97} k_B$) are the largest single-object entropies in the universe. GG-Theory predicts the Kerr shadow to the precision EHT measures; the induced-gravity correction is suppressed by some $(\ell_{\text{lock}}/r_s)^2 \lll 1$ and is wholly unobservable.

The stellar black hole. A $10 M_\odot$ remnant of the kind LIGO observes in mergers has $S_{BH} \sim 10^{79} k_B$ — to be compared with the $\sim 10^{58} k_B$ thermodynamic entropy of the progenitor star, an illustration that horizon formation is the most entropy-generating process known.

The primordial black hole. A black hole of mass $\sim 5 \times 10^{11} \text{ kg}$ ($5 \times 10^{14} \text{ g}$) has an evaporation time of order the present age of the universe (with the multi-species emission rate; the photon–graviton-only clock of Appendix K, Step 3 runs an order of magnitude slower): such objects, if they formed, are completing their evaporation now, at $T_H \sim 10^{11} \text{ K}$ ($\sim$ tens of MeV). This is the one black hole for which the Page transition (§36) is, in principle, a near-term observable: its late-time radiation spectrum should turn over from thermal to information-restoring at the Page time $t_{\text{Page}} \approx (1 - 2^{-3/2}) t_{\text{evap}}$. This is the sharpest empirical handle the framework’s unitarity claim has.

47 EHT observations and horizon-scale geometry

47.1 What EHT measures

The Event Horizon Telescope images the shadow of supermassive black holes at $\sim 20 \mu\text{arcsec}$ resolution. The shadow is the projection of the photon sphere (at $r_p = 3G_N M/c^2$ for Schwarzschild) plus the dark region inside it. Its observed angular diameter sets constraints on the geometry at horizon scale.

For M87*, the 2019 image showed a shadow of angular diameter $\sim 42 \mu\text{arcsec}$, consistent with

the Kerr prediction within $\sim 10\%$ errors (which are dominated by astrophysical uncertainties, not by the imaging system).

47.2 GG-Theory prediction for the shadow

In GG-Theory, the 4D metric near a macroscopic black hole is the Kerr metric to very high precision, because the scale profile corrections are suppressed by the compactification scale $1/\ell_{\text{lock}} \sim M_{\text{lock}} = 3.85 \text{ TeV}$ (Arisaka, 2026, P1-KK), i.e. the compactification length is $\ell_{\text{lock}} \equiv \hbar/(M_{\text{lock}}c) \simeq 5 \times 10^{-20} \text{ m}$ (not the much shorter ultraviolet length $1/E_{\text{GUT}} \sim 10^{-32} \text{ m}$). The leading correction to the shadow size is suppressed by the square of the ratio of this length to the horizon radius:

$$\frac{\Delta r_{\text{shadow}}}{r_{\text{shadow}}} \sim \left(\frac{\ell_{\text{lock}}}{r_s} \right)^2 \sim \left(\frac{M_{\text{Pl}}}{M} \cdot \frac{M_{\text{Pl}}}{M_{\text{lock}}} \right)^2 \sim 10^{-66} \text{ for M87*}. \quad (65)$$

This is utterly below the EHT resolution. The GG-Theory prediction is: **the shadow is Kerr to parts in 10^{66} for the supermassive holes EHT images** (and to better than 10^{-46} or below for any astrophysical hole). There should be no observable deviation.

A deviation at the EHT level would *kill* GG-Theory: the theory would be required to produce Kerr exactly, so any sub-10% deviation would falsify it.

47.3 Future prospects: ngEHT

The next-generation EHT (ngEHT) program planned for the 2030s would improve resolution to $\sim 5 \mu\text{arcsec}$ and could detect polarization signatures, axion effects, and multiple photon sub-rings; the proposed space-VLBI Black Hole Explorer (BHEX), targeting launch in the early 2030s, would extend the array to Earth–space baselines and resolve the photon ring directly. The GG-Theory prediction remains: Kerr to all observable precision.

48 LIGO/LISA ringdown spectroscopy

48.1 Quasi-normal modes

When two black holes merge, the final state rings down via gravitational-wave emission of quasi-normal modes (QNMs). The QNM frequencies and damping times are characteristic of the final black-hole geometry: for Kerr, they depend only on M and $a = J/M$. The spectrum of QNMs is therefore a direct fingerprint of the black hole’s nature. The runtime vocabulary of §42 gives this fingerprint a structural identity: the post-merger ringdown is the *within-basin* piece of GRIP-In — relaxation toward the floor of the Kerr attractor — and the QNM spectrum is the *curvature of the basin*: the Hessian eigenvalues of the gravitational free-energy ledger at the attractor. In one sentence: *LIGO measures the second derivative of GFF at the Geometric Attractor* — the same functional whose minima, on other strata, are the electron, the proton, and the dark matter (Arisaka, 2026, A11-Origin). The merger that precedes the ringdown is basin coalescence on the fixed landscape; nothing in a binary inspiral is a GRIP-Out.

For the fundamental $\ell = m = 2$ QNM, the frequencies are approximately:

$$f_{\text{QNM}} \approx \frac{c^3}{2\pi G_N M} \left[1 - \frac{63}{100} (1 - a^*)^{3/10} \right], \quad (66)$$

where $a^* = ac^2/G_N M$ is the dimensionless spin parameter. For the $\sim 60 M_\odot$ remnant of a GW150914-like merger, the QNM frequency is ~ 250 Hz — in the LIGO band.

48.2 KK corrections to QNMs

In GG-Theory, the bulk has compact dimensions with radius $\ell_{\text{lock}} \equiv \hbar/(M_{\text{lock}} c) \sim 5 \times 10^{-20}$ m. The Kaluza–Klein (KK) spectrum of the compact direction has modes at frequencies

$$\omega_n^{\text{KK}} = n c / \ell_{\text{lock}} = n M_{\text{lock}} / \hbar \sim n \cdot 6 \times 10^{27} \text{ Hz}, \quad (67)$$

which is utterly outside the LIGO/LISA band. No direct KK signal is expected.

The leading effect on the ordinary QNM spectrum is through radiative corrections from virtual KK modes, giving

$$\frac{\Delta f_{\text{QNM}}}{f_{\text{QNM}}} \sim \left(\frac{\ell_{\text{lock}}}{r_s} \right)^2 \sim 10^{-48} \quad (\text{stellar-mass}), \quad (68)$$

again utterly below LIGO sensitivity.

48.3 LIGO/LISA prediction

The GG-Theory prediction for LIGO/LISA ringdown spectroscopy is: pure Kerr QNMs with sub- 10^{-48} corrections for stellar-mass remnants. No observable KK modes, no firewalls, no “echoes” from the horizon. This is in sharp contrast with the alternatives. Fuzzball proposals predict “echoes” from the fuzzball radius at r_f outside the horizon — post-ringdown gravitational-wave signals. Firewalls would potentially modify the effective boundary condition at the horizon. And extra-dimension models with large extra dimensions could modify the QNM spectrum at an observable level.

A detection of gravitational-wave echoes, or of QNM spectrum deviations from Kerr beyond astrophysical uncertainties, would be difficult to accommodate in GG-Theory.

49 Primordial black holes and the early universe

49.1 Primordial black-hole formation

Primordial black holes (PBHs) are black holes that may have formed in the very early universe from large density perturbations. Their mass spectrum depends on the power spectrum of primordial perturbations and the thermal history of the early universe. In the regime grammar that the companion cosmology paper applies to the whole early universe (Arisaka, 2026, A10-Cosmos), PBH formation is a *nested* event: the landscape-reorganizing transitions of the first instants (inflationary exit, reheating) are the GRIP-Outs, and PBH capture is GRIP-In on the landscapes those events left

behind — which is why the PBH mass spectrum is a fossil of the thermal history rather than an independent sector.

In standard cosmology, PBHs can form with masses ranging from near-Planck (10^{-8} kg) to many solar masses. Those with mass $\lesssim 5 \times 10^{11}$ kg (5×10^{14} g) would have evaporated by now via Hawking radiation (Arisaka, 2026, A10-Cosmos); those with mass $\gtrsim 5 \times 10^{11}$ kg would still be around and could in principle contribute to dark matter — though within the cohort the dark-matter slot is saturated by the Kaxion of the companion C1 paper (Arisaka, 2026, C1-Kaxion), which bounds any primordial-black-hole contribution to a subdominant share (Arisaka, 2026, A10-Cosmos, T-Cosmos-14).

49.2 GG-Theory PBH mass spectrum

In GG-Theory the early universe has a derived thermal history: a smooth inflationary plateau one e-fold below the GUT edge $E_{\text{GUT}} \simeq 2.3 \times 10^{16}$ GeV, a smooth exit, and a prompt, barrier-free crossing of the lock scale (Arisaka, 2026, A10-Cosmos). That history is what makes primordial black holes scarce. The minimal module’s scalar spectrum is red and smooth, the collapse fraction $\beta(M) \sim \exp[-\delta_c^2/2\sigma^2(M)]$ is exponentially small, and any population is at most a trace of the dark matter — the cosmology companion’s T-Cosmos-14, which this paper consumes and does not re-derive. If a narrow small-scale feature nonetheless arises, the OGN-5 integer structure constrains where it sits.

A conjectural prediction, fenced as model-level: the scale structure may imprint preferred mass *ratios* on the PBH spectrum. Successive bulk eigenmodes of the bulk are spaced geometrically, so if PBH formation is modulated by that structure the preferred masses form a ladder of *ratios*

$$\frac{M_n^{\text{PBH}}}{M_{n+1}^{\text{PBH}}} \sim e^{k \Delta r_n} \sim e^7 \simeq 1.1 \times 10^3, \quad (69)$$

with Δr_n a unit step of the radial coordinate. Two cautions are essential and we state them plainly. First, this fixes only *ratios*: the absolute mass scale of any rung is set by the horizon mass at the formation epoch, $M_{\text{PBH}} \sim M_{\text{Pl}}(t_{\text{form}}/t_{\text{Pl}})$, which is *far above* the Planck mass and is *not* predicted by the scale integers — so the spectrum floats freely up and down the mass axis until an absolute anchor is supplied. Second, the exponent k (and hence whether the ladder step is exactly e^7) is not derived here. The whole construction is therefore a model-level conjecture, not a theorem of the cohort, and is labeled as such throughout.

49.3 Observational tests

Searches for PBH dark matter run on four fronts: femtolensing and microlensing constraints (OGLE, Subaru HSC); gravitational-wave signatures of PBH mergers (LIGO); effects on the CMB, through Thomson scattering by evaporation products; and the disruption of stellar populations. The minimal module predicts that primordial black holes are scarce, so the ladder says where a rare population would sit, not that one exists. A well-populated PBH mass function whose peak spacing follows the e^7 ratios would be strong support for the conjecture; one whose spacing contradicts them

fires F-BH-4 and retires the ladder — the conjecture, not the module (a single detection does not fire it); and the absence of any population is the minimal module’s own expectation, not a kill.

50 Hawking spectrum corrections

50.1 Standard Hawking spectrum

The standard Hawking radiation spectrum (for a Schwarzschild black hole) is thermal:

$$\frac{dN}{dt d\omega d\Omega} = \frac{\Gamma_\ell(\omega)\omega^2}{(2\pi)^3(e^{\hbar\omega/k_B T_H} - 1)}. \quad (70)$$

The temperature T_H is set by the horizon’s surface gravity, and the graybody factors $\Gamma_\ell(\omega)$ depend on the scattering of radiation off the effective potential outside the horizon.

50.2 GG-Theory corrections

In GG-Theory, there are three classes of corrections to this spectrum:

(1) High-frequency cutoff. The spectrum should show an ultraviolet cutoff at the compactification scale:

$$\hbar\omega_{UV} \sim \hbar c/\ell_{\text{lock}} \sim M_{\text{lock}}. \quad (71)$$

Beyond this frequency, the bulk KK modes become active and the effective spectrum turns over. For astrophysical black holes, $M_{\text{lock}} \gg T_H$ so this correction is unobservable.

(2) Low-frequency small-mass cutoff. For very small black holes (near-Planck-mass), the back-reaction of Hawking radiation on the geometry becomes important, and the scale profile may not be stable. The spectrum may show modifications in the final stages of evaporation.

(3) KK overtones in partial-wave decomposition. For black holes large enough that KK modes are relevant (e.g., primordial near-Planck-mass BHs), the angular partial-wave decomposition should include contributions from KK modes in addition to ordinary 4D partial waves.

50.3 Observable prospects

None of the GG-Theory corrections to the Hawking spectrum are observable for astrophysical black holes in the near future. The Hawking temperature for a $10 M_\odot$ black hole is $\sim 6 \times 10^{-9}$ K, which is far below the CMB and undetectable.

For primordial black holes near the present evaporation threshold, $M \sim 5 \times 10^{11}$ kg, $T_H \sim 2 \times 10^{11}$ K — some 20 MeV — and the Hawking spectrum peaks at tens of MeV, climbing into the GeV range only in the final stages. The last burst is a target for gamma-ray telescopes (Fermi-LAT, CTAO). The GG-Theory prediction for the spectrum of evaporating PBHs is indistinguishable from the standard Hawking prediction at the relevant energies (MeV to GeV), because the compactification scale is $M_{\text{lock}} = 3.85 \text{ TeV} \sim 4 \times 10^{12} \text{ eV}$, far above the observable range.

51 Entanglement-entropy tests via analog black holes

51.1 Analog gravity

Since the 1980s, “analog gravity” systems have been explored: condensed-matter or optical systems that reproduce aspects of black-hole physics. Acoustic black holes realize a sonic horizon where a transonic flow reaches the speed of sound, and have been observed to emit analog Hawking radiation (Unruh 1981; experimental confirmation in the 2010s). Optical black holes use pulses in nonlinear optical fibers to mimic horizon physics. And Bose–Einstein-condensate black holes realize the same structure in trapped cold-atom flows. These analog systems cannot test quantum gravity, but they can test the *structure* of Hawking radiation: thermal spectrum, entanglement-entropy scaling, correlations between emitted modes and their partners.

51.2 Implications for GDOS

In GDOS, the analog-black-hole physics should match the astrophysical black-hole physics at the level of the partial-trace structure: the analog Hawking radiation should be thermal at the analog Hawking temperature; the entanglement entropy of the analog radiation with its partner modes should follow the appropriate analog Page curve; and no analog firewalls or exotic horizon effects should appear. Analog-black-hole experiments through mid-2026 are consistent with all three predictions. This is indirect confirmation of the GDOS picture.

52 Cosmological black holes and de Sitter horizons

52.1 The de Sitter horizon

In a cosmology with positive cosmological constant Λ , there is a cosmological horizon at distance $r_H = \sqrt{3/\Lambda}$. This is a causal surface: events beyond it cannot influence the observer. Like a black-hole horizon, it has an associated Gibbons-Hawking temperature $T_{\text{GH}} = H/(2\pi)$ where H is the Hubble rate.

In GG-Theory, the de Sitter horizon is also an OGGE port: the observer inside sees thermal radiation at T_{GH} from the horizon, and the entropy of the de Sitter horizon is $S_{\text{dS}} = A_H/(4G_N\hbar)$, with $A_H = 4\pi r_H^2$.

For our universe, with vacuum energy density $\rho_\Lambda \sim 10^{-122} M_{\text{Pl}}^4$, the numbers are set: the horizon radius is $r_H \sim 10^{26}$ m (the Hubble radius); the Gibbons–Hawking temperature is $T_{\text{GH}} \sim 10^{-30}$ K, utterly unobservable; and the entropy is $S_{\text{dS}} \sim 10^{122}$ — the maximum entropy the observable universe can hold.

This 10^{122} number is the famous “dark energy entropy.” It sets a universal upper bound on the information content of the observable universe.

52.2 The de Sitter entropy from induced gravity

The Gibbons–Hawking entropy $S_{\text{dS}} = A_H/4G_N$ (Gibbons and Hawking, 1977) is not merely *analogous* to the black-hole result — in GG-Theory it is the *same theorem*. The Euclidean de Sitter saddle is the round four-sphere S^4 , on which the cosmological horizon is the codimension-two bifurcation surface; the replica construction of §31.5 applies verbatim — on the same licensed arena, for the static-patch clock carries the identical four endowments and the Euclidean four-sphere is its analytic shore, the cap’s period being the geometry’s datum — with the conical-deficit identity $\int \sqrt{g} R \supset 4\pi(1-n)A_H$ on that surface and the induced Einstein–Hilbert action delivering

$$S_{\text{dS}} = \frac{A_H}{4G_N}, \quad \frac{1}{4} = \frac{4\pi}{16\pi}, \quad (72)$$

with the identical coefficient. The named-statement deck registers this as the text key C-BH-2: the de Sitter horizon entropy is Theorem 1 evaluated on the cosmological OGGEp port, inheriting its premises and its grade. The microstates are again the hidden-fiber holonomies (§31.6), now traced over the cosmological horizon. The cosmological horizon is the third panel of Figure 8: the OGGEp trace forced by cosmic expansion, with the observer’s accessible state the partial trace over modes beyond $r = c/H$. The exact-1/4 result of this paper is therefore not a black-hole accident; it is a statement about *any* OGGEp horizon, of which de Sitter is the cosmological instance.

52.3 Open de Sitter puzzles, stated honestly

De Sitter quantum gravity is less settled than the black-hole case, and we do not claim to resolve its specific difficulties. There is no asymptotic boundary on which to anchor a dual theory (the AdS/CFT crutch is unavailable); the meaning of the finite $S_{\text{dS}} \sim 10^{122}$ as a Hilbert-space dimension (the finite-dimensional-Hilbert-space conjecture) is open; and dS/CFT remains speculative. GG-Theory contributes the *same* structural reading it gives the black hole — the cosmological horizon as a geometrically-forced OGGEp port, the entropy as the entanglement of that forced trace, the coefficient from induced gravity — which is a candidate organizing principle, not a resolution of these dS-specific puzzles. We flag this explicitly so the de Sitter claims are not over-read: the coefficient and the OGGEp-port reading carry over; the deeper de Sitter quantum-gravity questions remain open.

52.4 Connection to the Hubble-tension companion (C2)

The companion Hubble paper (Arisaka, 2026, C2-Hubble) reads the local-versus-cosmological H_0 discrepancy on one derived Friedmann background: a single present rate — 66.6 ± 0.2 km/s/Mpc on the two-road reading the companions carry, 66.8 ± 0.2 from the geometry alone (Arisaka, 2026, A10-Cosmos) — read by the microwave background and by every ladder-free late probe, with the calibrated distance ladder’s 73.5 the predicted outlier, a calibration systematic the framework predicts and does not derive. The de Sitter horizon is where this paper meets that reading: every pipeline reads the background through a projection, and the cosmological horizon is the port through which the background is read at all. The present black-hole analysis and the C2 cosmological analysis

are thus two evaluations of one bulk-to-boundary projection — the black hole at a local horizon, the Hubble tension at the cosmological one — with the shared $S = A/4G_N$ entropy ledger (§52.2) as the common thread. This also connects to the $\Omega_m = 1/\pi$ inventory of (Arisaka, 2026, A10-Cosmos), which lives in the same boundary-projection framework. In the regime vocabulary the connection closes a circle: the de Sitter vacuum toward which the universe relaxes is the cosmological stratum’s GGA — one of the exactly two points where cosmic GRIP-In terminates (Arisaka, 2026, A10-Cosmos) — so the cosmological horizon of this section is not merely an analogue of the black-hole horizon: it is the port of the basin in which every black hole’s own descent (§42) finally ends.

53 The numerical table of predictions

Table 7: GG-Theory quantitative predictions for black-hole observables. Where observations exist, the pull is given. Most of the status column says consistent, and that should be read carefully rather than proudly: on every observable where a measurement exists the framework predicts Kerr, and so does general relativity, so the table earns nothing there. What it does earn is the location of the exposure — the rows marked not tested — so that agreement elsewhere cannot be mistaken for confirmation.

Observable	GG prediction	Measurement (or prospect)	Status
M87* shadow diameter	Kerr (42 μ arcsec)	EHT 2019: $42 \pm 3 \mu$ arcsec	consistent
Sgr A* shadow diameter	Kerr (51 μ arcsec)	EHT 2022: $51.8 \pm 2.3 \mu$ arcsec	consistent
GW150914 ringdown QNM	Kerr (250 Hz, $\tau=4$ ms)	LIGO 2015: consistent with Kerr	consistent
GW250114 ringdown (O4)	Kerr (two QNMs)	LVK 2025: fundamental + overtone, Kerr-consistent	consistent
PBH mass spectrum	Conjecture (model-level): mass ratios spaced by e^7 ; exponent not derived, absolute anchor open (§49)	None detected	not tested
Hawking temperature for $M_\odot$ BH	6×10^{-8} K	Unobservable	not tested
Bekenstein–Hawking entropy	$A/(4G_N\hbar)$	Consistent with horizon area in merger simulations	consistent
Page curve	Shape forced (T-BH-2); turnover dated (L-BH-Fresh); quantitative curve open	Not yet tested in real BH	not tested
Firewall signature	None predicted	None observed	consistent
Echoes from horizon	Not predicted	No convincing detection	consistent
Information paradox	Dissolved at the level of the reading (bulk closed, boundary open); quantitative Page curve open	(Theoretical consistency)	consistent

The Part in one breath, and the door to the next. Run the status column of Table 7 top to bottom and the shape of the module's exposure is plain: every tested row reads *consistent*, every untested row is dated or declared unfeasible, and no row is fitted. The physics intuition to carry forward: *for this framework, silence is not the absence of a prediction — it is the prediction*, and every sharper instrument raises the stakes of the bet. But predictions are cheap until they are bonded. The next Part posts the bond: the record of every shot fired at the module so far, the kill criteria under which it agrees to die, and the dated timeline on which the verdicts arrive.

54 Part VII in plain English: what the framework predicts where instruments can look, which is less than one would like

Here is the awkward fact this Part is built around. Everywhere an instrument can currently look, the framework predicts exactly what general relativity predicts.

The shadow of M87* is Kerr. The shadow of Sgr A* is Kerr. The ringdown of GW150914 is Kerr, and so is the two-mode ringdown resolved in 2025. There are no echoes, no firewall signature, and no departure from the thermal spectrum at any frequency anyone can reach. Every one of those rows reads consistent, and not one of them is evidence for this framework, because a theory that agrees with general relativity outside the horizon was always going to pass them.

So the Part is arranged around what that means rather than around the agreement. If the prediction is silence, then silence is what is at stake: every sharper instrument narrows the space in which a departure could hide, and the prediction gets riskier without ever getting more impressive. That is a real position to hold and a poor one to advertise, and the tables are laid out so a reader can see which it is.

The rows where the framework does say something are the ones nobody can test yet — the spacing of a primordial black-hole mass spectrum, the de Sitter horizon read as the same kind of port, corrections suppressed by a scale no instrument reaches. They are printed with their status attached rather than folded in among the successes.

Part VIII

Experiments — Past, Now, and Future

The Kill Criteria and the Evidence Base of the Black-Hole Module

This Part is the module’s confrontation record. Part VII stated what the framework predicts; this Part states what would kill it — the kill criteria, declared before the instruments report — and what the instruments have said so far, in the cumulative evidence base.

Physics intuition. A predictions Part lists what the theory says; an experiments Part posts the bond — the pre-registered conditions under which the theory agrees to die, and the running account of every shot fired at it so far. The milestones, the dated timeline, and the tables below are that bond for the black-hole module.

55 The record so far: what has been measured, and what it did to the module

Four instrument classes have already fired at this module, and the record of each shot is one-sided in the same direction.

Gravitational waves. Three hundred and ninety compact binary coalescences stand confirmed since 2015, roughly three quarters of them from the completed fourth observing run (§56). The module’s exposure is the ringdown: every quasi-normal-mode analysis to date — from GW150914’s fundamental mode to GW250114’s resolved fundamental-plus-overtone pair and its inspiral-to-ringdown horizon-area budget — is Kerr-consistent, and the area theorem is now an observationally confirmed statement. For this module that last point carries special weight: the generalized second law whose *direction* Lemma 2.2 derives has had its classical half (the area half) confirmed on a real merger. No echo candidate has survived scrutiny. F-BH-2 and F-BH-3 stand unfired at every sensitivity yet reached.

Horizon-scale imaging. Two shadows — M87* and Sgr A* — imaged, re-imaged across epochs, and now polarization-resolved; the shadow diameters sit on Kerr within errors dominated by astrophysics, and the epoch-to-epoch changes live entirely in the magnetized plasma, never in the shadow. F-BH-1 stands unfired.

Analog horizons. Laboratory horizons — acoustic, optical, condensate — have confirmed the structural content the port reading requires: thermal spectra at the analog temperature and the mode-pair entanglement of the trace. These platforms cannot kill the module (F-BH-5 is fenced

honestly: real-horizon spectra are unobservable), but they stress its *structure*, and the structure has held in every laboratory that has looked.

Primordial black holes. Microlensing surveys, subsolar gravitational-wave searches, and evaporation-photon limits continue to carve the PBH mass function without yet populating it. The module’s model-level ratio-ladder conjecture (F-BH-4) therefore remains untested rather than surviving: an honest “not yet”.

56 Now: the instruments in mid-2026

The LIGO–Virgo–KAGRA network completed its fourth observing run in November 2025 and is between runs: the GWTC-4.0 and GWTC-5.0 catalog releases (128 and 161 new events) are the run’s harvest, the final O4 data segment is still under analysis, and a six-month *intermediate run* (IR1) is announced to begin in late 2026, with the fifth observing run’s plans being refined toward A+–class sensitivity. The Event Horizon Telescope is in its multi-epoch era — annual campaigns, polarimetry on both shadows — with its next-generation extension in staged planning. The CTAO gamma-ray observatory, the instrument class that would see the final bursts of evaporating primordial holes, is under construction with its early-science era opening. And the module’s *external* hostages are on their own schedules: the collider search for the lock cluster, the haloscope scan of the Kaxion window, and the neutrino-ordering program — none of them black-hole instruments, all of them able to kill this paper’s reading from outside (§61.7).

57 The decade ahead: milestones, with the verdicts pre-registered

The verdicts below are signed now, so that no future reading of this module can be accused of moving after the shot.

2026–2027 — the intermediate run. IR1 returns the network to the sky for six months at O4-class-or-better sensitivity; tens of further ringdown-capable events are the expectation. *If right:* Kerr-consistent quasi-normal modes at every signal-to-noise ratio, echo searches silent — the null tightening, undramatically. *If wrong:* a confirmed QNM deviation above the 1% class or a reproducible echo train — F-BH-2 or F-BH-3 fires and the module retires, with no fallback.

Late 2020s — the fifth run. At A+–class sensitivity the loudest events reach the signal-to-noise class where two-mode spectroscopy is routine rather than exceptional, and the GW250114 measurement becomes a population statement. *If right:* the stacked ringdown catalog is Kerr to well below the percent class; the area theorem confirmed on dozens of mergers. *If wrong:* same rows, same firing rule — the bet does not change, only the stakes.

Early 2030s — the sharpened shadow. The next-generation EHT program and, if selected, the BHEX space-VLBI mission (targeting an early-2030s launch) push shadow astrometry toward the percent class and resolve the photon ring. *If right:* the shadow stays Kerr at every resolution — the module predicts that the sharper the image, the more boring it gets, and stakes itself on that boredom. *If wrong:* a confirmed shadow deviation at the 5% class — F-BH-1 fires.

Mid-2030s — LISA. The adopted space interferometer (launch targeted mid-2030s) observes extreme-mass-ratio inspirals: year-long orbits that map the central object’s multipole structure to the 10^{-4} class. This is the sharpest one-sided null the module will ever face — five orders below the current ringdown exposure. *If right:* Kerr multipoles, exactly, tens of thousands of cycles at a time. *If wrong:* any confirmed non-Kerr multipole retires the module from a direction no ground-based instrument can reach (the registered row is F-BH-10, Table 9).

Mid-to-late 2030s — the third generation. Einstein Telescope and Cosmic Explorer, in program definition through this decade, would return thousands of ringdowns per year. At that rate the ringdown null becomes a precision discipline, and F-BH-2’s 1% threshold is left far behind; the module’s exposure grows by the same factor as the catalog. *If right:* thousands of clean Kerr ringdowns per year, the null hardened into a precision law. *If wrong:* a reproducible spectral deviation at the 10^{-3} class fires F-BH-2 with statistics no single event can contest.

Decidable at a desk, any time — the internal row. F-BH-9 needs no instrument: the sign-resolved species sum either reproduces the induced $1/G_N$ or it does not. It is chartered as its own computation, and it is the one row on this timeline whose date the cohort itself controls. *If right:* the closure claim is promoted and O-BH-7’s answer becomes arithmetic. *If wrong:* F-BH-9 fires, retiring the closure reading and leaving the $1/4$ standing on its own premises (Table 24).

Physics intuition. A theory whose falsifiable content is carried by the null does not wait nervously for the data; it stands still on the rifle range, in writing, while the instruments buy longer barrels. Every entry above is the same single bet — *the horizon is exactly Kerr and exactly thermal, because the projection is exactly what the constitution says it is* — placed once, priced in advance, and re-exposed to every sharper instrument the decade delivers. What the module can never do is win by silence alone; what it can always do is die by one confirmed shot. That asymmetry is not a weakness. It is the falsification standard of the series, worn on the sleeve.

58 The falsifier timeline, dated

Table 8 closes the Part’s forward apparatus: the schedule-centric companion to the kill table below, one line per confrontation, with the verdict type each instrument can return — *kill* (can falsify outright), *support* (can strengthen, cannot confirm alone), *watch* (tension monitor). Every kill line is pre-declared and dial-free.

Table 8: The dated falsifier timeline of the black-hole module. “Kill” lines are pre-declared and dial-free; the claims-level register of each row is the kill table below, and the external hostages are the cohort falsifiers this module inherits by construction (§61.7). Under the geometry being right, every kill line returns silence: the module’s signature is the ever-sharper null.

Forced claim	Instrument(s)	Decision win-dow	Verdict type	Kill condition
Kerr shadow (null)	EHT, ngEHT, BHEX	permanent; sharpening early 2030s	kill	confirmed $> 5\%$ deviation (F-BH-1)
Kerr ringdown (null)	LVK IR1/O5, LISA, ET/CE	permanent; sharpening 2026–2040	kill	confirmed $> 1\%$ QNM deviation (F-BH-2)
No echoes	LVK, ET/CE	permanent	kill	reproducible post-ringdown echoes (F-BH-3)
Kerr multipoles	LISA EMRIs	mid-2030s	kill	non-Kerr multipole at 10^{-4} class (F-BH-10)
PBH ratio ladder (model-level)	microlensing, sub-solar GW, CTAO	2030s	kill (model) / support	populated mass function off the e^7 spacing (F-BH-4)
Trace thermality (structure)	analog platforms	continuous	support / watch	— (structure evidence only; the spectral kill F-BH-5 lives at true horizons)
Species-sum closure	dedicated computation	decidable at will	kill (internal)	sum misses $1/G_N$ (F-BH-9)
External hostages	HL-LHC; haloscopes; JUNO/-DUNE	2029–2041	kill (external)	lock cluster excluded; Kaxion window excluded; inverted ordering

Read the timeline as an experimentalist would, cadence first. The permanent lines are already running: the shadow and ringdown nulls sharpen continuously, with the next discrete jumps at ngEHT and BHEX in the early 2030s and the LVK runs before them. The one dated cliff is LISA’s: extreme-mass-ratio inspirals in the mid-2030s read the Kerr multipole structure at the 10^{-4} class, and F-BH-10 fires or it does not. Two lines need no telescope at all — the species-sum closure (F-BH-9) is decidable at a desk whenever the dedicated computation is run, and the external hostages (the lock cluster at HL-LHC, the Kaxion window in the haloscopes, the neutrino ordering at JUNO/DUNE) can retire this module from neighboring fields on their own schedules. Note also what the verdict column does *not* offer: no row returns “confirm.” Support lines strengthen the reading; only kill lines are decisive; and under the geometry being right, every kill line returns silence.




Figure 14: The kill-test timeline, drawn. Bands mark the decision windows of Table 8; every band is the same one-sided bet at increasing stakes, and the one red dot is the row the cohort itself controls — the internal species-sum computation, decidable without any instrument.

59 Refresh rules: the module’s data discipline

Four standing rules govern every number this Part and the evidence base carry; they are part of the module’s honesty apparatus, stated once so that no future edit can drift silently.

Rule 1 (measured values). Measured values follow the instruments’ latest published results and carry their own dates; the evidence base is re-dated at each version of record, and a “through year Y ” null advances only on a verified sweep of the literature.

Rule 2 (predictions). The GG-Theory prediction column never moves under refresh. Any change to a predicted value is a version-of-record event of the owning paper, never a data refresh of this one.

Rule 3 (kill rows are dial-free). A fired kill retires the module. No refresh, reinterpretation, or auxiliary hypothesis may be introduced after the fact to weaken a kill row; the rows fire as written or not at all.

Rule 4 (pulls). Printed pulls are recomputed programmatically from the printed prediction and the refreshed measurement at every refresh, never carried forward by hand.

Kill criteria

The kill table is the module’s contract, stated in the field’s own observables. Each row names the observable, the deviation threshold at which the row fires, and the timeline on which the firing could happen; the refresh rules of the preceding section apply to every row — they fire as written or not at all.

Table 9: Black-hole-specific kill criteria. Any one, if experimentally established, would retire the GG-Theory reading of black holes — except F-BH-4, the one model-level row, which retires the ladder conjecture and not the module. Ten criteria, and the last column is the honest one: two of them need instruments that do not exist, and two are theory-side questions with no instrument at all. The ones that can bite this decade are the shadow and the ringdown, and both campaigns are already running. A registry that listed only the reachable tests would look stronger and would be a different document.

Label	Observable	Deviation threshold	Timeline
F-BH-1	EHT shadow	$> 5\%$ deviation from Kerr on M87* or Sgr A*	Ongoing, ngEHT 2030s
F-BH-2	GW ringdown	$> 1\%$ QNM deviation from Kerr in LIGO/LISA sample	Ongoing
F-BH-3	GW echoes	Post-ringdown signal at $\Delta t \sim M \log M$	Ongoing analysis
F-BH-4	PBH ratio structure	A well-populated PBH mass function whose peak <i>spacing</i> is inconsistent with the e^7 scale ladder (model-level; a single detection does not trigger it)	Ongoing
F-BH-5	Hawking spectrum	Deviation from thermal form at observable frequencies	Near-Planck BHs required; unfeasible
F-BH-6	Information loss	Demonstration of $ \psi\rangle \rightarrow \rho_{\text{mixed}}$ in gravitational setting	Theoretical consistency; no direct test
F-BH-7	Bekenstein coefficient	$S_{\text{BH}} = cA/(4G_N\hbar)$ with $c \neq 1$	Theoretical; no direct test
F-BH-8	Firewall signature	High-energy quanta at horizon accumulating in stationary BH	Accessible only with ultrafine BH spectroscopy
F-BH-9	Species sum	The sign-resolved $\sum_i \chi_i$ over the GG-A9-SM spectrum with the Kaluza–Klein tower fails to reproduce the $1/G_N$ of Step 1 (Eq. (179); L-BH-Tower; internal)	Theory-side; decidable by a dedicated computation, no instrument required
F-BH-10	EMRI multipoles	Confirmed non-Kerr multipole structure at the 10^{-4} class in a LISA extreme-mass-ratio inspiral	LISA era, mid-2030s

The ten rows sort by reach, and reading them that way is how the table is meant to be digested. Four are live instrument rows on astrophysical data: the shadow (F-BH-1), the ringdown (F-BH-2), the echoes (F-BH-3), and the LISA multipoles (F-BH-10) — the rows the coming decade will keep squeezing. One is model-level by construction: the PBH spacing row (F-BH-4) fires only on a well-populated mass function off the e^7 ladder, never on a single detection. Four are registered but presently beyond instrumental reach — the Hawking-spectrum form (F-BH-5), demonstrated information loss (F-BH-6), a Bekenstein coefficient different from one (F-BH-7), and the firewall signature (F-BH-8) — kept on the books because the register requires the exposure to be stated even where no instrument yet exists. And one is internal: the species sum (F-BH-9), the only row whose firing date the theory itself controls. Nothing in the table waits on a definition; every threshold is a number.

Cumulative UCLA-Arisaka evidence base

The evidence base closes the Part, and its two halves should be read differently. The upper rows are the cohort’s cross-ledger campaigns — predictions made elsewhere in the program, confronted with 2025–2026 data — and they are the module’s external bond: the black-hole reading runs on the same bulk that produced them, so their fate is its fate. The lower rows are the black-hole benchmarks this paper adds to the cohort’s running account.

Table 10: Cumulative evidence for the GG-Theory program as of July 2026, including the black-hole qualitative consistency checks added by the present paper. The upper block belongs to the cohort and not to this paper, and is reprinted because a black-hole module quoting only its own consistency checks would be quoting a very short list. The lower block is what this paper adds, and every row of it says consistent with Kerr — a check passed, not a prediction confirmed, since general relativity says the same. That is why the two blocks are separated rather than summed.

Campaign / observable	GG prediction	Measurement (2025–2026)	Pull
H_0 (one background; C2, A10)	$H_0 = 66.6 \pm 0.2$ km/s/Mpc (two-road reading; 66.8 ± 0.2 from geometry alone)	67.4 ± 0.5 (CMB, the background read through the sound-horizon calibration); ladder 73.50 ± 0.81 , the predicted outlier — a calibration systematic, not derived	-1.5σ
W -mass (world avg)	$m_W = 80.3700$ GeV (A9, QD)	80.3692(133)	-0.06σ
Neutrino $\sum m_\nu$	≈ 0.059 eV (NH floor)	< 0.072 eV (DESI+CMB)	consistent
Higgs mass	125.20 GeV	125.20(11)	0.0σ
Di-Higgs rate	$\mu_{HH} = 1$	$0.9^{+1.4}_{-1.1}$	+0.08σ
CHSH (loophole-free)	$2\sqrt{2}$	2.697(15)	consistent
Born rule precision	exact	$\lesssim 10^{-13}$ for a_e	consistent
Arisaka Constant Q_B	0.1555 eV	ATP ~ 0.17 eV	consistent
<i>Black-hole benchmarks added by present paper:</i>			
M87* shadow	Kerr: 42 μ as	EHT 2019: 42 ± 3 μ as	consistent
Sgr A* shadow	Kerr: 51 μ as	EHT 2022: 51.8 ± 2.3 μ as	consistent
GW150914 ringdown	Kerr QNM	LIGO 2015: Kerr-consistent	consistent
GW250114 ringdown + area law	Kerr QNMs; $\Delta A > 0$	LVK 2025: overtone resolved; area theorem confirmed	consistent
Bekenstein entropy $A/(4G_N\hbar)$	Derived from scale profile	Consistent with merger area	consistent
Page curve	Shape forced; quantitative curve open	Not yet directly tested	consistent
Firewall	Not predicted	Never detected	consistent
GW echoes	Not predicted	None convincingly detected	consistent

Three features of the table deserve the reader’s attention. First, the campaign block is honest about its one visible strain: the H_0 row sits at -1.5σ against the microwave background, printed as

measured; the distance ladder’s 73.5 is not a second value of the rate but the outlier the framework predicts, a calibration systematic it names and does not derive (the earlier partial derivation of a local lift is withdrawn). Second, the black-hole block contains no number the module fitted: the shadows, the ringdowns, and the area theorem are pure-Kerr consistency checks, exactly what a framework whose signature is silence must show. Third, “consistent” is doing modest, clearly-priced work throughout: it strengthens the reading and confirms nothing — confirmation is not on offer in this Part, only survival.

The Part in one breath, and the door to the next. The record is one-sided: every shot fired at the module since 2015 has landed on *consistent*, and the module’s response is not celebration but escalation — each upgraded instrument makes the promised silence harder to keep. The physics intuition to carry forward: *a theory is worth exactly the risks it signs for*, and the tables above are its signatures. What remains is judgment: how this reading compares, framework by framework, with the alternatives that half a century of black-hole physics has put on the table — and what the bottom line is. That reckoning is the business of the next Part.

60 Part VIII in plain English: the record of every shot fired, and the dates the next ones are scheduled for

A prediction is worth what it risks, so this Part is the accounting.

The record first, and it is one-sided in a way that should be read carefully rather than proudly. Since 2015 every measurement that could have contradicted the module — two shadows, a decade of ringdowns, the area theorem confirmed in 2025 — has landed on consistent. That is not a run of successes. It is a run of tests the framework was always going to pass, and the right response is to say so and then post a harder bond.

The bond is the falsifier timeline, with dates. Ten named criteria, each with the instrument that would apply it and the window in which it will. Two of them need instruments that do not exist. Two are theory-side questions with no instrument at all, including the species-sum computation this paper leaves open. The rest run on the next-generation Event Horizon Telescope, on LISA in the middle 2030s, and on the ringdown catalog accumulating now.

The Part also states its refresh discipline, which is a small thing that matters more than it looks. Every number in these tables comes from a primary record with a date attached, and there is a written rule for when each is re-read. A comparison table that quietly ages is worse than no table at all, because it goes on looking current.

Part IX

Summary and Comparison

Scorecards, Framework Comparisons, and the Bottom Line

This Part is the fast reader’s exit and the referee’s desk: the whole paper weighed, in one place. It opens by stating plainly what **GG-Theory** claims about the black hole and what would falsify the claims; scores the classical puzzles one by one; sets the module against every leading framework — first in two tables, then in a section of arguments; and closes with the concept dictionary, the axioms actually used, and the bottom line in a single row. Everything comparative the paper has to say lives in this Part; everything forward-looking — the inventions, the open questions, the future directions — is deliberately kept for the next.

Synthesis: what **GG-Theory** claims about the black hole

Before the summary tables, let us state plainly what **GG-Theory** claims about the black hole, and what would falsify these claims.

The positive claims. The black-hole horizon in **GG-Theory** is a specific locus in 4D spacetime where the bulk degrees of freedom accessible to the external observer become highly compressed into compact-dimension and radial-interior sectors that are not directly observable. The **OGGEP** functor, applied at this locus, produces a CPTP channel whose output is the Hawking radiation spectrum. The entropy stored at the horizon is literally the entanglement entropy between the 4D external Hilbert space and the inaccessible bulk sector. The coefficient $1/4$ is *derived* from induced gravity (§31.5): the regulator-independent ratio $4\pi/16\pi$, with no free parameter. The Page curve’s shape follows from bulk unitarity plus CPTP projection; the quantitative curve is open. No firewall is required because the apparent monogamy violation of AMPS is an artifact of 4D-only thinking; at the 6D level, the “interior partner” is a bulk state shared with the entire evaporation history.

What this predicts. For astrophysical black holes, all observables (EHT shadow, LIGO ringdown, Hawking spectrum if it were observable) should match Kerr to extraordinary precision (better than parts in 10^{46} , and $\sim 10^{-66}$ for the EHT-imaged supermassive shadows), because the corrections are suppressed by $(\ell_{\text{lock}}/r_s)^2$ with $\ell_{\text{lock}} \equiv \hbar/(M_{\text{lock}}c) \sim 5 \times 10^{-20}$ m. For primordial black holes, the scale structure may imprint preferred mass *ratios* (a geometric ladder spaced by e^7); the absolute scale is set by the formation-horizon mass and is not yet derived (§49). The Bekenstein–Hawking entropy coefficient $1/4$ is derived (§31.5); the one residual task is the explicit species sum that fixes the induced G_N numerically (§31.10).

What would falsify the claims. Any of the following observations would retire the GDOS reading of black holes: a Kerr deviation at the $> 5\%$ level in EHT shadows; a QNM deviation at the $> 1\%$ level in LIGO/LISA ringdowns; a detection of gravitational-wave echoes from outside the horizon; a demonstrated information-loss-in-gravitation at any energy scale; a violation of the Bekenstein entropy coefficient via microscopic counting experiments.

What is theoretical vs. empirical. The information paradox itself is a theoretical consistency problem, not an observational one. No current or near-future experiment can directly test the Page curve in a real astrophysical black hole. The GDOS resolution is empirically consistent because it predicts Kerr + small corrections for all feasible observations, and it is theoretically consistent because it is grounded in a specific UV theory (GG-6) rather than an abstract duality (AdS/CFT).

What is new in GDOS. Compared to the existing resolutions (AdS/CFT, islands, ER=EPR, fuzzball), the GDOS reading brings a Standard-Model-compatible UV theory (GG-6 with SU(5) GUT); a specific, finite microstate count via $U(1)_{Q_i}$ holonomies on S^1_θ , with the coefficient fixed regulator-independently by induced gravity (and OGN-5-independent, Theorem 1); a Page curve whose shape is forced by unitarity plus projection (the quantitative curve open), with no replica wormholes and no quantum-extremal-surface extremization as separate postulates; a geometric reading of the “interior partner” as a bulk state, dissolving the AMPS monogamy paradox; and concrete Kerr corrections priced at the $(\ell_{\text{lock}}/r_s)^2$ level — no observable deviation for astrophysical black holes, with the falsifiable content carried by the *null* rather than by a positive signature. Each of these is developed as a claimed invention, with its provenance and its limits, in the Discussions Part.

The role of the present paper in the program. This paper is the black-hole companion (C3-BH) of the twenty-two-paper GG-Theory cohort and completes its gravitational branch. Its place in the program, and the honest division between what is established here and what is registered open, are the business of the Discussions Part.

Puzzle-by-puzzle scorecard

Table 11: Classical puzzles of black-hole physics and their resolution in **GG-Theory**. Read the last column first. Resolved and reframed are different claims: resolved means the puzzle does not arise in the six-dimensional reading, reframed means it does arise and is answered differently. Neither means the standard treatments were wrong about their own domain. The third column is included so a reader can see exactly which lift does the work in each row, and judge the row on that rather than on the verdict.

Puzzle	Standard question	6D lift	4D resolution	Status
Information paradox (Hawking 1976)	Pure $\rightarrow$ mixed; unitarity violated?	Bulk unitary; 4D CPTP projection	Radiation encodes info in bulk-to-boundary correlations	Resolved
Page curve	Why should S_R turn over at t_P ?	Bulk bipartition symmetric around t_P	Shape forced by bulk unitarity + partial trace; quantitative curve open	Reframed
Bekenstein–Hawking entropy	What states count?	$U(1)_{Q_i}$ holonomies on S_θ^1 at horizon	$S_{\text{BH}} = A/(4G_N\hbar)$, exact $4\pi/16\pi$ (induced gravity)	Resolved
Firewall paradox (AMPS 2012)	Monogamy vs no-drama	A partner is bulk, not independent 4D	No firewall forced; the monogamy premise fails at the 6D level	Resolved
ER=EPR (Maldacena–Susskind 2013)	Entangled BHs = wormholes?	Entanglement encoded in shared bulk sector	Automatic; no ad hoc ER=EPR postulate needed	Reframed
Islands formula (Penington et al. 2019)	How to get Page curve from gravity?	Island = bulk region traced by OGGEF	Formula recovered; derivation physically concrete	Reframed
Singularity	What’s at $r = 0$?	$r = 0$ is small-scale-profile region, not singular in 6D	Coordinate artifact; bulk is smooth	Resolved
No-hair theorem	Why only 3 parameters?	4D observables are zero-mode projection of 6D state	Hair is in bulk; 4D sees only (M, J, Q) zero modes	Resolved
Species problem	Why coefficient $1/4$ not $N/4$?	Induced gravity; no bare G_N	$1/4 = 4\pi/16\pi$, species-indep.	Coeff. exact
Why black holes radiate?	Quantum effects in curved ST	Horizon = OGGEF port = CPTP channel	Thermal output automatic	Resolved
de Sitter entropy	$S_{\text{dS}} = A_H/(4G_N\hbar)$?	Cosmological horizon = OGGEF port	Same mechanism as BH entropy	Resolved
Holographic principle	Why area-extensive entropy?	Bulk-to-boundary projection has A -scaled bandwidth	Holography automatic; $1/(4G_N\hbar)$ from scale profile	Resolved
Complementarity	Inconsistent exterior vs infalling views?	Different partial traces of same 6D state	No inconsistency; no firewall required	Resolved

Framework comparison

Table 12: Six frameworks for black-hole physics compared across twelve desiderata. Y = has the property, N = does not, $\sim$ = partial/ambiguous. No column here is clean, including the last one, and the two entries where this framework does not score — the microstate count and the reach of its quantitative predictions — are the two a specialist would ask about first. They are left where they fall in the row order rather than moved to the bottom. A framework built to count microstates in a controlled sector should be judged on that, which is why the comparison is by property and not by total.

Property	Classical GR	AdS/CFT	Fuzzball	Firewall	Islands	GG-Theory
Resolves info paradox	N	Y	Y	Y (partially)	Y	Y
Bulk unitarity preserved	N	Y	Y	N	Y	Y
Page curve derived	N	Y (limited)	$\sim$	N	Y (2D)	Y (shape, 4D; curve open)
No firewall needed	Y	Y	Y	N	Y	Y
Standard Model compatible	Y	N	$\sim$	$\sim$	$\sim$	Y
Real-universe applicable	Y	N	$\sim$	Y	$\sim$	Y
Smooth horizon (no drama) compatible	Y	Y	N (fuzzball replaces)	N	Y	Y
Singularity resolved	N	$\sim$	Y	N	$\sim$	Y
BH microstates counted	N	Y (specific cases)	Y	N	$\sim$	$\sim^\dagger$
Quantitative predictions beyond Kerr	$\sim$	Limited	Some	Some	Limited	Limited but specific
No free parameters	N	N	N	N	N	Y
UV theory specified	N	Y (string)	Y (string)	N	N	Y (GG-6)

These desiderata are necessarily schematic, and the table is best read as a map of explanatory roles rather than a scoreboard. On it GG-Theory scores well; the entry that remains in progress — the explicit numerical microstate enumeration ($\sim^\dagger$) — is flagged honestly in §31.10. († The coefficient $1/4$ and the identity of the microstates are settled for generic black holes in §31.5–31.6; the outstanding item is the numerical species sum, not the coefficient.)

Comparative landscape: what each framework takes as primary

The most useful comparison is not which framework “wins” but what each takes as primary, explains well, and finds difficult. Table 13 records this in deliberately neutral terms; the string-

theoretic entries in particular are *internally forced* by consistency, not arbitrary, and we present them as such.

Table 13: Comparative landscape of major black-hole information approaches. Schematic by design: the aim is to distinguish explanatory roles, not to force false exclusivity. The last column is the one worth reading, because it is where each approach is weakest by its own lights rather than by this paper’s, and every entry in it is drawn from that approach’s own literature. The table is schematic by design and the caption says so: these approaches are not exclusive, and a reader could hold several of these rows simultaneously with this one.

Approach	Primary ontology	Main strength	Main open issue
Pure Hawking semiclassicals	Quantum fields on a fixed curved background	Temperature and approximate thermal flux	Monotonic radiation entropy and information loss if taken literally (Hawking, 1976).
Complementarity	Observer-dependent but unitary horizon	Retains unitarity and a smooth horizon together	Hard to formulate once entanglement monogamy is imposed.
Fuzzballs / string microstates	Explicit microstate geometries or string states	Genuine microscopic entropy accounting in special (BPS) sectors (Strominger and Vafa, 1996)	Generality beyond controlled, supersymmetric sectors.
Island / QES program	Generalized entropy of the radiation subsystem	Quantitative Page-curve recovery in controlled models (Almheiri et al., 2019; Penington, 2020)	Deeper ontology of islands; extension to realistic, non-AdS settings.
GG-6 / GDOS / OGGEF	Closed 6D bulk plus open boundary projection	Exact $A/4G_N$ for <i>generic</i> black holes (§31.5); a universal reason generalized entropy and branch switching arise	Numerical species closure of the induced G_N (§31.10); benchmarking against detailed island computations.

The entropy problem, head to head. The sharpest single test is the one this paper targets: reproduce $S = A/4G_N$, *with the exact 1/4, by counting*. Stated neutrally: Strominger–Vafa achieves the exact coefficient by an explicit D-brane state count — but only for supersymmetric, extremal or near-extremal black holes (Strominger and Vafa, 1996). Loop quantum gravity reproduces the area law from spin-network states, with the 1/4 fixed by tuning the Barbero–Immirzi parameter. Fuzzballs count horizonless microstate geometries, sharpest in controlled (largely supersymmetric) sectors. The island/QES program recovers the *Page curve* in controlled models, but does not by itself supply a first-principles microstate count for a real Schwarzschild black hole. And GG-Theory obtains the exact $1/4 = 4\pi/16\pi$ from induced gravity for the *generic, non-supersymmetric* black hole (§31.5), with the microstates identified as the hidden-fiber holonomies (§31.6) — the outstanding item the numerical species sum, not the coefficient. The honest verdict is therefore narrow and strong: on *generality* GG-Theory reaches the case the explicit string count cannot, and it fixes the coefficient without tuning; on *explicit enumeration* the string result remains the more complete in its own supersymmetric domain. The two are best read as complementary — one maximal in a protected sector, the other general but still closing its species sum.

61 Comparison with the leading frameworks: the deeper analysis

The tables above give the map; this section gives the arguments. We take the leading alternative frameworks in turn — classical general relativity, the AdS/CFT correspondence, the fuzzball program, the firewall proposal, and the islands formula — and close with what the runtime taxonomy adds that none of them has. Everything comparative the paper has to say lives in this Part; what the module *adds* is deliberately kept for the Discussions.

61.1 Classical general relativity

Classical GR is the starting point. It gives us the horizons, the singularity, the no-hair theorem, and the Bekenstein–Hawking entropy as an analogy to thermodynamics. But it does not explain what the horizon really is, what the microstates are, or why the entropy is area-extensive. The information paradox is, in classical GR, simply a statement of breakdown: quantum mechanics demands something that classical GR cannot provide.

GDOS augments classical GR with the bulk-to-boundary projection structure. The classical 4D metric is the zero-mode projection of the 6D bulk. All of classical GR is recovered in the low-energy limit. The Hawking radiation, the entropy, and the Page curve emerge as consequences of the projection rather than as postulates added on top of classical GR.

61.2 The fuzzball program

The fuzzball program, developed primarily by Samir Mathur and collaborators since 2003, proposes that the interior of a black hole is not empty but rather a quantum superposition of horizonless geometries. In this picture, there is *no horizon*: the region inside the would-be Schwarzschild radius is a complex quantum-gravitational state that differs from vacuum by order-unity amounts at horizon scale.

The fuzzball proposal resolves the information paradox by eliminating the horizon: if there is no horizon, there is no Hawking radiation in the standard sense, and no information loss. But this requires that external observations — the shape of the black-hole shadow, the ringdown spectrum — be dominated by order-unity quantum-gravity effects near what would classically be the horizon. As of 2026, EHT and LIGO/LISA observations are consistent with Kerr to high precision, which strongly constrains fuzzball proposals: they must reproduce Kerr behavior at macroscopic scales to extraordinary precision, which is hard to motivate from string-theory fundamentals.

The strengths of the fuzzball program include explicit microstate geometries (in specific string-theory backgrounds), a natural resolution of the information paradox (no horizon means no horizon-based information loss), and a concrete mechanism by which microstates differ from one another. The weaknesses include: (i) horizonless geometries are hard to reconcile with high-precision EHT and LIGO observations consistent with Kerr; (ii) fuzzball constructions work best in highly supersymmetric backgrounds not directly relevant to real black holes; (iii) the physical mechanism by which the classical horizon emerges from the fuzzball superposition is unclear.

GDOS takes a different position: there *is* a horizon, and it is smooth at the 6D bulk level. The microstates are not distinct horizonless geometries but distinct bulk configurations with the same

4D zero-mode projection, so observational agreement with Kerr is automatic at the 4D level, to all feasible precision. Fuzzballs, by contrast, must either reproduce Kerr exactly — in which case they add no testable content — or predict small deviations, in which case they risk falsification as EHT precision improves. Information, here, is preserved via the OGGEF CPTP channel, not via horizonless microstate geometries.

61.3 AdS/CFT correspondence

The AdS/CFT correspondence is a specific duality between Type IIB string theory on $\text{AdS}_5 \times S^5$ and $\mathcal{N} = 4$ super-Yang-Mills in 4D Minkowski space. It is a genuine duality — each side is a complete quantum theory, and physical quantities are translated via a precise dictionary. In this framework, black-hole evaporation is manifestly unitary because the boundary CFT is unitary.

AdS/CFT provides a deep and rigorous setting for studying quantum gravity in anti-de-Sitter space. Its limitations for real-universe physics are, however, substantial. Our universe is de Sitter, not anti-de Sitter; the actual gauge group is not $\mathcal{N} = 4$ SYM; and the bulk theory is not type IIB string. Asymptotically-flat black holes in 4D Minkowski space do not have a direct AdS/CFT dual — their treatment requires cavity or brane constructions that are not rigorous. Put plainly: AdS/CFT needs the box and the dual, and our universe declines to provide either. What is missing is an anchor for the boundary description that does not depend on asymptotics; GG-Theory’s anchor is the *observer* — the port is defined by causal access, not by a wall at infinity — which is why the machinery moves to flat space and de Sitter without modification.

GDOS occupies a complementary position: it is less rigorous than AdS/CFT (it is not a proved duality but a proposed constitutional framework), but it is *physically more relevant* to our universe because its bulk content is the Standard Model. GG-Theory and AdS/CFT should be thought of as two holographic frameworks at different levels of formalization and different domains of applicability. The bulk-to-boundary projection OGGEF is the analogue of the AdS/CFT dictionary, defined for the physical universe.

61.4 Firewall proposal

The firewall proposal (AMPS 2012) sacrificed horizon smoothness to preserve unitarity plus Bekenstein–Hawking entropy. Infalling observers, in this picture, encounter a wall of high-energy quanta at the horizon.

The firewall proposal is consistent with unitarity but violates the equivalence principle of general relativity. It has few defenders as of 2026; the islands formula has superseded it as the preferred resolution in the AdS/CFT community.

GDOS agrees with the post-firewall consensus that the horizon is smooth, but provides a different mechanism: the apparent monogamy violation dissolves at the 6D bulk level, so no firewall is needed.

61.5 The islands formula in GDOS

The islands formula, established in 2019–2020, is currently the most widely-accepted resolution of the information paradox. It reproduces the Page curve in two-dimensional JT gravity via the

quantum-extremal-surface prescription combined with replica wormhole saddles in the Euclidean gravitational path integral.

The strengths of the islands framework are its mathematical precision (in 2D) and its direct contact with AdS/CFT. The weaknesses are: (i) it is rigorously established only in 2D, and its extension to 4D is partial; (ii) the replica wormhole mechanism is non-perturbative and lacks a clean physical interpretation; (iii) the extremization prescription is an added computational ingredient rather than a primary physical principle.

GDOS reproduces the islands formula in the appropriate limit but does not require it as a separate postulate. The bulk partial-trace structure of OGGEF produces the same entanglement entropy as the islands formula, and the quantum-extremal-surface is identified with the geometric boundary between the traced and un-traced bulk regions. The GDOS version has the advantage that the ingredients (unitary bulk evolution plus CPTP projection) are physically concrete, rather than formal path-integral manipulations.

The islands formula (9) is a precise computational rule for the entanglement entropy in two-dimensional JT gravity coupled to matter. In GDOS, the analogous statement is that the 4D entanglement entropy is

$$S(R) = S_{\text{vN}}(\rho_{\partial}^R) = \min_{\Sigma} \left[\frac{\text{Area}(\Sigma)}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I_{\Sigma}) \right], \quad (73)$$

where Σ is the quantum extremal surface in the bulk, I_{Σ} is the island bounded by Σ , and the minimization is over all admissible Σ .

In the JT limit (two-dimensional bulk with specific coupling to matter), this reduces to the islands formula of Penington and Almheiri et al.

61.6 ER=EPR in GDOS

The ER=EPR proposal of Maldacena and Susskind (2013) asserted that two entangled black holes are connected by a wormhole. In GDOS, this is automatic: two black holes whose interior degrees of freedom are entangled in the bulk share a common bulk sector, which geometrically is a wormhole-like connection through the bulk compact dimensions.

The GDOS picture makes ER=EPR computationally tractable: given two bulk regions B_1, B_2 with entangled bulk states, the OGGEF projection to the 4D exteriors produces two 4D black holes whose radiation is correlated. The ER bridge is the shared bulk sector; the EPR correlations are the resulting 4D radiation correlations.

61.7 What the runtime taxonomy adds that no leading framework has

The comparisons above were drawn at the level of the information question, where every leading framework has an answer of some kind. The runtime classification developed in this paper (§42, §42.3) permits a sharper statement of what is unique here, because it operates at a level the other frameworks do not occupy at all: the *dynamical classification* of black-hole events. Four deliverables — collected in Table 14 — have no counterpart in any leading approach.

(i) A two-regime taxonomy with theorems behind it. Formation, merger, accretion, evaporation, the Hawking–Page transition, and the Page turnover are sorted by a single question — *does the landscape change?* — into capture, coalescence, traffic, channel exit, the one true GRIP-Out, and a milestone, respectively. Each label is backed by a proved theorem (Lyapunov descent, Out-necessity, No-Exit (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A11-Origin)), not by a metaphor. No other framework classifies black-hole dynamics at all; they classify *states* (microstates, fuzzball geometries, code subspaces) and remain silent about which events reorganize the problem and which merely evolve it.

(ii) The conditional-ground-state statement, with an expiration date. That every astrophysical black hole is today a *conditional* ground state — exitless only because $T_H < T_{\text{CMB}}$, with the condition expiring on the calculable schedule $M_{\text{open}}(t) \propto a(t)$ (§42.2) — is a new *kind* of statement about black holes: their thermodynamic finality is boundary-condition-relative and cosmology-scheduled. The ingredients (Hawking temperature, CMB cooling) are textbook; the statement is unwritable without the open-system vocabulary, which is why it does not appear in the fuzzball, AdS/CFT, or islands literatures.

(iii) Gravitational-wave observables tied to the same functional as particle physics. The ringdown spectrum is the Hessian of the free-energy ledger at the Kerr attractor (§48) — the same functional whose minima, on other strata, are the electron, the proton, and the dark-matter candidate. Islands and AdS/CFT say nothing about LIGO; fuzzballs predict horizon-scale deviations that EHT and ringdown data increasingly disfavor. Here the observed Kerr-ness of ringdowns is not an embarrassment to be explained away but a *confirmation*: clean basin curvature is exactly what a deep, simple attractor should show.

(iv) Interlocking falsifiability. The machinery that resolves the paradox here was not built for black holes: the same runtime, with the same theorems, is the engine that selects the electroweak vacuum, stabilizes the electron (Arisaka, 2026, P3-Electron), runs decoherence and measurement, and predicts the dark-matter window of the companion papers. The account is therefore falsifiable *from outside black-hole physics*: empty the 1.59–2.13 GHz Kaxion window at DFSZ depth or find a *second* line inside it, exclude the 3.85 TeV resonance cluster, or measure an inverted neutrino ordering, and this paper’s reading of black holes dies with them — an inverted ordering, measured in Illinois, would retire a theorem about horizons. (The dark-matter signature is one line, standing alone. An earlier reading of the cohort predicted a twin-peak doublet at $\Delta\nu/\nu = 1/35$; the registered mode count of the anomaly companion leaves one light angle, Arisaka, 2026, C1-Kaxion withdraws the doublet, and its absence is now part of the prediction, not a kill.) No black-hole-only framework offers hostages of that kind; a framework that cannot be shot from a neighboring field is, by the falsification standard of this series, the weaker for it.

Table 14: The leading frameworks against the runtime criteria. “Dyn. taxonomy” = a classification of black-hole *events* (not just states); “cross-domain engine” = the same proved machinery operating outside gravity; “external falsifiers” = kill criteria located in other fields.

Framework	Real verse?	uni-	Info answer	Dyn. taxon- omy	Cross- domain engine	External falsifiers
Classical GR	yes		none (paradox)	no	no	no
Fuzzball	in principle		no horizon	no	no	no
AdS/CFT	AdS only		duality	no	no	no
Islands / QES	toy models		Page curve	no	no	no
GG-Theory runtime	yes		OGGEP channel (Parts IV–VI)	yes (two regimes)	yes (A6–A11)	yes (one Kaxion line; 3.85 TeV; ν -ordering)

Concept dictionary: 4D to 6D for black holes

For the canonical Dictionary of program-defined acronyms (45 entries in ten categories), see GG-A2-GG6 Appendix B. The table below translates only the black-hole-specific $4D \leftrightarrow 6D$ vocabulary.

Table 15: Dictionary translating 4D black-hole concepts to their 6D bulk origin in GG-Theory. A dictionary is a claim about translation rather than about physics, and the third column is where the claim can be checked. Each row says which bulk object the four-dimensional concept is a projection of, and by what operation; a row whose mechanism is vague is a row where the translation has not yet been earned. Nothing is derived in the table — every entry points at the section that derives it.

4D BH concept	6D bulk origin	Mechanism
Event horizon	OGGEP port (one-way causal boundary)	Locus separating accessible from inaccessible bulk
Hawking radiation	CPTP channel output	Partial trace over bulk interior
Hawking temperature T_H	Scale profile at horizon	$k_B T_H = \hbar \kappa / (2\pi c)$ automatic
Bekenstein–Hawking entropy	von Neumann entropy of $\rho_\partial^{\text{out}}$	Entanglement of 4D exterior with bulk
Microstate count e^{S_{BH}/k_B}	Phase-space volume of $U(1)_{Q_i}$ holonomies on S_θ^1	One per horizon pixel of area $4\ell_P^2$
Page curve	Bipartite pure-state entanglement entropy	Shape forced by bulk unitarity; curve open
Singularity at $r = 0$	Small-scale-profile region of 6D bulk	Coordinate artifact in 4D projection
No-hair theorem	Zero-mode projection of bulk state	Higher-mode content traced over
Holographic principle	OGGEP bandwidth is area-extensive	Fundamental structural property
de Sitter horizon	Cosmological OGGEP port	Same mechanism as BH horizon
Interior partner (AMPS)	Bulk degrees of freedom, not 4D subsystem	Dissolves monogamy puzzle
Quantum extremal surface	Boundary of OGGEP-traced region	Geometric locus of partial trace
Replica wormhole	Bulk path-integral saddle connecting replicas	Natural in GG-6 path integral
Island I	Bulk region traced by OGGEP at given time	Physical realization of formal island

Program axioms used

Table 16: GG-Theory axioms and postulates used in the black-hole derivations. Seven entries, and the point of listing them is to bound what the derivations are allowed to use. The last column says where each is actually spent, so a reader who doubts a result can find the premise it rests on without leaving the paper. Nothing here is derived; it is the list of what the derivation may assume, and the exact 1/4 uses fewer of them than the row order suggests.

Label	Name	Statement	Role in BH physics
AX-1	Local causality	Forward-time causation in bulk	Horizon is a forward-causal boundary
AX-2	Primitive connection	Gauge-gravity composite connection on M_6	Bulk metric includes gauge sector
AX-3	OGGEP	Bulk $\rightarrow$ boundary functor	Horizon is a port for $\mathcal{D}$; Hawking radiation is output
AX-4	MSMF	Max sym, min freedom	Forces Kerr on boundary; no freedom for firewalls
POS-CP	Complete positivity	CPTP structure preserved	Hawking channel is CPTP
POS-OGN	OGN-5	Integer ledger (1, 3, 5; 7, 35)	Field content and induced- G_N species sum; spectrum structure (the 1/4 itself is OGN-5-independent, Thm 1)
POS-RG	Scale consistency	Commutes with coarse-graining	Area law at all scales

Bottom-line resolution

Table 17: Single-row summary of the 1975–2026 information-paradox debate. Two columns, one question per row, and the differences between them are the whole of the information paradox. Read down, the bulk column answers no loss and no firewall throughout; the boundary column agrees, with a parenthesis each time saying why the naive reading does not. The one row where the two genuinely differ is the singularity, and that difference is a statement about coordinates rather than about physics.

Question	6D Bulk	4D Boundary
Does info get lost in evaporation?	NO	NO (appears so at naive Hawking level)
Is the evaporation unitary?	YES	YES (via bulk-to-boundary correlations)
Does the Page curve hold?	YES (shape; curve open)	YES (shape; curve open)
Are there firewalls?	NO (none forced)	NO (none forced)
Is the horizon smooth?	YES in the canonical bulk	YES (4D-apparent); nothing forces otherwise
Is there a singularity?	NO (6D-smooth)	YES (4D-coordinate)
Is information stored on boundary?	YES (holography)	YES (Bekenstein area law)
<i>The bulk is unitary and smooth; the 4D boundary appears paradoxical because the 4D observer is seeing only a partial trace of the 6D state.</i>		

The Part in one breath, and the door to the next. Weighed against half a century of frameworks, the module’s case is generality without dials: it reaches the astrophysical hole the protected counts cannot, recovers the shape of the islands result without its auxiliary machinery, forces no firewall on the horizon while keeping unitarity, and prices its one open computation as a registered falsifier rather than a footnote. The physics intuition to carry forward: *where others must engineer protection, a forced trace and a fully induced G_N make the results structural*. The comparison says where the module stands. What it *adds* — the inventions, what remains honestly open, and the road ahead — is the business of the next Part.

62 Part IX in plain English: how the reading stands against five rivals, including where it loses

Six frameworks, twelve properties, and no column that scores on all of them.

The comparison is arranged by property rather than by winner, because the frameworks were built for different jobs. String microstate counting produces genuine microscopic entropy in supersymmetric sectors and does not reach an astrophysical hole; that is a limitation of scope rather than a failure, and the table says which. The island program recovers the Page curve quantitatively in controlled models and leaves the ontology of an island open. AdS/CFT is exact, and is about a universe that is not this one.

Where this reading claims an advantage it is generality: the coefficient is derived for generic holes with no protected sector, the Page curve follows from a ledger rather than from auxiliary machinery, and the horizon stays smooth without giving up unitarity. Where it does not score, it does not score — and two of those places are the first a specialist would ask about. The microstate count is partial. The quantitative predictions beyond Kerr are limited and specific rather than broad.

The honest summary is that this framework buys generality and pays for it in sharpness. Whether that is the right trade is a judgment, and this Part sets out the terms rather than announcing the answer.

Part X

Discussions

Inventions, Open Questions, and Future Directions

The judgment Part behind us, this Part is the reflective one: what is genuinely new here, where it sits in fifty years of literature, what role it plays in the wider program, what remains honestly open, and how the module can be stressed from the theory side alone. Nothing in this Part re-argues the comparisons of Part IX; its business is the module’s own ledger — the inventions on one page, the debts on the facing page — and the directions in which the work continues.

63 What is new

The preceding Parts together establish that the black-hole information paradox admits a complete resolution within the GG-Theory/GDOS framework. This is not the first claim of resolution in the history of the subject; the islands formula of Penington and Almheiri et al. in 2019 already gave a concrete derivation of the Page curve within two-dimensional JT gravity, and the AdS/CFT correspondence has for 25 years provided a setting in which black-hole evaporation is manifestly unitary. What, then, does GG-Theory add that was not already present?

The claim of this paper is that GG-Theory brings four specific novelties beyond the existing literature.

Novelty 1: a Standard-Model-compatible UV completion. Every existing resolution of the information paradox is formulated in a setting that is *not* the physical universe. The Strominger-Vafa microstate count works for extremal black holes in specific string-theory backgrounds. The AdS/CFT framework operates in anti-de-Sitter space with $\mathcal{N} = 4$ super-Yang-Mills boundary, neither of which describes our universe. The islands formula is proved rigorously only in two-dimensional Jackiw-Teitelboim gravity coupled to matter, not in four-dimensional asymptotically-flat Schwarzschild. Fuzzball proposals require specific string-theoretic structures not yet mapped to Standard-Model gauge content.

GG-Theory is different in this respect: the 6D bulk $M_4 \times I_r \times S^1_\theta$ carries the SU(5) GUT group plus the two new Abelian symmetries $U(1)_B$ and $U(1)_{Qi}$, and this content reduces on the 4D boundary to exactly the Standard Model (as established in the constitutional papers GG-A1-Constitution–GG-A4-Euclid and the Standard-Model paper GG-A9-SM). The black-hole physics developed in the present paper therefore takes place in a UV theory that *is* the theory of our universe. The resolution is not a proof of principle in a toy model; it is a claim about how real astrophysical black holes work.

Novelty 2: the exact Bekenstein–Hawking coefficient from induced gravity. The coefficient $1/4$ has resisted first-principles derivation for over fifty years. Strominger–Vafa obtained it

for supersymmetric, extremal black holes in string theory; the loop-quantum-gravity community recovered it at the cost of fitting the Barbero–Immirzi parameter; the fuzzball program aims at it through special microstate geometries. **GG-Theory** derives it differently, and for the *generic* black hole: because **GGEP** leaves no bare four-dimensional Newton constant, the horizon entanglement entropy *is* the induced $A/4G_N$, and the conical replica fixes the coefficient as the regulator-independent ratio $1/4 = 4\pi/16\pi$ (Theorem 1), with no supersymmetry and no extremality. Where the string-theoretic count needs symmetry protection to keep the number of states rigid, here the coefficient is not a count at all but a ratio of two normalizations fixed upstream by non-negotiable physics; a ratio cannot drift. The microstates are the hidden-fiber holonomies traced at the horizon (§31.6); the species sum that fixes the induced G_N numerically closes *structurally* through the Kaluza–Klein tower (Appendix R), with the precise leading coefficient the one item still outstanding. The coefficient itself carries no adjustable parameter.

Novelty 3: the Page curve’s shape from bulk unitarity alone, with no auxiliary machinery.

The islands formula derivation requires evaluating the gravitational path integral with replica boundary conditions, identifying saddles that include wormhole topologies connecting the replicas, and extremizing over quantum extremal surfaces. These steps are mathematically non-trivial and are rigorously justified only in specific low-dimensional models.

In GDOS, the Page curve’s shape is forced. Bulk unitarity (AX-1; the custody GG-A5-GDOS states) guarantees that the total bulk state remains pure. The OGGEF functor (AX-3) is a CPTP channel whose output is the 4D radiation. For a bipartite pure state, the entanglement entropy of either side is a function of the bipartition only; it follows the Page curve by Page’s 1993 theorem applied to a typical bulk state. No replica wormholes, no QES extremization, no island formula as a separate postulate. The ingredients reduce to (i) unitary bulk evolution and (ii) partial trace at the horizon.

Novelty 4: a geometric dissolution of the firewall paradox. The AMPS firewall argument runs as follows: for a late-time Hawking mode B , unitarity requires entanglement with early radiation R ; smooth-horizon physics requires entanglement with an interior partner A ; entanglement monogamy forbids both simultaneously; therefore either unitarity fails, smoothness fails, or some postulate of complementarity fails. The standard options — firewalls, ER=EPR, non-locality — each have their own difficulties.

GDOS resolves this by identifying the “interior partner A ” not as an independent 4D subsystem but as a bulk degree of freedom that is part of the OGGEF-traced region. Since the bulk is unitary and the boundary is CPTP, the apparent tripartite decomposition $A \otimes B \otimes R$ is an artifact of the 4D projection; at the 6D level there is no monogamy violation because A and R are not independent subsystems. The infalling observer is not forced to meet a firewall, because the 6D geometry is smooth at the horizon and the monogamy argument that demanded one has no independent interior partner to bind; the apparent discontinuity is a 4D coordinate artifact, resolved by the same mechanism that resolved the original Schwarzschild “singularity” at $r = r_s$ via Finkelstein’s 1958 coordinates.

64 Related work and precursors

Because this paper will be read against five decades of black-hole literature, we state plainly where it sits. The thermodynamic foundation is Bekenstein’s proposal that horizon area is physical entropy (Bekenstein, 1973), Hawking’s derivation of thermal radiation (Hawking, 1975), the Bardeen–Carter–Hawking four laws (Bardeen et al., 1973), and Page’s unitarity curve (Page, 1993). Our derivation of the coefficient belongs squarely to the *induced-gravity* lineage — Sakharov’s hypothesis that the Einstein–Hilbert term is generated by matter loops (Adler, 1982; Sakharov, 1968; Visser, 2002), its black-hole specialization by Jacobson (Frolov and Fursaev, 1997; Jacobson, 1994), and the Susskind–Uglum / Callan–Wilczek / Demers–Lafrance–Myers identification of the horizon entanglement entropy with the induced $1/4G_N$ (Callan and Wilczek, 1994; Demers et al., 1995; Susskind and Uglum, 1994).

What GG-Theory adds to that lineage is constitutional: GGEP makes the four-dimensional Newton constant *fully* induced (no bare term to split against), which removes the chief ambiguity of the Sakharov program and makes the entanglement entropy the entropy itself. The projection half of the story belongs to the holographic lineage — ’t Hooft and Susskind’s holographic principle (Susskind, 1995; ’t Hooft, 1993) and Maldacena’s AdS/CFT (Maldacena, 1998) — of which OGGE is the open-system, real-universe generalization; the islands program (Almheiri et al., 2019; Penington, 2020) is recovered as the GFF ledger at the horizon stratum (§33), and Wald’s Noether-charge entropy (Wald, 1993) and Jacobson’s “Einstein equation as equation of state” (Jacobson, 1995) are complementary geometric readings of the same area law.

The information-theoretic anchors are the AMPS firewall (Almheiri et al., 2013), Hayden–Preskill scrambling (Hayden and Preskill, 2007), and the log-area corrections of Carlip and others (Carlip, 2000; Solodukhin, 2011), all of which we engage below. The honest one-line placement: GG-Theory is induced-gravity for the coefficient and holographic for the projection, set in the Standard-Model-containing universe, with the generic (non-supersymmetric) exact $1/4$ as its distinctive new result.

65 The role of GG-Theory in the broader program

This paper is the black-hole companion (C3-BH) of the twenty-two-paper GG-Theory cohort. The preceding papers established the foundations of the theory (GG-A1-Constitution–GG-A4-Euclid; Arisaka, 2026, A4-Euclid), its quantum-mechanical reading (GG-A8-QM), its open-system/thermodynamic reading (GG-A5-GDOS–GG-A6-GRIP), and its precision tests at specific observables (Hubble tension, W mass, neutrino masses, Higgs). The present paper extends the program into the gravitational domain by addressing what is arguably the most celebrated problem in theoretical physics.

The significance of the black-hole reading in the broader program is threefold.

First, it provides a non-trivial consistency check of the theory in a regime far from its initial motivation. The GG-Theory framework was developed primarily to address puzzles in particle physics (the Hubble tension, the W mass, the Higgs sector, the neutrino masses), and only secondarily to address conceptual issues in quantum theory and thermodynamics. That the

same framework, applied to black holes, forces the shape of the Page curve and removes the premise of the firewall argument without additional input, is non-trivial. It suggests that the underlying geometric structure is capturing something real.

Second, it unifies disparate observables under a single framework. The Bekenstein–Hawking entropy, the Hubble tension, the Higgs mass, and the W -boson mass are, in the standard reading, unrelated phenomena. In GG-Theory they are all manifestations of the same 6D bulk geometry, seen through different 4D projections. The coherence of the predictions across scales — from the Higgs at 125 GeV to the one derived Hubble rate near 67 km/s/Mpc to the Bekenstein entropy of a $M_\odot$ black hole at $10^{77}k_B$ — is itself a piece of evidence for the framework.

Third, it opens testable predictions at observational frontiers. As the Event Horizon Telescope and LIGO/LISA continue to improve, and as new gravitational-wave detectors (Einstein Telescope, Cosmic Explorer) come online in the 2030s, the horizon-scale physics of black holes will become accessible at unprecedented precision. The GG-Theory predictions of Kerr-to- 10^{-46} -or-better for astrophysical black holes (Kerr-to- 10^{-66} for the EHT supermassive shadows) constitute one-sided falsifiers: the framework predicts exact Kerr far beyond any foreseeable sensitivity, so any observed deviation, at any accessible precision, retires it. The coming two decades will tighten this null test by orders of magnitude.

66 What the module explains, what remains open, and future directions

A scientifically useful synthesis must be honest about both strengths and limits. The present module explains well why pure-to-mixed exterior behavior is unsurprising once the exterior state is recognized as a projected open subsystem; why generalized entropy should have the area-plus-bulk-entanglement form (§33); why naive interior/exterior factorization is too crude in gravity; why a Page-like turnover arises through branch competition rather than miraculous late-time repair (§36); and why black holes are the cleanest testbeds for a bulk-deterministic, boundary-probabilistic theory.

What remains open is equally important, and we state it plainly. The induced-gravity derivation fixes the *exact* $1/4$ coefficient, the area law, and the identity of the microstates; what is not yet complete is the closed species-by-species sum that recovers the induced G_N *numerically* for the full GG-Theory spectrum — the explicit benchmark against the most controlled string-theoretic counts (Strominger and Vafa, 1996). GG-Theory does not yet reproduce, model by model, the detailed outputs of every replica-wormhole or island computation (Almheiri et al., 2019; Penington, 2020). It does not yet offer sharp horizon-scale observational signatures beyond the expectation that smooth semiclassical exteriors survive until regimes where independent evidence indicates failure. And at present it interprets the structure of controlled semiclassical calculations as much as it supersedes their technical machinery. These limits are not fatal; they define the next stage of work: if the framework is more than a reinterpretation, it must demonstrate explicit quantitative reproduction

of generalized-entropy calculations in tractable models, and derive distinctive corrections that competing frameworks do not naturally imply.

The specific work orders follow.

Quantitative calculations. While the structural claims of this paper — horizon as port, radiation as CPTP channel output, entropy as entanglement entropy — are developed in detail, the explicit quantitative predictions (KK corrections to the Hawking spectrum, precise PBH mass spectrum peaks, higher-order graybody factors) require further computation. Several of these are in progress.

The curved-space existence residue, adopted. The cosmology companion’s observer-vacuum reading closes the rotating and collapse cases at model level — entries R14–R15 of the GG-A8-QM/GG-P4-Real real-bulk series, of which Theorems 3 and 4 are this paper’s instances — and leaves one item conjectural: the exact curved-space existence theory, which it imports from the field rather than owns (Arisaka, 2026, A10-Cosmos). Its natural home is this paper, where the controlled model class is in hand. We adopt it here as a named open problem of the black-hole module.

Full 4D asymptotically-flat case. Most of the rigorous results on the islands formula are for 2D JT gravity. The full 4D asymptotically-flat Schwarzschild case remains less well-studied. GG-Theory proposes that the 4D case follows directly from the bulk structure; working out the explicit 4D computation and comparing with numerical- relativity simulations of Hawking evaporation is an open project.

The gravitational sector beyond black holes. The full GG-Theory reading of general relativity — non-black-hole spacetimes, gravitational dynamics beyond the horizon sector — lies beyond this paper’s scope and is registered here as open territory rather than promised work. The cosmological applications (inflation, dark energy, structure formation) are developed in the cosmology companion (Arisaka, 2026, A10-Cosmos) — a cohort companion drawing on the same GG-A1-Constitution–GG-A9-SM inheritance as the present paper, not a downstream consequence of it.

The generalized second law beyond the exchange class. Lemma 2.2 derives the second law’s direction and computes its production exactly in the Gaussian exchange class, with the linearized induced first law as the area price. The promotion to a Wall-class statement — arbitrary horizon cuts at the generality of Wall’s theorem (free and superrenormalizable fields, linearized backreaction) and beyond it to the interacting case (Wall, 2012) — is registered open: it requires half-sided modular machinery and ANEC-grade input that this module imports rather than owns. The registered path runs through the same two theorems the exchange class used (the KMS face, T-BH-5; the induced first law), which is what makes the promotion a sharpening rather than a rebuild.

Cross-checks with numerical gravity. The GG-Theory predictions for ringdown spectra, shadow geometries, and post-merger relaxation are, in principle, testable against numerical-relativity simulations. Work on embedding the GG-Theory corrections in numerical-relativity codes is a natural next step.

Analog gravity tests. Tabletop analog-gravity systems (acoustic black holes, Bose-Einstein condensates, optical black holes) provide controllable settings in which to test the structural features of GG-Theory. Specifically, the partial-trace structure of the horizon should produce characteristic entanglement statistics in analog Hawking radiation, which are experimentally accessible.

67 Falsifiability and kill criteria

The observational kill table is the business of Part VIII (Table 9); what belongs here is the *theory-side* exposure — the ways the module can be stressed with no instrument at all. The module would be weakened or falsified if a gauge-consistent proof established that naive factorization across the horizon is exact in the relevant quantum-gravitational regime, with no need for constrained-subsystem reconstruction; if controlled Page-curve calculations systematically resisted reinterpretation as a branch-minimized GFF ledger; if monotonic thermal entropy of the accessible radiation were shown exact all the way to evaporation while preserving global unitarity, with no broader subsystem semantics; or if the induced-gravity species sum failed to converge to the G_N that appears in the 4D Einstein equations — the direct internal kill criterion for the entropy claim.

Table 18: A compact falsification dictionary for the black-hole module. Five probe classes, each with what the module expects and what would stress it, and the middle column is deliberately weaker than a prediction. The module is minimal by construction, so most of what it says about a probe class is that ordinary physics should hold there; the last column is where it is exposed. A dictionary of this kind is worth less than a numbered falsifier, and it is here because not every stress test reduces to one.

Probe class	What the minimal module expects	What would stress it
Classical exterior geometry	Ordinary GR-like ringdown and exterior relaxation in semiclassical regimes	Robust evidence that smooth exteriors fail generically, far from any endpoint or UV scale, without an open-system explanation.
Horizon bookkeeping	Area-plus-bulk-entanglement entropy controls the relevant radiation branch	A contradiction between generalized-entropy calculations and the branch-ledger reading.
Quantum simulators / analog horizons	Open-channel entropy growth followed by branch-limited or correlation-restoring behavior	Persistent exact thermal monotonicity even in setups designed to mimic unitary late-time recovery.
Subsystem structure	Gauge-constrained, non-naive effective factorization	A fully gravitational exact factorization theorem leaving no room for DGI-style correction.
Microscopic closure	The induced- G_N species sum reproduces the Bekenstein–Hawking coefficient	A consistent spectrum in which the species sum misses $1/4G_N$.

The Part in one breath, and the door to the next. The inventions of this module are, in the end, one invention used four times: take openness seriously — let the projection be forced, let gravity be induced, let the ledger rank the branches — and the horizon’s classic paradoxes resolve

without new postulates, while the open items become work orders with registered kill conditions instead of loose ends. The physics intuition to carry forward: *a discussion should read like an audit, and this one is meant to*. The arguments are now all on the table — the derivations, the record, the comparisons, the debts. What remains is to step back and say what it means: for horizons, for information, and for the view from the six-dimensional bulk. That is the Conclusions.

68 Part X in plain English: what is actually new here, what is borrowed, and what is still owed

Four inventions, and the Part's own view is that they are one invention used four times: take openness seriously.

Let the projection be forced by the causal structure rather than chosen, and the thermality follows. Let gravity be fully induced, and the entropy coefficient stops being a count and becomes a ratio. Let a ledger rank the branches, and the Page curve becomes a crossover rather than a mechanism. Let the same ledger price each exchange, and the direction of the second law acquires a cost that can be computed. Not one of the four needs a postulate the cohort did not already have.

The related-work section is longer than a paper of this kind usually carries, and deliberately so. Sakharov's induced gravity, Jacobson's thermodynamic derivation of the field equations, complementarity, the island program — the module stands on all of them, and the Part names what each contributed instead of positioning against it.

What remains open is set out as work orders with kill conditions attached rather than as loose ends. The species sum. The factorization question. The Wall-class general statement. Each has a name, a home, and a condition under which it would end the reading rather than merely embarrass it. That is the form the Part argues an honest discussion should take, and it is why this one reads more like an audit than a summary.

Part XI

Conclusions

Horizons, Information, and the View from the Six-Dimensional Bulk

The argument of the paper is now complete. The theorems are proved, the numbers are on record, the comparisons are drawn, and the ledger of inventions and open debts is signed. What remains is the step every long argument owes its reader at the end: to put the machinery down, step back, and say plainly what it means — for black holes, for information, and for the picture of nature that emerges when the six-dimensional bulk is taken seriously. This closing Part is written for the broadest audience of the paper. Nothing in it is new physics, and nothing in it adds hypotheses to the record; each section restates, in ordinary language, something the earlier Parts established in equations, and a reader who wants the derivation behind any sentence will find the pointers in Parts IV–X. We proceed from the physics outward: first the six-dimensional picture itself, then what it does to the fifty-year paradox, then what a horizon really is, and finally what all of this suggests about determinism, observation, and the reach of physical description.

69 The 6D picture, one more time

Let us step back from the technical detail and state, in one paragraph, what GG-Theory says about black holes.

Nature, at its most fundamental level, operates in six dimensions: the four we perceive, plus a radial direction I_r along which fields stretch and compress, plus a compact angular direction S^1_θ around which gauge fields wrap. In this 6D bulk, the laws are simple and deterministic: a unitary evolution of field configurations on a smooth manifold, governed by a single action principle with gauge-gravity unification and the BI-6 modified Bianchi identity. Matter, energy, and information flow freely through all six dimensions.

Observers, however, are 4D creatures. We exist on the boundary of the 6D bulk and can only measure zero-mode projections of the bulk fields. What we see is the *OGGEP functor applied to the 6D state*: a partial trace over the compact and radial dimensions, yielding a 4D density matrix on the zero-mode Hilbert space. Because partial trace is a CPTP channel, the effective 4D dynamics is open, dissipative, and often non-unitary in appearance — even though the underlying 6D dynamics is unitary.

A black hole is, in this picture, a *region of the 4D boundary where the OGGEP projection becomes especially severe*. Near the horizon, the bulk fields develop structures that are extremely scale-extended into the compact and radial directions, so that the 4D observer loses access to correspondingly more of the bulk state. The horizon itself is the geometric locus at which the 4D observer’s accessible Hilbert space factorizes into an “accessible exterior” sector and a “traced-over interior” sector. Hawking radiation is the operational consequence of this partial trace: the effective

4D description is a thermal output from a CPTP channel whose input is the bulk state. Information is not lost; it flows into dimensions the 4D observer cannot see.

It is worth pausing on how ordinary — and how radical — this structure is. An open quantum system is simply a system whose description omits part of the world: an atom coupled to the photon field, a molecule in a solvent, a qubit leaking into its substrate. In the laboratory, the experimenter chooses the cut. Trace out the photon modes and the atom decays; trace out the solvent and the molecule decoheres; what is ignored is a matter of practical convenience, and the resulting open-system equations are honest approximations. The single structural observation on which this paper rests is that near a horizon the cut is no longer chosen — it is imposed. Causal access to the compact and radial sectors ends at the port, for every 4D observer, no matter how well instrumented. A black hole is therefore the universe’s most perfect open quantum system: the one whose environment is hidden not by experimental limitation but by the causal structure of spacetime itself. And that is precisely why this paper could be written with theorems where the subject usually offers metaphors. Openness imposed by geometry is exact and lawful in a way that openness chosen for convenience never quite is: the channel is forced, its properties are calculable, and the familiar apparatus of open-system physics — CPTP maps, entropy production, relaxation — applies to the black hole not as analogy but as description.

70 The paradox was never about information

Seen in this light, the information paradox of 1976 appears in a new perspective. Hawking asked: if the initial state is pure and the final state is mixed, where did the information go? The debate of the past fifty years has largely been about whether information is really lost, or hidden in subtle correlations, or encoded holographically on the boundary.

GG-Theory suggests a different framing. The information was never in danger of being lost; it simply flowed into bulk dimensions invisible to the 4D observer. The paradox was an artifact of assuming that the 4D description captures the complete physical content. Once one recognizes that 4D observation is a partial trace of a 6D state, the apparent paradox dissolves into a statement about how much of the bulk the observer has access to. Fifty years of anguish about information loss were, on this reading, anguish about mistaking a shadow for the object.

There is a quieter way to say the same thing. Every working physicist already writes down mixed states without alarm. When a laboratory qubit entangles with its surroundings and the experimenter traces the surroundings out, the qubit evolves from a pure state to a mixed one, and nobody declares a crisis: the purity has gone into correlations with degrees of freedom the description no longer tracks. Hawking’s 1976 calculation produced exactly such a mixed state. The alarm arose because, for a black hole, there seemed to be no environment for the purity to have gone into — the radiation was all there was, and it was thermal. The 6D bulk supplies precisely the missing register. The compact and radial sectors behind the port are the environment, as physically real as the photon modes of a laboratory, and the mixedness of Hawking radiation becomes the same benign bookkeeping that every open system exhibits. What was extraordinary about the black hole was never the physics; it was that geometry had hidden the environment so completely that

physics spent five decades looking for it.

This is a structurally similar move to previous reconceptions in physics. When Copernicus moved the observer off the Earth to see that the “apparent” motion of the Sun was really due to Earth’s rotation, he did not change the data; he changed the interpretive frame. When Einstein moved the observer into a freely-falling elevator to see that the “apparent” gravity of the elevator was really acceleration, he did not change the data; he changed the frame. When Copenhagen-school physicists insisted that the wave function describes what we know rather than what is, they did not change the experimental outcomes; they changed the interpretive category.

GG-Theory makes a similar move: the 4D information paradox dissolves when we recognize that the 4D observer sees only a partial trace of a 6D unitary reality. The 4D frame is correct as far as it goes, but it is incomplete, and the paradox is in the incompleteness rather than in the underlying physics.

The missing register, named by the anatomy. The cohort’s recent campaign (GG-W1-Symmetry; [Arisaka, 2026, P5-CPT](#); [Arisaka, 2026, A8-QM](#)) lets the “missing environment” be named with topological precision. The two sectors behind the port are fibers of opposite kind — *holonomy requires closed loops, and an interval has none*: the compact circle carries everything quantized and protected, a group, the fiber that remembers; the scale direction carries only one-way transport, a semigroup, the fiber that forgets. What the paradox of 1976 asked of the black hole was, in this language, to be both fibers at once: to *keep* — purity, unitarity, the conserved ledgers — and simultaneously to *spend* — thermality, irreversibility, the one-way evaporation. The anatomy divides the labor instead. Everything the evaporation must conserve is circle-side: the charges, the ledger integers, the holonomy-class records — deposited on the fiber no boundary operation can continuously deform, which is why the information was never destructible in the first place. Everything thermal and one-way about Hawking’s calculation is interval-side: the temperature is the shore periodicity of the CPT companion’s KMS reading ([Arisaka, 2026, P5-CPT](#)) — *temperature is where the interval meets the shore*, and a horizon is where that meeting is sharpest — while the evaporation itself is the semigroup doing the only thing a semigroup does. The paradox was thus a billing error: it demanded that the interval’s one-way channel also carry the circle’s books, and fifty years of debate audited the wrong fiber. The Page curve, in this reading, is the reconciliation schedule — the circle’s ledger read out through the interval’s channel, at exactly the rate the channel permits. In the campaign’s shorthand: *the circle keeps what the hole must return; the interval prices the return; the paradox was billing one fiber for the other’s job.*

71 What is a horizon?

This paper has used the word “horizon” many times, in several senses. The word deserves a moment of reflection.

Classically, a horizon is a causal boundary: a surface in 4D spacetime that can be crossed in one direction but not the other. The event horizon of a black hole is the most famous example; the cosmological horizon in an accelerating universe is another.

In GG-Theory, horizons take on an additional structural meaning: they are *ports of the OGGEF functor*, loci at which bulk information flows into dimensions inaccessible to the 4D observer. The black-hole horizon is one example. The cosmological horizon is another. The apparent “inside” of a chemical compound, the unobservable wave function of Schrödinger’s cat, the hidden thermal bath of a Brownian particle — all of these are, in the GDOS reading, different kinds of horizons. Each marks a boundary between what the observer can access and what has been traced over.

Stated this way, the horizon stops being a place and becomes a relation. Nothing locally remarkable happens at the Schwarzschild radius; an infalling observer crosses it without incident, as relativists have rightly insisted for a century. What changes at the horizon is not the local physics but the relation between one particular observer and the bulk state — which registers of the world that observer’s port can still reach. A surface can therefore be a horizon for one observer and not for another, just as a hill can block one walker’s view and not a neighbor’s, without anyone doubting that the landscape behind it is real.

The lesson generalizes to the properties themselves. Is the black hole hot? For the distant observer, yes — measurably, calculably. For the infalling one, no. Both are right, because temperature lives in the relation between the bulk and a port, not in the bulk alone (§28); the same is true of entropy, and even of when the information comes back. Physics has met this before — simultaneity became frame-dependent in 1905 without becoming unreal — and the horizon extends the lesson to thermodynamics: *objective* does not mean *observer-free*; it means lawfully related across observers.

The unification of these diverse phenomena under the single concept of OGGEF port is one of the deeper structural claims of GG-Theory. It suggests that the mathematics of black holes — entanglement entropy, thermodynamic radiation, information flow — applies very generally to any open quantum system, with the black hole being simply the most gravitationally spectacular example. If that claim is right, a student who has understood decoherence in a quantum-optics laboratory already holds, perhaps without knowing it, the conceptual core of black-hole thermodynamics. The horizon adds geometry’s signature — the exact area law, the thermal clock, the branch ledger — but not a new rulebook.

72 On determinism, probability, and the view from nowhere

The GG-Theory picture restores a form of determinism to fundamental physics. In the 6D bulk, the evolution is unitary and deterministic: given the initial bulk state, the later bulk state is fixed. Randomness, probability, and irreducible uncertainty arise only at the 4D boundary, as consequences of the partial-trace structure of OGGEF.

This position has historical echoes. Einstein famously believed that “God does not play dice with the universe” and resisted the Copenhagen interpretation’s insistence on irreducible randomness. His colleague Schrödinger, too, was uncomfortable with the Born-rule probabilities. Both would have been, we suspect, sympathetic to the GG-Theory picture, in which the 4D probabilities emerge from a deeper 6D determinism through the projection functor. In the 6D bulk, evolution is unitary and nothing is ever lost — Einstein’s instinct. In every 4D description, outcomes are probabilistic and entropy grows — Bohr’s. The horizon is where their century-old argument is staged most

starkly, and the resolution offered here is that they were describing different levels of the same world, each correctly.

It is worth being precise about the kind of probability that survives this reading. The randomness a 4D observer encounters is not a die thrown by nature at each event; it is the lawful ignorance of an observer whose port delivers only part of the state. This is closer in kind to the probability of statistical mechanics, which quantifies a coarse-graining, than to irreducible chance — with one decisive difference. In statistical mechanics the coarse-graining is optional: a better microscope, in principle, refines it away. Here it is not. The missing information sits behind a causal boundary rather than a practical one, and no 4D instrument, however perfect, can retrieve it. Probability, in this picture, is what determinism looks like through a port. (How the quantitative Born weights themselves emerge from the projection is the business of the quantum-mechanics companion (Arisaka, 2026, A8-QM), where the status of each step in that derivation is registered explicitly; nothing in the present paper depends on it.)

But this deterministic restoration comes at a price. The 6D bulk is not directly observable; it is defined by its projections. Its “existence” is, in a metaphysical sense, posited rather than demonstrated. The question “does the 6D bulk really exist?” is not a simple factual question in the way “does Alpha Centauri exist?” is; it is a question about what we should take as the fundamental description of nature.

The pragmatic reply is that the 6D bulk is as real as any unobservable scientific entity. Atoms, quarks, the electromagnetic field, the Higgs boson, dark matter — none of these is directly seen; all are posited as the best explanation of observable phenomena. The 6D bulk, in GG-Theory, is a posited structure whose existence is warranted by its success in reading the Hubble tension, the W mass, the Higgs sector, the neutrino structure, quantum measurement, thermodynamic arrow of time, and — in this paper — the black-hole information paradox. If the 6D bulk explains all of this without adjustable parameters, the inference to its existence is as warranted as the inference to any other unobservable.

The philosophical view this suggests is sometimes called the “view from nowhere”: a perspective not tied to any particular observer, in which the fundamental description is the 6D bulk, and every 4D description is a particular partial-trace projection. No 4D observer sees the bulk directly; every 4D description is incomplete; yet the bulk grounds and unifies all of them.

The phrase has a philosophical history: it names the ideal of a description of the world not indexed to any particular point of view. Physics has pursued that ideal since Copernicus, and modern physics has twice discovered its limits — special relativity showed there is no observer-free fact about simultaneity, and quantum mechanics seemed to show there is none about measurement outcomes. GG-Theory proposes a precise division of the territory. There *is* a view from nowhere: the 6D bulk state and its unitary law. But it is not a view any observer can occupy. Every actual point of view is a port, and every port truncates — not as a defect of the observer, but as a statement about the projection itself. A reader may hear in this an echo of the old philosophical distinction between the world as it is and the world as it appears. The difference is that here the relation between the two is not left as a mystery: it is a specific, computable map, and the appearances it produces are constrained strongly enough to be falsifiable. That is what lifts the picture from

metaphysics to physics — not that it closes the gap between reality and observation, but that it gives the gap equations.

73 Black holes as Rosetta stones

For a generation of physicists, black holes have been the Rosetta Stones of quantum gravity. They are the cleanest setting in which general relativity, quantum mechanics, and thermodynamics meet head-on. Their entropy is gravitational (area) and informational (microstate count) and thermodynamic ($A/4G_N\hbar$) simultaneously. Their horizons are geometric boundaries, information-theoretic partial-trace surfaces, and thermodynamic heat baths at the same time. If you understand black holes, you have a good claim to understanding quantum gravity.

The metaphor deserves to be taken literally for a moment. The Rosetta Stone deciphered Egyptian because a single decree was carved three times — in hieroglyphs, in Demotic, in Greek — so that a reader fluent in one script could break into the others. The black hole is nature’s analogue: one object on which the same statement is written in three languages at once. $S_{\text{BH}} = A/(4G_N\hbar)$ is a single equation that reads as geometry (an area), as information (a count of microstates), and as thermodynamics (a heat divided by a temperature). Whoever learns to parse all three scripts at once holds a key that opens far more than black holes.

GG-Theory takes this Rosetta-Stone status seriously. The same 6D geometry that explains the Higgs boson, the fine-structure constant, the cosmic neutrino background, and the accelerating expansion of the universe also explains the black-hole entropy and the Page curve. The black hole becomes the setting in which all the diverse observables of the program are tied together, because its extreme geometry amplifies the bulk-to-boundary projection to the point of visibility. And in the cohort’s reading there is a fourth script on the stone: the particle-physics script, in which the same integer ledger and the same lock scale that organize the Higgs sector and the dark-matter prediction reappear in the horizon’s bookkeeping. A comparison table can register that as consistency; the physics registers it as unity.

For readers outside the specialist community, this is perhaps the central message of the paper. The fifty years of information-paradox literature, the many-volumed books on black-hole physics, the vast output of the AdS/CFT community since 1997 — all of this is not separate from the physics of everyday particles and forces. It is continuous with it. The horizons of black holes obey the same partial-trace structure that makes your laptop lose information to the environment when it heats up, that makes sugar dissolve in tea, that makes a photograph fade in sunlight. Everyday irreversibility and Hawking radiation are one mechanism at two volumes; the black hole is simply the place where geometry performs the projection that nature otherwise leaves to circumstance. Black holes are just the most dramatic examples of a universal structure in how 4D observers see a 6D reality.

74 A closing reflection

We close with a reflection on what the work reported here does and does not achieve.

It does not prove that GG-Theory is correct. The framework makes predictions — some already

consistent with observations, some awaiting future experiments — but correctness of a physical theory can never be proved; it can only be progressively supported by evidence or falsified.

It does not resolve every open question about black holes. The quantitative details of the microstate counting, the precise form of the KK corrections to the Hawking spectrum, the connection to numerical-relativity simulations, and many other technical issues remain open. This paper establishes a framework in which these questions can be asked; it does not answer all of them.

It does not replace the existing literature on the information paradox. The AdS/CFT correspondence, the islands formula, the fuzzball program, the ER=EPR proposal — each of these captures part of the truth about black holes. GG-Theory synthesizes rather than replaces them, situating each within a larger framework grounded in a complete 6D theory.

What the work does achieve, we believe, is the following: it *reframes* a fifty-year-old problem in a way that dissolves the paradox without introducing new physical mysteries. The information paradox, in the GDOS reading, is not a puzzle about Nature; it is a puzzle about which description of Nature is fundamental. Once the 6D bulk is taken as fundamental and the 4D boundary as its partial-trace projection, the paradox dissolves automatically, and the many fragmented resolutions of the past five decades (islands, wormholes, fuzzballs, firewalls) are revealed as different pieces of a single picture. And it does so under the ledger’s discipline: no parameter was tuned, no mechanism was invented for the occasion, and every load-bearing claim carries a pre-registered way to kill it. Readers who wish to test that discipline step by step will find the machine room in the Technical Appendices and, in the objections Part, thirty questions from the general reader answered in plain English followed by fifty-three adversarial readings; the paper asks to be checked, not believed.

The black hole was called, by John Wheeler in the 1960s, a “mass without matter.” In GG-Theory it becomes something more specific: *a window from our 4D world into the 6D bulk*. What falls into a black hole does not disappear; it simply enters dimensions we cannot see. What comes out as Hawking radiation is the echo of that crossing, projected back into our 4D observation. The horizon is not a wall; it is the boundary of what we can access.

For a century the black hole has been physics’ memento mori — the place where stars die, where clocks stop, where theories are sent to fail. For half of that century it carried a darker charge: the fear that the horizon devours information, and with it the lawfulness of nature itself. The reading defended in this paper ends on a gentler verdict. The horizon destroys nothing. It is not where physics breaks down but where a perspective does — the surface on which nature states, exactly and quantitatively, how much of the world one observer’s window can carry. What falls in is not erased but kept, in registers we do not see, and returned, slowly, in correlations we are only beginning to learn to read. If this framework is right, the fifty-year crisis of the black hole was never a crisis of nature. It was the moment a species of four-dimensional observers began to discover that its window is not the world — and the same discovery, carried home from the horizon, reaches into everything the window frames: why a measurement yields one outcome, why heat flows one way, why the past feels fixed and the future open. Seen from the bulk, nothing is ever lost. The universe keeps perfect books; what we call paradoxes are the shadows thrown on our four-dimensional wall by a reality two dimensions richer than our sight. That, in the end, is what the black hole has been telling us for fifty years, patiently, in a language of area and temperature we have only now learned

to read: that the wall has a door, that the books balance, and that the world is larger than the window through which we watch it.

75 Part XI in plain English: what a horizon turns out to be, if this reading is right

If the reading is right, a horizon is not primarily a place. It is a boundary of access, and everything strange about black holes follows from that one property rather than from anything exotic in the geometry.

The temperature is what a traced-over state looks like. The entropy is the size of what was traced. The coefficient is a ratio between two normalizations that had to be fixed anyway. The information was never lost, only held on the far side of an assumption that expires on schedule. And the singularity is where a four-dimensional coordinate system fails, not where the universe does.

That reframing has a consequence worth stating plainly, because it is what this Part is really about. The paradox was never about information. It was about assuming that a description written from one side of a one-way boundary ought to be complete, and then being surprised when it was not. Every resolution proposed in fifty years can be read as a way of relaxing that assumption; this one relaxes it at the level of the geometry rather than at the level of the formalism.

What would end it is short, and printed: a coefficient other than one quarter, a departure from Kerr outside the horizon, a species sum that does not close. None of the three is far-fetched, and one of them is a calculation somebody could begin this year.

The Part closes on the reason black holes are worth this much work. They are the only objects where gravity, quantum mechanics and thermodynamics all have to be right at the same time, and they have been sending the same three-number message for a century. The claim of this paper is that the message was never as short as it looked.

Author's Note

This paper honors four theorists whose work on black-hole physics this paper builds upon:

Jacob Bekenstein (1947–2015) proposed in 1972–1973 that black-hole horizon area is a physical entropy. His argument, initially controversial, was the first step toward the modern understanding of black-hole thermodynamics. He died in 2015, five years before the Penington-Almheiri-et-al. papers vindicated his framework at the deepest level. Bekenstein's $S_{\text{BH}} = A/(4G_N\hbar)$ becomes, in GG-Theory, the entanglement entropy of a 4D observer with the 6D bulk — a precise geometric meaning that Bekenstein glimpsed but could not formalize in his own lifetime.

Stephen Hawking (1942–2018) in 1974 derived black-hole thermal radiation, identifying the key paradox that has driven 50 years of work. Hawking's intellectual honesty — he initially argued for information loss, then publicly changed his position in 2004 — is a model for how controversial questions should be approached. His 2016–2018 work on soft hair suggested horizon-level degrees of freedom that we now understand as $U(1)_{Q_i}$ holonomies. Hawking did not live to see the islands formula or the present work, but his framework provides the foundation on which they build.

Don Page (b. 1948) in 1993 derived the entanglement entropy curve that bears his name. Page's remarkable insight — that thinking carefully about typical quantum states gives quantitative predictions about black-hole evaporation — anticipated by 25 years the islands-formula developments

and provided the target that any theory of quantum gravity must hit. Page's curve is, in **GG-Theory**, a consequence in shape of bulk unitarity plus **OGGEP** projection; its quantitative form is open work.

Juan Maldacena (b. 1968), in 1997, introduced the **AdS/CFT** correspondence and transformed theoretical physics for a generation. His 2013 **ER=EPR** proposal (with Susskind) and his work on the islands program have been central to the recent progress. Although **GG-Theory** does not use the full **AdS/CFT** machinery, it inherits the central lesson of Maldacena's program: that gravitational physics in a bulk is dual to/encoded in a boundary theory, and that information flows are fundamentally holographic. **GG-Theory** makes this precise in the real, Standard-Model-containing universe rather than in the abstract $\mathcal{N} = 4$ SYM setting.

We also acknowledge the broader string-theory, holography, and numerical-relativity communities whose work over five decades made the synthesis reported here possible.

Acknowledgments

This work was developed through extensive iterative collaboration with three large language models: ChatGPT (OpenAI) for exploratory drafting, narrative development, and broad-scope ideation; Perplexity for literature synthesis, citation verification, and recent experimental cross-referencing; and Claude (Anthropic) for structural organization, mathematical consistency checking, LaTeX typesetting, and final manuscript preparation. The sequential refinement of the GG-Theory framework (GG-Constitution / GG-6 / BI-6 / GDOS / Gi-4 Pillars / Qi-4 Pillars) across hundreds of conversation sessions with these three systems was essential to the coherence of the present paper.

The author thanks Aaron Blaisdell and Hirotaka Iwasaki for stimulating discussions on the program's interdisciplinary scope and for critical feedback on the accessibility of the presentation. The author also thanks hundreds of collaborators in the Elegant Mind Club over the past decade, including Javier Carmona, Elizabeth Milles, Mingda He, Syed Bashir Hydari, Brian Ta, Chris Dao, Umaima Afifa, Isabella Bestanoby, Caominh Le, and Natsuko Yamaguchi.

All scientific concepts (especially the foundation of the GG-Constitution, GGEP/GG-6/BI-6, and OGGE/ GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

The author thanks the UCLA physics department for institutional support, colleagues at UCLA and elsewhere for critical feedback, and the broader gravitational, cosmological, and string-theory communities for five decades of theoretical progress on black-hole physics. The present synthesis would not have been possible without the preceding work of Bekenstein, Hawking, Page, Maldacena, Susskind, Strominger, Vafa, Witten, 't Hooft, Penington, Almheiri, Engelhardt, Marolf, Maxfield, and many others whose contributions are cited throughout.

We note also the contributions of the Event Horizon Telescope collaboration and the LIGO/LISA collaborations, whose observational work provides the empirical context against which any theoretical framework must be tested.

Part XII

Technical Appendices

A graduate-level course in black-hole physics read through **GG-Theory** — full from-scratch derivations of the Schwarzschild geometry, the four laws of black-hole mechanics, Bekenstein’s entropy, the Hawking effect, the Euclidean section, the Hawking–Page transition, the Page curve, and the modern holographic era from state counting to islands — preceded by the reference dictionaries and followed by the source ledger, benchmark audit, and registries

The body of this paper has emphasized the theorem chain and kept the heaviest calculation out of the reader’s way. The present appendices supply that technical content in full, and taken together they are intended to be read as *a graduate-level course in black-hole physics seen through GG-Theory*, addressed to an incoming graduate student in physics or astronomy who will carry the program forward. We organize them into three groups. The *reference* appendices (A–D) tabulate the abbreviations, the key symbols, the cross-cohort naming convention, and the numerical reference card, and serve as the reader’s quick lookup. The *course* appendices (E–R) are the course proper: a self-contained sequence that takes up black-hole physics in the order the subject was actually built — the Schwarzschild geometry and its horizon; the four laws of classical black-hole mechanics; Bekenstein’s entropy and the generalized second law; the Hawking effect derived with the full Bogoliubov machinery; the Euclidean section and the Hawking–Page transition; the Page curve; the modern holographic era from the Strominger–Vafa count to the islands formula; and, in the closing chapters, the channel identities, the extremal and de Sitter cases, and the induced-gravity species sum. Wherever the material allows, each chapter follows *three steps*. First comes a full from-scratch derivation of the standard result, in the form a graduate student would meet it in Wald or Misner–Thorne–Wheeler, carried through with no skipped manipulations and closed by a worked numerical example. Then comes a frank statement of what that standard derivation *assumes but cannot explain* — the questions it leaves open: what the horizon is made of, what states the entropy counts, where the information goes. Finally comes the **GG-Theory** reading, in which the same quantity descends as a theorem of the open-system boundary law — and because those derivations are the business of the main text, the third steps mostly *point*: into Parts IV–VI for the port, the exact 1/4, and the Page turnover, and into Part IX for the comparisons. The *ledger-and-registry* appendices (S–Y) that close the Part record the internal source ledger, the cross-paper inheritance table, the consolidated benchmark audit, the named-statement deck, the promotion roadmap, and the falsification-test catalog, closing with an annotated reading guide.

Throughout these appendices we adhere to two stylistic commitments, shared with the cohort’s reference implementation (**GG-A10-Cosmos**). First, every derivation is presented step by step with no skipped intermediate manipulations: a graduate student should be able to reproduce each calculation with paper and pencil, and every number quoted anywhere in this Part traces to a verification script named in the benchmark audit (Appendix U). Second, every formal step is accompanied by a

physics intuition note that translates the algebra into a mental picture, on the spot rather than in a separate section. Each course appendix opens with a boxed *What this appendix teaches* abstract and closes with a boxed *take-away* — the standard result, what it assumed, and what the geometry replaced the assumption with — so that the Part can be used as a one-semester reading course, one appendix per week, each self-contained. And every derivation chapter closes with *problems and complete worked solutions* — placed before the take-away box, numbered by chapter and Step, each solution worked in explicit Steps and ending with its physics intuition, the sets cross-referencing one another into a single course network. They carry full solutions because, in the cohort’s phrase, they are meant to be done, not admired: the reader who works them owns the derivations, not merely the conclusions.

Prerequisites and recommended textbooks. The course assumes advanced undergraduate mechanics, electrodynamics, statistical physics, and quantum mechanics, plus a first exposure to general relativity (the Schwarzschild solution is derived from scratch in Appendix E, so a student meeting black holes for the first time can begin there). It is meant to be read *alongside* a canonical textbook, not in place of one: the recommended way to use each appendix is to read the corresponding textbook chapter first, return here for a compact but complete recap, and only then proceed to the later Steps, where the standard result is re-examined and re-read as a statement of the geometry. For self-study we particularly recommend the following, each cited again in the appendix to which it is best matched:

- **General relativity foundations.** Wald (1984), *General Relativity*, and Misner et al. (1973), *Gravitation* — the two canonical graduate texts, covering the differential geometry, the Einstein equations, and the classical black-hole solutions in the conventions this course follows; Carroll (2004), *Spacetime and Geometry*, as the gentler modern first pass.
- **Black-hole mechanics at depth.** Poisson (2004), *A Relativist’s Toolkit* — the standard source for congruences, the Raychaudhuri equation, and horizon mechanics (Appendices E–F); Chandrasekhar (1983), *The Mathematical Theory of Black Holes*, and Frolov and Novikov (1998), *Black Hole Physics*, for the complete classical theory including the Kerr–Newman sector (Appendix P).
- **Quantum fields in curved spacetime.** Birrell and Davies (1982), *Quantum Fields in Curved Space* — the classic source for the Bogoliubov machinery of Appendix H; Parker and Toms (2009), *Quantum Field Theory in Curved Spacetime*, for the effective-action and heat-kernel methods behind Appendix R; Wald (1994), *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*, for the algebraic formulation and the generalized second law.
- **The information era.** Susskind and Lindesay (2005), *An Introduction to Black Holes, Information and the String Theory Revolution* — the most readable account of the information paradox as the field understood it before the islands era; Harlow (2016), the Jerusalem lectures, for the quantum-information toolkit (Appendices K and O); and Almheiri et al. (2021), *The entropy of Hawking radiation*, for the modern islands synthesis reviewed in Appendix L.

Throughout, GG-A1-Constitution–GG-A11-Origin, GG-P1-KK–GG-P5-CPT, and GG-C1-Kaxion–GG-C3-BH refer to the papers of the GG-Theory cohort, to which the third steps repeatedly point.

A Summary of abbreviations

For the canonical Dictionary of program-defined acronyms (45 entries in ten categories) see GG-A2-GG6 Appendix B; the entries below are the working set of *this* paper, one line each.

AX-n / POS-n The constitution’s axioms and postulates (GG-A1-Constitution); AX-1 causality, AX-2 GGEP, AX-3 OGGEF, AX-4 MSMF.

GGEP The geometric equivalence principle: gravity is bulk geometry; no bare boundary G_N (the premise of Theorem 1).

OGGEF The bulk-to-boundary projection — observation itself; at a horizon, geometrically forced (L-BH-Port).

GDOS The open-system boundary law (GG-A5-GDOS); its freshness clause L-GDOS-Fresh is consumed with a validity window by Lemma 2.1.

GFF / GEP / CEP The universal free-energy ledger and its geometric and curvature–entropy pieces; $GFF = F_{GEP} + F_{CEP}$.

GRIP / GRIP-In / GRIP-Out The runtime descent and its two regimes (fixed landscape / landscape reorganization).

GA / GGA Geometric Attractor; Ground Geometric Attractor (the basin with no exit channel).

CPTP Completely positive, trace preserving — the channel class of every boundary map in this paper.

GSSB Geometric Spontaneous Symmetry Breaking (GG-A6-GRIP): the landscape-reorganization stage of the GRIP cascade; a GSSB event in the entropy ledger (T-BH-2) is a reorganization of its minimizing branch.

DGI Dynamical Gauge Invariance (owner: GG-A1-Constitution, stratum POS-1; the initial stage of GG-A6-GRIP’s GRIP cascade). Used in this paper for its constrained-subsystem content in the factorization puzzle (§33).

qes Quantum extremal surface (the island literature’s minimization locus).

KMS The thermal (Kubo–Martin–Schwinger) condition — the clock’s thermal face (Theorem 5).

QNM Quasi-normal mode (the ringdown spectrum).

EMRI Extreme-mass-ratio inspiral (the LISA-era Kerr multipole probe).

BI-6 / OGN-5 / GU-5 The bulk anomaly identity; the integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$; the five universal constants $\{c, G_N, \hbar; k_B, Q_B\}$.²

²Notation, kept uniform across the cohort: N is the OGN-5 ledger’s second entry, and every ledger identity is written with it. $N_c = 3$ is the color triplicity, the three colors of $SU(3)_C$ — what that entry counts, named where the

T-BH- n / L-BH-* / C-BH-* / F-BH- n This module’s keyed theorems, lemmas, corollaries, and falsifier rows; the full deck is Appendix V.

B Glossary of key symbols

κ, β Surface gravity; inverse clock period $\beta = 2\pi/\kappa$ (the geometry’s one contribution to thermality).

$A, S_{\text{BH}}, S_{\text{gen}}$ Horizon area; Bekenstein–Hawking entropy $A/4G_N$; generalized entropy $A/4G_N + S_{\text{out}}$.

r, θ, K The scale coordinate on I_r with depth $K = 35$; the phase angle on S_θ^1 .

ℓ_{lock} The compactification length $\hbar/(M_{\text{lock}}c)$; the suppression scale of every Kerr deviation, $(\ell_{\text{lock}}/r_s)^2$.

$\alpha_\omega, \beta_\omega$ Bogoliubov coefficients; $|\beta_\omega/\alpha_\omega| = e^{-\pi\omega/\kappa}$ is the branch-cut datum (T-BH-4) and the KMS face displayed (T-BH-5).

$J, |A|$ The clock’s rotor (derived complex structure) and speed (modulus) — the polar split of the Killing generator.

$\Gamma_\ell(\omega), \Gamma$ Graybody factor; exchange-channel transmissivity.

$\bar{n}_H, g(n)$ KMS occupation $(e^{\hbar\omega/k_B T_H} - 1)^{-1}$; thermal entropy function $(n+1)\ln(n+1) - n\ln n$.

$S(\rho||\sigma)$ Relative entropy — the monotone of the data-processing inequality, and the exact GSL production of Lemma 2.2.

$I(q;R)$ Mutual information of the next engaged mode with the radiation record — the freshness meter of Lemma 2.1.

$t_{\text{Page}}, M_\times$ The freshness-budget expiry time; the conditional-GGA crossover mass (C-BH-1) — also the GSL’s equality point.

a_1, a_2, ϵ Seeley–DeWitt heat-kernel coefficients; proper-time cutoff (the shared ancestor of $1/G_N$ and dS_{ent}/dA).

$\chi_i, c_{\log}$ Per-species induced-gravity weights (the open sum, F-BH-9); the universal logarithmic-correction coefficient.

C Cross-cohort naming

Cohort papers are cited by designator (GG-A1-Constitution through GG-P5-CPT, GG-C1-Kaxion–GG-C3-BH) with natbib labels of the form (Arisaka 2026, XX-Name); keyed statements are cited by their permanent keys (T-BH- n , L-GDOS-Fresh, T-P5-B1, L-Constitution-1), which never change across versions of record; and no two Arisaka keys share one `citep`. The full convention is the cohort’s, stated once in GG-A1-Constitution.

color reading is the point. $N_g = 3$ is the generation (family) count; it is *not* a ledger entry and is not fixed by the ledger (GG-P2-N3). $N_f = 2N_g = 6$ is the quark-flavor count of a beta function at high scale. Three of these are three and one is six, and they are three different quantities.

D Numerical parameters for reference

- Planck mass: $M_{\text{Pl}} \approx 2.18 \times 10^{-8}$ kg.
- Planck length: $\ell_P \approx 1.62 \times 10^{-35}$ m.
- Schwarzschild radius of solar-mass BH: $r_s(M_\odot) = 2G_N M_\odot / c^2 \approx 2.95$ km.
- Hawking temperature of $M_\odot$ BH: $T_H(M_\odot) = \hbar c^3 / (8\pi G_N M_\odot k_B) \approx 6.17 \times 10^{-8}$ K.
- Evaporation time of $M_\odot$ BH: $t_{\text{evap}} = 5120\pi G_N^2 M^3 / (\hbar c^4) \approx 2 \times 10^{67}$ yr; Page time $t_{\text{Page}} \approx 0.646 t_{\text{evap}}$.
- Entropy of $M_\odot$ BH: $S_{\text{BH}} \approx 10^{77} k_B$.
- M87* mass: $M \sim 6.5 \times 10^9 M_\odot$; shadow $\sim 42 \mu\text{as}$.
- Sgr A* mass: $M \sim 4.3 \times 10^6 M_\odot$; shadow $\sim 51.8 \mu\text{as}$.
- GG-Theory compactification length: $\ell_{\text{lock}} \equiv \hbar / (M_{\text{lock}} c) \simeq 5 \times 10^{-20}$ m (with $M_{\text{lock}} = 3.85$ TeV), longer than the unification length $\hbar c / E_{\text{GUT}} \simeq 8.7 \times 10^{-33}$ m ($E_{\text{GUT}} = M_{\text{Pl}} e^{-2\pi} \simeq 2.3 \times 10^{16}$ GeV, the GG-P1-KK ladder) by the factor $2e^{K-2\pi} \simeq 6 \times 10^{12}$; the depth- K hierarchy itself is $M_{\text{Pl}} / 2M_{\text{lock}} = e^K \simeq 1.6 \times 10^{15}$.
- OGN-5 integer $K = 35$; scale hierarchy $e^K \sim 10^{15}$.

E The Schwarzschild geometry from Einstein's equations

What this appendix teaches. The complete derivation of the Schwarzschild solution from the vacuum Einstein equations — the ansatz, the two field equations that survive, and their integration — followed by the geometry of the horizon (why $r = r_s$ is a coordinate artifact and $r = 0$ is not), the Kruskal extension and the Penrose diagram, and the orbits: photon sphere, ISCO, and the shadow, closed by a worked computation of the M87* shadow angle against the EHT measurement. Then: what the classical solution assumes and cannot answer, and the GG-Theory reading in which the horizon becomes a port and Finkelstein's 1958 lesson repeats one level up.

Read alongside: [Carroll \(2004\)](#) chs. 5–6 for a first pass; [Wald \(1984\)](#) ch. 6 or [Misner et al. \(1973\)](#) chs. 23–25 at course level; [Chandrasekhar \(1983\)](#) ch. 3 at depth.

Every course on black holes begins where the subject began: with the solution Karl Schwarzschild found within weeks of Einstein's November 1915 field equations ([Einstein, 1915](#); [Schwarzschild, 1916](#)), computed in a wartime field post. This appendix derives it with every manipulation displayed, then walks the geometry until the two central surprises — the removable singularity at $r = r_s$ and the irremovable one at $r = 0$ — are theorems the reader owns.

E.1 Step 1 — the vacuum solution, worked

We seek the general static, spherically symmetric metric. The most general such line element can be brought to the form

$$ds^2 = -e^{2\alpha(r)} c^2 dt^2 + e^{2\beta(r)} dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2), \quad (74)$$

with two unknown functions $\alpha(r)$, $\beta(r)$. (*Physics intuition*: staticity means a timelike Killing vector orthogonal to constant- t slices — nothing moves, nothing rotates; spherical symmetry uses the angular part to define r so that a sphere at coordinate r has area exactly $4\pi r^2$. Both assumptions will be relaxed at the end of the Step and nothing will change.)

The nonvanishing Christoffel symbols of (74) are computed from $\Gamma_{\mu\nu}^\lambda = \frac{1}{2}g^{\lambda\rho}(\partial_\mu g_{\nu\rho} + \partial_\nu g_{\mu\rho} - \partial_\rho g_{\mu\nu})$:

$$\Gamma_{tr}^t = \alpha', \quad \Gamma_{tt}^r = c^2 \alpha' e^{2(\alpha-\beta)}, \quad \Gamma_{rr}^r = \beta', \quad \Gamma_{\theta\theta}^r = -r e^{-2\beta}, \quad \Gamma_{r\theta}^\theta = \Gamma_{r\phi}^\phi = \frac{1}{r}, \quad (75)$$

plus $\Gamma_{\phi\phi}^r = -r e^{-2\beta} \sin^2 \theta$, $\Gamma_{\phi\phi}^\theta = -\sin \theta \cos \theta$, $\Gamma_{\theta\phi}^\phi = \cot \theta$, with primes denoting d/dr . Assembling the Ricci tensor $R_{\mu\nu} = \partial_\lambda \Gamma_{\mu\nu}^\lambda - \partial_\nu \Gamma_{\mu\lambda}^\lambda + \Gamma_{\lambda\rho}^\lambda \Gamma_{\mu\nu}^\rho - \Gamma_{\nu\rho}^\lambda \Gamma_{\mu\lambda}^\rho$ and demanding vacuum, $R_{\mu\nu} = 0$, three independent equations survive:

$$R_{tt} = c^2 e^{2(\alpha-\beta)} \left[\alpha'' + (\alpha')^2 - \alpha' \beta' + \frac{2\alpha'}{r} \right] = 0, \quad (76)$$

$$R_{rr} = -\alpha'' - (\alpha')^2 + \alpha' \beta' + \frac{2\beta'}{r} = 0, \quad (77)$$

$$R_{\theta\theta} = 1 - e^{-2\beta} [1 + r(\alpha' - \beta')] = 0. \quad (78)$$

Multiply (76) by $e^{2(\beta-\alpha)}/c^2$ and add (77): everything cancels except

$$\frac{2}{r}(\alpha' + \beta') = 0 \quad \implies \quad \beta(r) = -\alpha(r), \quad (79)$$

choosing the integration constant so that the metric is Minkowski at infinity ($\alpha, \beta \rightarrow 0$). Substituting into (78):

$$1 = e^{2\alpha} (1 + 2r\alpha') = (r e^{2\alpha})', \quad \implies \quad r e^{2\alpha} = r - r_s \quad (80)$$

with r_s an integration constant. Hence

$$\boxed{ds^2 = -\left(1 - \frac{r_s}{r}\right) c^2 dt^2 + \left(1 - \frac{r_s}{r}\right)^{-1} dr^2 + r^2 d\Omega_2^2} \quad (81)$$

The constant is fixed by the Newtonian limit: for weak fields $g_{tt} = -(1 + 2\Phi/c^2)$ with $\Phi = -G_N M/r$, so

$$r_s = \frac{2G_N M}{c^2} \approx 2.95 \text{ km} \times \frac{M}{M_\odot}. \quad (82)$$

(*Physics intuition*: r_s is exactly the radius at which the Newtonian escape velocity reaches c — the coincidence Michell and Laplace noticed in the 1780s. It is a coincidence: the Newtonian calculation has light slowing and returning, while in relativity nothing inside r_s moves outward at all. But it correctly signals that r_s is where escape becomes the issue.)

Two remarks complete the Step. First, nowhere did we impose time-independence on the two field equations beyond the ansatz; if one repeats the computation with $\alpha(r, t)$, $\beta(r, t)$, the tr -component of the field equations forces $\dot{\beta} = 0$ and the rest forces (81) again. This is *Birkhoff's theorem*: the vacuum outside any spherical mass — even a collapsing or exploding one — is exactly static

Schwarzschild. (*Physics intuition*: spherical pulsation has no monopole gravitational radiation to emit, so the exterior cannot learn the star is moving.) Second, the solution is unique: one exterior geometry per mass, with no other adjustable property — the first appearance of the rigidity that Part I called the no-hair property and that the main text (§61.7) reads as the boundary face of MSMF.

E.2 Step 2 — the horizon is not a place where physics breaks: Kretschmann, Finkelstein, Kruskal

At $r = r_s$ the metric component g_{rr} diverges, and for half a century (1916–1958) the field wavered over what that meant. The resolution is a two-line invariant computation. The curvature scalar that cannot be transformed away is the Kretschmann invariant, and for (81)

$$K \equiv R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} = \frac{48 G_N^2 M^2}{c^4 r^6} = \frac{12 r_s^2}{r^6}. \quad (83)$$

At the horizon $K(r_s) = 12/r_s^4$ — *finite*, and for a supermassive hole absurdly small: tidal stresses at the horizon scale as $1/M^2$, so a freely falling observer crossing the horizon of M87* feels tidal stresses some ten thousand times weaker than those of standing on the Earth. At $r \rightarrow 0$, by contrast, K diverges: that is a genuine curvature singularity, unremovable by any chart. (*Physics intuition*: the mean density inside the horizon, $\bar{\rho} = 3c^6/(32\pi G_N^3 M^2)$, *falls* as $1/M^2$ — about $2 \times 10^{19} \text{ kg m}^{-3}$ for a solar-mass hole but roughly the density of water for $10^8 M_\odot$. A supermassive horizon is not a place of extreme matter; it is a place of extreme bookkeeping.)

The coordinate cure is Finkelstein (1958): define the tortoise coordinate and the advanced time

$$r_* = r + r_s \ln \left| \frac{r}{r_s} - 1 \right|, \quad v = ct + r_*, \quad (84)$$

so that (81) becomes

$$ds^2 = -\left(1 - \frac{r_s}{r}\right)dv^2 + 2 dv dr + r^2 d\Omega_2^2, \quad (85)$$

manifestly regular at $r = r_s$. Ingoing light rays are $v = \text{const}$; outgoing rays obey $dr/dv = \frac{1}{2}(1 - r_s/r)$, which is positive outside, zero *at* the horizon, and negative inside: the light cones tip over until, inside the horizon, even the outward branch points to smaller r , and reaching $r = 0$ becomes as compulsory as reaching next Tuesday. The maximal analytic extension (Kruskal, 1960) doubles the chart into the four regions of Figure 15; the physical spacetime formed by collapse (Oppenheimer and Snyder, 1939) contains only regions I and II, with the star’s surface replacing the white-hole sector. Penrose (1965) then proved the singularity is not an artifact of exact symmetry: once a trapped surface forms, some geodesic terminates. Incompleteness is generic (Hawking and Penrose, 1970).

The lesson of 1958–1965, which this paper asks the reader to hold on to, is grammatical: the question “is $r = r_s$ singular?” had been answered *relative to a chart*, and the chart was the problem. What is invariantly true at the horizon is not divergence but *one-way causal structure*: it is the boundary of the region from which null rays reach $\mathcal{I}^+$.




Figure 15: Penrose diagram of the maximally extended (eternal) Schwarzschild solution. Light rays run at 45° ; the horizons $\mathcal{H}^\pm$ are the diagonals through the bifurcation point; the singularities $r = 0$ are the wavy horizontal boundaries — *spacelike* surfaces, i.e. moments of time, not places in space. A physical black hole formed by collapse contains only regions I and II (the star’s interior replaces III and IV); the evaporating case, with the OGGEF port drawn in, is Figure 12 of the main text.

E.3 Step 3 — orbits, the photon sphere, the ISCO, and the shadow

Geodesics of (81) conserve $E = (1 - r_s/r) c^2 \dot{t}$ and $L = r^2 \dot{\phi}$ (dots: proper time; equatorial plane). Substituting into $g_{\mu\nu} \dot{x}^\mu \dot{x}^\nu = -c^2$ gives the radial energy equation $\frac{1}{2} \dot{r}^2 + V_{\text{eff}}(r) = \text{const}$ with

$$V_{\text{eff}}(r) = -\frac{G_N M}{r} + \frac{L^2}{2r^2} - \frac{G_N M L^2}{c^2 r^3}, \quad (86)$$

the first two terms exactly Newtonian, the third the relativistic correction that changes everything at small r (Figure 16). Setting $V'_{\text{eff}} = 0$ gives circular orbits at $r_\pm = \frac{L^2}{2G_N M} [1 \pm \sqrt{1 - 12(G_N M/cL)^2}]$; the inner root is unstable, the outer stable, and the two merge when the square root vanishes, at $L = 2\sqrt{3} G_N M/c$:

$$r_{\text{ISCO}} = \frac{6G_N M}{c^2} = 3r_s, \quad (87)$$

the innermost stable circular orbit — the inner edge of thin accretion disks and the endpoint of every quasi-circular inspiral. For photons the same computation with $ds^2 = 0$ leaves a single unstable circular orbit, the *photon sphere*,

$$r_{\text{ph}} = \frac{3G_N M}{c^2} = 1.5r_s, \quad (88)$$

and rays from infinity are captured if their impact parameter is below the critical value

$$b_c = 3\sqrt{3} \frac{G_N M}{c^2} = \sqrt{27} \frac{G_N M}{c^2} \approx 2.6r_s. \quad (89)$$

(*Physics intuition:* the $-1/r^3$ term turns the angular-momentum barrier into a cliff. In Newton, more L always means a higher wall and a safe perihelion; in Schwarzschild, any orbit that dips inside r_{ph} finds the wall replaced by a slide. The photon sphere is the crest of that slide, and the shadow is its silhouette cast on the sky.)

Worked numbers. For M87* ($M = 6.5 \times 10^9 M_\odot$, distance $D = 16.8$ Mpc), equation (89) gives a critical-curve angular diameter

$$\theta_{\text{sh}} = \frac{2b_c}{D} = \frac{2\sqrt{27} G_N M}{c^2 D} \approx 40 \mu\text{as}, \quad (90)$$

to be compared with the EHT’s measured emission-ring diameter of $42 \pm 3 \mu\text{as}$ (Event Horizon Telescope Collaboration et al., 2019) — the ring sits slightly outside the critical curve, as emission models predict. The 2022 Sgr A* image (Event Horizon Telescope Collaboration, 2022) repeats the test three orders of magnitude down in mass on the same formula. This is the sense in which Part VIII’s null tests are *geometry* tests: every number in them descends from (81) plus its Kerr generalization (Kerr, 1963).




Figure 16: The Schwarzschild effective potential (86) for two values of the specific angular momentum (curves scaled for display). The relativistic $-1/r^3$ term removes the Newtonian centrifugal barrier at small radius: for $L > 2\sqrt{3} G_N M/c$ there is an unstable maximum and a stable minimum; at $L = 2\sqrt{3} G_N M/c$ they merge into the ISCO at $r = 6G_N M/c^2$; below, nothing orbits.

E.4 Step 4 — what the classical solution assumes, and cannot answer

The Schwarzschild derivation consumed one physical input — $R_{\mu\nu} = 0$ with the Newtonian boundary condition — and it answers every question about *trajectories* outside and inside the hole. It cannot answer four questions it raises itself. (*i*) What happens at $r = 0$, where its own curvature

invariant (83) diverges and geodesics end (Penrose, 1965)? (ii) What, if anything, is the horizon made of — the surface has locally unremarkable geometry, yet globally it is an absolute one-way membrane? (iii) Why does the exterior forget everything about the star except M (and, with Kerr (1963) and Reissner and Nordström (1916), J and Q) — where does the rest of the star’s description go? (iv) Nothing in the solution is thermodynamic: no temperature, no entropy, no arrow. The next two appendices show how questions (iii) and (iv) forced themselves on the field anyway, and Appendix H shows how the quantum vacuum answers (iv) uninvited.

E.5 Step 5 — the GG-Theory reading

In the main text’s reading the Schwarzschild exterior is untouched — it is the 4D zero-mode projection of the bulk solution, exact to the measured precision (Part VII prices the corrections at $(\ell_{\text{lock}}/r_s)^2$, unobservably small). What changes is the grammar of the two singular loci. The horizon, which Finkelstein (1958) taught us is regular *in the right 4D chart*, is read as the locus where the OGGEp projection becomes causally forced — a port, not a place (§25); and the singularity at $r = 0$, which no 4D chart can cure, is read as curable by the right *bulk* chart: a small-scale-profile region of the 6D geometry, smooth in six dimensions, singular only in the projection (§38.2). Finkelstein’s lesson — what looks pathological may be your coordinates — is thus applied twice, once inside 4D and once from 6D down, and the second application is falsifiable through the same null tests the shadow computation above feeds (Part VIII).

Problems and worked solutions. E.1.1. Prove Birkhoff’s theorem: drop the staticity assumption and show the vacuum equations restore it.

Solution. Step 1: repeat the ansatz (74) with $\alpha(r, t)$, $\beta(r, t)$. The only new nonvanishing Christoffel symbols are $\Gamma_{tt}^t = \dot{\alpha}$, $\Gamma_{rr}^r = \dot{\beta} e^{2(\beta-\alpha)}/c^2$, and $\Gamma_{tr}^r = \dot{\beta}$ (dots: $\partial/\partial(ct)$). Step 2: the one component of the Ricci tensor that has no static counterpart is $R_{tr} = \frac{2}{r} \dot{\beta}$, and vacuum forces $\dot{\beta} = 0$: the spatial geometry cannot depend on time. Step 3: with $\beta = \beta(r)$, equation (78) gives $\alpha' + \beta'$ as a function of r alone, so $\alpha(r, t) = \alpha_0(r) + h(t)$; the free function $h(t)$ rescales the time coordinate, $dt \rightarrow e^{-h} dt$, and is pure gauge. What remains is exactly (81). Physics intuition: a pulsating spherical star has no monopole radiation channel, so its exterior cannot receive news of the pulsation — staticity is not assumed, it is enforced by spherical symmetry itself.

E.1.2. Derive the Newtonian correspondence $g_{tt} = -(1 + 2\Phi/c^2)$ used to fix r_s .

Solution. Step 1: for a slow particle, $dx^i/d\tau \ll c dt/d\tau$, the geodesic equation reduces to $\frac{d^2 x^i}{d\tau^2} = -\Gamma_{tt}^i (\frac{c dt}{d\tau})^2$. Step 2: for a static, weak field, $\Gamma_{tt}^i = -\frac{1}{2} \delta^{ij} \partial_j g_{tt}$ (metric derivatives are small; indices raised with η). Step 3: with $t \approx \tau$, $\frac{d^2 x^i}{dt^2} = \frac{c^2}{2} \partial_i g_{tt}$, and matching Newton, $\ddot{\vec{x}} = -\nabla\Phi$, gives $g_{tt} = -(1 + 2\Phi/c^2)$ up to a constant fixed by flatness at infinity. With $\Phi = -G_N M/r$, comparison with (81) yields $r_s = 2G_N M/c^2$ (82). Physics intuition: only the *time-time* component matters for slow motion — Newtonian gravity is time running at different rates at different heights, which is why the horizon, where $g_{tt} \rightarrow 0$, is first of all a place where the exterior’s clocks see time stop.

E.2.1. Find the black-hole mass at which an infalling body of height 2 m feels a head-to-foot tidal acceleration of one Earth gravity at the horizon crossing.

Solution. Step 1: the Newtonian estimate of the tidal (gradient) acceleration across a height h at radius r is $\Delta a = 2G_N M h / r^3$, and (83) confirms the relativistic scaling: curvature at the horizon $\propto 1/M^2$. Step 2: evaluate at $r = r_s = 2G_N M / c^2$: $\Delta a = 2G_N M h \cdot c^6 / (8G_N^3 M^3) = h c^6 / (4G_N^2 M^2)$. Step 3: set $\Delta a = g$ and solve:

$$M = \sqrt{\frac{h c^6}{4G_N^2 g}} \approx 9.1 \times 10^{34} \text{ kg} \approx 4.6 \times 10^4 M_\odot.$$

Below $\sim 5 \times 10^4 M_\odot$ the crossing is violent (a stellar-mass horizon dismembers the visitor well outside r_s); above it, the crossing is gentler than standing on Earth. Physics intuition: this single number explains why the supermassive holes are the *safe* ones and why “falling in unharmed” (Q4 of the general-reader questions) is a statement about galactic centers, not X-ray binaries.

E.3.1. Derive the photon sphere and the critical impact parameter (89) from the null effective potential.

Solution. Step 1: for null geodesics the radial equation reads $\dot{r}^2 = E^2 - L^2 W(r)$ with $W(r) = (1 - r_s/r)/r^2$ — the $-G_N M L^2 / c^2 r^3$ cliff of (86) with no Newtonian $-1/r$ term to soften it. Step 2: circular null orbits sit at $W'(r) = 0$: $-2r^{-3} + 3r_s r^{-4} = 0 \Rightarrow r_{\text{ph}} = \frac{3}{2}r_s = 3G_N M / c^2$, a maximum of W — hence unstable, (88). Step 3: a ray from infinity with impact parameter $b = L/E$ is captured iff $b^{-2} > W(r_{\text{ph}}) = \frac{1}{27} (c^2 / G_N M)^2$, i.e. $b < b_c = \sqrt{27} G_N M / c^2$ (89). Physics intuition: the photon sphere is a mountain pass for light — rays with b just above b_c orbit many times before escaping, which is why the observed ring is sharp: it is the pass photographed from below.

E.3.2. Predict the Sgr A* shadow diameter from the star-orbit mass and compare with the 2022 EHT ring.

Solution. Step 1: the interferometric star-orbit campaign gives $M = 4.30 \times 10^6 M_\odot$ at $D = 8.28$ kpc, so $G_N M / c^2 = 6.3 \times 10^9$ m. Step 2: the critical-curve angular diameter is $\theta = 2\sqrt{27} G_N M / (c^2 D) \approx 53 \mu\text{as}$. Step 3: the EHT ring (Event Horizon Telescope Collaboration, 2022) is $51.8 \pm 2.3 \mu\text{as}$ — agreement at the few-percent level, with the residual within the known offset between the emission ring and the critical curve (cf. the M87* case worked in Step 3 of the text). Physics intuition: the mass in Step 1 comes from Keplerian star orbits at $\sim 10^3 r_s$; the ring tests the *same* mass at $\sim 2 r_s$. Agreement across three orders of magnitude in radius is a stronger statement than either number alone.

E.3.3. Show that the ISCO binding energy is $1 - \sqrt{8/9}$ of the rest mass, and interpret astrophysically.

Solution. Step 1: for a circular orbit at radius r , conserved energy per unit rest mass is $\tilde{E} = (1 - r_s/r) / \sqrt{1 - \frac{3}{2}r_s/r}$ (combine $V'_{\text{eff}} = 0$ with the normalization $g_{\mu\nu} \dot{x}^\mu \dot{x}^\nu = -c^2$). Step 2: at the ISCO (87), $r = 3r_s$: $\tilde{E} = (1 - \frac{1}{3}) / \sqrt{1 - \frac{1}{2}} = \frac{2}{3}\sqrt{2} = \sqrt{8/9} \approx 0.9428$. Step 3: the binding energy released by slow inspiral to the ISCO is therefore $1 - \sqrt{8/9} \approx 5.7\%$ of mc^2 — against 0.7% for hydrogen fusion. Physics intuition: this factor-of-eight edge over nuclear burning is why accreting black holes power the most luminous persistent sources in the sky, and why quasars were the first observational hint that collapsed objects exist at all. (For Kerr the ISCO moves inward and the efficiency rises further — Appendix P.)

Take-away. *Standard result:* the unique static spherical vacuum solution (81), with a regular one-way horizon at $r_s = 2G_N M/c^2$, a genuine singularity at $r = 0$, photon sphere at $1.5 r_s$, ISCO at $3 r_s$, and shadow diameter $2\sqrt{27} G_N M/c^2$, confirmed at M87* and Sgr A*. *What it assumed:* classical vacuum general relativity. *What it cannot answer:* the fate of $r = 0$, the nature of the horizon, where the star’s information goes, and why any of it should be thermodynamic. *The GG-Theory reading:* the exterior survives untouched as a projection; the horizon becomes a port; the singularity becomes a chart artifact one level up.

F Classical black-hole mechanics: the four laws

What this appendix teaches. The four laws of classical black-hole mechanics, derived rather than tabulated: the surface gravity and its constancy (zeroth law), the first law $\delta M = \kappa \delta A c^2 / (8\pi G_N)$ verified explicitly for Schwarzschild, the Raychaudhuri equation and Hawking’s area theorem (second law) with the GW150914 area audit as the worked example, and the third law — closed by the Bardeen–Carter–Hawking analogy table and the honest 1973 question the analogy could not answer: the temperature of what? The entropy of which states?

Read alongside: Poisson (2004) chs. 2–5 (congruences and horizon mechanics, the modern standard); Wald (1984) ch. 12; Bardeen et al. (1973) for the original.

By 1972 the classical theory of the previous appendix had matured into a set of theorems with an unsettling shape: they were, line for line, the laws of thermodynamics with (M, κ, A) in place of (E, T, S) (Bardeen et al., 1973). This appendix derives the four laws and then states precisely why, in 1973, the resemblance had to be called an analogy — the question the rest of the course answers.

F.1 Step 1 — surface gravity and the zeroth law

For a stationary horizon generated by the Killing field ξ^μ , the surface gravity κ is defined on the horizon by $\xi^\nu \nabla_\nu \xi^\mu = (\kappa/c) \xi^\mu$: the failure of the generator to be affinely parametrized. Operationally, κ is the force per unit mass that an observer at infinity must supply, through an idealized massless string, to hold a unit mass just above the horizon. For Schwarzschild the computation is three lines: with $f(r) = 1 - r_s/r$, the local proper acceleration of a static observer is $a(r) = \frac{G_N M}{r^2} f^{-1/2}$, which diverges at the horizon; the redshift factor is $f^{1/2}$; their product stays finite,

$$\kappa = \lim_{r \rightarrow r_s} f^{1/2} a(r) = \frac{G_N M}{r_s^2} = \frac{c^4}{4G_N M} = \frac{c^2}{2r_s}. \quad (91)$$

(*Physics intuition:* the local force diverges but the string’s mechanical advantage — the redshift — discounts it by exactly the diverging factor. κ is the price of holding position *as billed at infinity*, which is why it, and not the local acceleration, will turn out to set the temperature seen at infinity.) The *zeroth law* states that κ is constant over the entire horizon of any stationary black hole — even a rotating, flattened one, where the local geometry varies from pole to equator. Constancy of an intensive quantity over a system in equilibrium is the defining mark of a temperature, and this is the first rung of the analogy.

F.2 Step 2 — the first law, verified

For Schwarzschild, $A = 4\pi r_s^2 = 16\pi G_N^2 M^2 / c^4$, so

$$\delta A = \frac{32\pi G_N^2 M}{c^4} \delta M \implies \frac{\kappa c^2}{8\pi G_N} \delta A = \frac{c^4}{4G_N M} \cdot \frac{c^2}{8\pi G_N} \cdot \frac{32\pi G_N^2 M}{c^4} \delta M = c^2 \delta M, \quad (92)$$

which is the *first law*

$$c^2 \delta M = \frac{\kappa c^2}{8\pi G_N} \delta A \longrightarrow c^2 \delta M = \frac{\kappa c^2}{8\pi G_N} \delta A + \Omega_H \delta J + \Phi_H \delta Q \quad (93)$$

in its full Kerr–Newman form, derived together with the Smarr formula in Appendix P (Step 3). The structure is exactly $dE = T dS + \text{work terms}$, with κ sitting where T sits and A sitting where S sits.

F.3 Step 3 — the Raychaudhuri equation and the area theorem

The second law is the deep one, and it needs one tool. Consider the congruence of null geodesics generating the horizon, with tangent k^μ , affine parameter λ , and expansion $\theta = \nabla_\mu k^\mu$ (the fractional rate of change of the cross-sectional area: $\theta = \frac{1}{\delta A} \frac{d\delta A}{d\lambda}$). Decomposing $\nabla_\nu k_\mu$ into trace, symmetric-traceless (shear $\sigma_{\mu\nu}$), and antisymmetric (twist, zero for hypersurface generators) parts, the geodesic equation and the Ricci identity give, in four lines of index algebra, the *Raychaudhuri equation*

$$\frac{d\theta}{d\lambda} = -\frac{1}{2}\theta^2 - \sigma_{\mu\nu}\sigma^{\mu\nu} - R_{\mu\nu}k^\mu k^\nu. \quad (94)$$

Every term on the right is ≤ 0 provided $R_{\mu\nu}k^\mu k^\nu \geq 0$, which by the Einstein equations is the *null energy condition* $T_{\mu\nu}k^\mu k^\nu \geq 0$ — satisfied by all classical matter. (*Physics intuition*: gravity focuses. Shear focuses. The only thing that could defocus a pencil of light is negative energy density, and classical physics has none to offer.) Suppose now some horizon generator had $\theta < 0$ at one point. Integrating $d\theta/d\lambda \leq -\theta^2/2$ gives $\theta \rightarrow -\infty$ within finite affine parameter $\lambda \leq 2/|\theta_0|$: the generators cross, forming a caustic — but a theorem of Penrose forbids horizon generators from leaving the horizon, and crossing generators would have to. Contradiction; hence $\theta \geq 0$ everywhere on any future horizon, and therefore

$$\frac{dA}{dt} \geq 0 \quad (95)$$

in every classical process (Hawking and Penrose, 1970). Area, like entropy, never decreases.

Worked instance — the first observational audit. In GW150914 (Abbott et al., 2016), holes of 36 and $29 M_\odot$ (initial spins consistent with small) merged into $62 M_\odot$ with final spin $a_* \approx 0.67$. Using the Kerr area $A = 8\pi(G_N M/c^2)^2(1 + \sqrt{1 - a_*^2})$: initial area $A_i \approx 2.3 \times 10^5 \text{ km}^2$ (the two Schwarzschild areas), final area $A_f \approx 3.7 \times 10^5 \text{ km}^2$ — an increase of $\sim 57\%$, comfortably positive even while $3 M_\odot c^2$ of energy left the system as gravitational waves. Mass can be spent; area cannot. This is now a live experimental program (the ringdown-era area-theorem tests of Part VIII).

F.4 Step 4 — the third law, and the analogy table

The *third law* states that no finite sequence of physical processes can reduce κ to zero — extremality ($a_* \rightarrow 1$, or $Q \rightarrow M$ in geometrized units) is asymptotic, like absolute zero. The worked Kerr–Newman version, with the capture-orbit argument, is Appendix P (Step 4). Assembled, the four laws sit in exact formal correspondence with thermodynamics:

Table 19: The Bardeen–Carter–Hawking correspondence (Bardeen et al., 1973), as it stood in 1973: a formal analogy with, at that date, no license to set the constants α . The table is reprinted as it stood in 1973, before Hawking’s calculation, and the point is the one its authors made themselves: at that date the correspondence was formal. Nothing fixed the constants, and the analogy could have been a coincidence of algebra. What the sections that follow supply is the license the 1973 paper explicitly declined to claim.

Law	Thermodynamics	Black-hole mechanics
Zeroth	T constant throughout a body in equilibrium	κ constant over a stationary horizon
First	$dE = T dS + \text{work}$	$c^2 dM = \frac{\kappa c^2}{8\pi G_N} dA + \Omega_H dJ + \Phi_H dQ$
Second	$\delta S \geq 0$	$\delta A \geq 0$ (Hawking’s area theorem)
Third	$T = 0$ unattainable in finite operations	$\kappa = 0$ unattainable in finite operations

F.5 Step 5 — what the analogy could not explain

In 1973 the correspondence was called, by its own authors, a formal analogy — and they were right to be careful, because taken literally it was inconsistent. A body with a temperature radiates; a classical black hole absorbs everything and emits nothing, so its temperature is exactly zero, while $\kappa \neq 0$. An entropy counts microstates; the no-hair theorems say a stationary hole has *one* exterior state per (M, J, Q) , so its entropy should be $k_B \ln 1 = 0$, while $A \neq 0$. The analogy demanded constants of proportionality — $T = \alpha\kappa$, $S = A/(8\pi G_N \alpha)$ for some α — that classical physics not only failed to supply but positively forbade. Two years later Hawking’s quantum computation (Appendix H) supplied $\alpha = \hbar/(2\pi c k_B)$ uninvited, and the analogy became physics with one stroke — leaving as its residue the question this paper’s main text addresses: an entropy was now *priced*, but what does it *count*?

F.6 Step 6 — the GG-Theory reading

In the main text’s ledger, each law acquires a mechanism. The zeroth law is the KMS face of the port clock: one thermal clock per stationary port, hence one κ (§28). The first law is the differential form of the GFF ledger (43) on the stationary branch. The second law is promoted from the null energy condition to the generalized second law — now including the radiation’s entropy, which the classical theorem cannot see — with its production computed exactly in the exchange class (Appendix O, Step 3, after Wall, 2012). And the third law becomes the statement that the port never quite closes (Appendix P, Step 4). The analogy of 1973, in other words, is read as a dictionary fragment: the boundary bookkeeping of an open system whose bulk ledger the classical

theory could name but not open.

Problems and worked solutions. F.1.1. Derive the static observer’s proper acceleration

$a(r) = G_N M r^{-2} f^{-1/2}$ **used in (91).**

Solution. Step 1: the static worldline has four-velocity $u^\mu = (c f^{-1/2}, 0, 0, 0)$ (normalized: $g_{tt} u^t u^t = -c^2$). Step 2: the acceleration is $a^\mu = u^\nu \nabla_\nu u^\mu$; the only nonvanishing component is $a^r = \Gamma_{tt}^r (u^t)^2 = (\frac{c^2}{2} f f') \cdot f^{-1} = \frac{c^2}{2} f'$. Step 3: the invariant magnitude is $a = \sqrt{g_{rr}} a^r = \frac{c^2}{2} f' f^{-1/2} = \frac{G_N M}{r^2} f^{-1/2}$, diverging at the horizon. Multiplying by the redshift $f^{1/2}$ and taking the limit gives $\kappa = G_N M / r_s^2 = c^4 / 4 G_N M$ (91). Physics intuition: the local engine must work infinitely hard to hover, but every joule it reports is infinitely redshifted on the way out; the finite product is the only number the far bookkeeper ever sees — which is precisely why it, and not the local force, becomes the temperature.

F.2.1. From the first law and $S \propto A$ alone, derive the mass-scaling of the temperature — before any quantum computation.

Solution. Step 1: $A = 16\pi G_N^2 M^2 / c^4 \Rightarrow A \propto M^2$, so if $S = \eta k_B A c^3 / (G_N \hbar)$ then $S \propto M^2$. Step 2: the thermodynamic identity $T^{-1} = \partial S / \partial E$ with $E = M c^2$ gives $T^{-1} \propto M$, i.e. $T \propto 1/M$: big holes are cold. Step 3: restoring the constants, $T = \hbar c^3 / (16\pi \eta G_N M k_B \cdot 2)$ — with the coefficient left as η , exactly the state of knowledge of 1973. Physics intuition: the negative specific heat of Appendix I (Step 3) is already visible here — T falls as E grows — so the analogy predicted gravity’s strangest thermodynamic signature before anyone believed the temperature was real.

F.2.2. Show that the classical analogy fixes only the product of the two unknown constants.

Solution. Step 1: write $T = \alpha \kappa$ and $S = \gamma A$ for undetermined α, γ . Step 2: the first law (93) demands $T dS = \frac{\kappa c^2}{8\pi G_N} dA$, i.e. $\alpha \gamma = c^2 / (8\pi G_N)$: one equation, two unknowns. Step 3: no classical experiment can split the product — a rescaling $\alpha \rightarrow \lambda \alpha$, $\gamma \rightarrow \gamma / \lambda$ changes nothing observable while the hole neither radiates nor counts states. Hawking’s computation (Appendix H) fixes $\alpha = \hbar / (2\pi c k_B)$, and $\gamma = k_B c^3 / (4 G_N \hbar)$ follows *instantly* from the 1973 product. Physics intuition: this is why the temperature and the entropy coefficient were nailed by a single quantum calculation — the classical first law had already welded them together.

F.3.1. Fill in the derivation of the Raychaudhuri equation (94).

Solution. Step 1: define $B_{\mu\nu} = \nabla_\nu k_\mu$ for the null congruence; along each geodesic, $k^\rho \nabla_\rho B_{\mu\nu} = -B_{\mu\rho} B^\rho_\nu - R_{\rho\nu\mu\sigma} k^\rho k^\sigma$, which follows from the Ricci identity $[\nabla_\rho, \nabla_\nu] k_\mu = R_{\rho\nu\mu}{}^\sigma k_\sigma$ plus the geodesic equation. Step 2: trace over $\mu\nu$: with $\theta = B^\mu_\mu$, $\frac{d\theta}{d\lambda} = -B_{\mu\nu} B^{\nu\mu} - R_{\mu\nu} k^\mu k^\nu$. Step 3: decompose the (projected, transverse) $B_{\mu\nu} = \frac{1}{2} \theta h_{\mu\nu} + \sigma_{\mu\nu} + \omega_{\mu\nu}$ in the two-dimensional screen space; then $B_{\mu\nu} B^{\nu\mu} = \frac{1}{2} \theta^2 + \sigma^2 - \omega^2$ (the $\frac{1}{2}$ because the screen is two-dimensional), and $\omega = 0$ for horizon generators (hypersurface-orthogonal). Equation (94) follows. Physics intuition: the equation is pure kinematics plus one import — geometry only enters through $R_{\mu\nu} k^\mu k^\nu$, which Einstein’s equations convert into energy density. “Gravity focuses light” is the entire physical content, and the area theorem is its global echo.

F.3.2. For the GW150914 masses, find the maximum energy the area theorem would have permitted the merger to radiate, and compare with the observed $3 M_\odot c^2$.

Solution. Step 1: the final area must satisfy $A_f \geq A_i$, and for a (nearly non-spinning) final hole $A \propto M^2$, so $M_f \geq \sqrt{36^2 + 29^2} M_\odot = 46.2 M_\odot$: the irreducible mass of the pair. Step 2: the radiated energy is bounded by $E_{\text{rad}} \leq (65 - 46.2) M_\odot c^2 = 18.8 M_\odot c^2$ — an efficiency ceiling of 29%. Step 3: the observed $3 M_\odot c^2$ (4.6%) sits comfortably inside: the theorem constrains, the dynamics decides. (A spinning final hole lowers the ceiling slightly; the worked audit in Step 3 of the text used the measured $a_* = 0.67$.) Physics intuition: the area theorem is gravity’s second law acting as an energy regulator — mergers are the most luminous events in nature *and* they are throttled, both facts from one inequality.

F.5.1. Make the 1973 inconsistency quantitative: show that a cold classical hole beside a warm bath violates the ordinary second law if the analogy is read literally.

Solution. Step 1: give the hole any temperature $T_{\text{BH}} > 0$ by the analogy, and place it in a radiation bath at T_b with $0 < T_b < T_{\text{BH}}$. Step 2: classically the hole absorbs the bath’s radiation and emits nothing, so heat flows *from the colder bath into the hotter hole*; $dS_{\text{total}} = dQ (T_{\text{BH}}^{-1} - T_b^{-1}) < 0$ if S_{BH} is read via the first law. Step 3: conclusion — of the triple (nonzero T , no emission, first law), one member must go. Hawking’s radiation removes “no emission” and the system becomes consistent; the alternative surgeries (deny T , deny the first law) each destroy the analogy wholesale. Physics intuition: this is the cleanest way to see that Hawking radiation was *forced* — thermodynamics had already decided horizons must shine; the quantum computation only revealed the mechanism and the coefficient.

Take-away. *Standard result:* the four laws (Table 19), each a theorem of classical general relativity, jointly isomorphic to thermodynamics. *What they assumed:* the null energy condition and classical matter. *What they could not explain:* the constants of the correspondence — a nonzero temperature for a perfect absorber and a nonzero entropy for a one-state object were classical contradictions. *The GG-Theory reading:* the laws are the boundary bookkeeping of the port; the missing constants arrive with the quantum (Appendix H), and the missing microstates with the projection (§31.6).

G Bekenstein’s entropy and the generalized second law

What this appendix teaches. The argument that put entropy on the horizon before any quantum computation existed: Wheeler’s teacup, Bekenstein’s 1973 proposal $S_{\text{BH}} = \eta k_B A c^3 / (G_N \hbar)$ with the coefficient η honestly unknown, the Bekenstein bound $S \leq 2\pi k_B E R / (\hbar c)$ derived from the same gedanken experiment that motivated it, the generalized second law, and the species problem — with worked numbers showing how far ordinary matter sits from saturation, and why only horizons saturate.

Read alongside: [Bekenstein \(1973\)](#) for the original; [Wald \(1994\)](#) for the modern formulation; [Susskind and Lindesay \(2005\)](#) for the physical picture.

Between the four laws of the previous appendix (1973 — a formal analogy) and Hawking’s radiation (1975 — the analogy made physical), there is a short, decisive chapter that is often skipped in the historical telling and never in the logical one: Bekenstein’s argument that horizon area *is* entropy, made while the field still believed black holes were cold and dead.

G.1 Step 1 — Wheeler’s teacup: the crisis

The provocation, as Wheeler put it to his student, is domestic: I set a cup of hot tea beside a cold one and let them equilibrate; I have committed a crime against the second law — the entropy of the universe went up and I can never bring it back down. But if I drop both cups into a passing black hole, the evidence is gone. The exterior state of the hole is fixed by (M, J, Q) alone (no-hair, Appendix E); no observation of the hole will ever reveal how much entropy it swallowed. If S_{outside} is all the entropy there is, the second law is not violated but worse: unfalsifiable, dying “the death of a thousand qualifications.” (*Physics intuition*: the second law is bookkeeping, and a black hole with no entropy column is an account that shreds receipts. Either the law fails, or the hole keeps books we had not learned to read.)

G.2 Step 2 — Bekenstein’s 1973 proposal

Bekenstein’s move (Bekenstein, 1973) was to take the area theorem of Appendix F not as an analogy to the second law but as its missing column. Three observations converge. (i) A never decreases classically — the right monotonicity. (ii) When a hole swallows matter, A increases by at least a calculable minimum — the right responsiveness. (iii) Dimensional analysis: an entropy is k_B times a pure number, and the only way to make a pure number from A , the constants of gravity, and the quantum is

$$S_{\text{BH}} = \eta \frac{k_B A c^3}{G_N \hbar} = \eta \frac{k_B A}{\ell_P^2}, \quad \ell_P^2 = \frac{G_N \hbar}{c^3}, \quad (96)$$

with η a dimensionless constant of order unity that *no classical argument can fix* — Bekenstein’s own estimate was $\eta = \ln 2 / (8\pi)$, from counting one bit per photon dropped in at the minimum increment. Note what the quantum is doing here: $\hbar$ enters (96) through dimensional necessity *two years before* any quantum calculation about black holes existed. An entropy proportional to *area*, not volume, was also unprecedented — the first sighting of what ’t Hooft (1993) and Susskind (1995) would later name the holographic principle (Appendix L).

G.3 Step 3 — the Bekenstein bound, derived

The same gedanken experiment sharpens into a bound on *ordinary* matter. Lower a box of energy E (as measured at infinity when far away) and proper size R slowly toward the horizon on an idealized string. At blueshift factor $f^{1/2}(r)$, the energy the box delivers to the hole if released is $E f^{1/2}$ — the string has extracted the rest as work at infinity. The closest the box’s center can approach is one box radius, $\delta \approx R$, where (using the near-horizon form $f^{1/2} \approx \kappa \delta / c^2$ from Appendix F, Step 1)

$$E_{\text{delivered}} \geq E \frac{\kappa R}{c^2}. \quad (97)$$

The first law then prices the area increase: using $c^2 \delta M = E_{\text{delivered}}$ in (93),

$$\delta A = \frac{8\pi G_N E_{\text{delivered}}}{\kappa c^2} \geq \frac{8\pi G_N E R}{c^4}, \quad (98)$$

independent of κ — the hole’s details cancel. If the generalized second law (Step 4) is to survive the box’s disappearance, the entropy carried in, S , must not exceed the bookkeeping credit $\delta S_{\text{BH}} = \eta k_B \delta A c^3 / (G_N \hbar)$; with the coefficient $\eta = 1/4$ that Appendix H will supply, the requirement is the *Bekenstein bound*

$$S \leq \frac{2\pi k_B ER}{\hbar c}. \quad (99)$$

(The quantum buoyancy of the box in the horizon’s thermal atmosphere, pointed out by Unruh and Wald, modifies the argument but not the bound’s status for weakly gravitating systems.)

Worked numbers. For a human brain ($E \approx 1.4 \text{ kg} \times c^2$, $R \approx 0.1 \text{ m}$) the bound allows $\sim 4 \times 10^{42}$ bits — some thirty orders of magnitude above any realistic information content: ordinary matter is entropically *dilute*. Saturation requires $S \sim k_B ER / \hbar c$, and comparing (96) with $E = Mc^2$, $R = r_s$ shows the unique saturating object is the black hole itself. (*Physics intuition:* the bound says information needs energy and room. A black hole is what you get when you pack information as densely as causality permits — which is why, in Part IV’s language, its boundary is pure port: nothing about the packing is visible except the area of the wrapping.)

G.4 Step 4 — the generalized second law, and the species problem

Bekenstein’s repair of Wheeler’s crime is the *generalized second law* (GSL):

$$\delta(S_{\text{outside}} + S_{\text{BH}}) \geq 0, \quad (100)$$

the statement that the teacup’s entropy is not destroyed but transferred to the horizon column of the ledger. Classically it holds because $\delta A \geq 0$ dominates; quantum mechanically it survives Hawking evaporation — where A *decreases* — because the radiation’s entropy more than compensates, a theorem proved in generality by Wall (2012) and computed exactly for the exchange class in Appendix O, Step 3. One structural worry surfaced immediately, and it is the right worry: the *species problem* (Susskind and Uglum, 1994). The bound and the GSL involve η/G_N ; entanglement entropy across the horizon grows with the number of matter fields; yet S_{BH} knows nothing of the particle spectrum. Either the coefficient is fine-tuned against the field content, or G_N itself is built from the field content — the induced-gravity option, which is precisely the mechanism the main text adopts (§31.3) and whose remaining open computation is Appendix R.

G.5 Step 5 — what the argument assumed, and the GG-Theory reading

Bekenstein’s 1973 case rested on consistency, not mechanism: the coefficient η was unknown (Hawking’s calculation was needed to fix $\eta = 1/4$), the microstates were unnamed, and the area-scaling was a scandal without an explanation. In the main text’s reading, all three residues are addressed in order: the coefficient is derived, regulator-independently, from induced gravity plus the conical replica (§31.5, Theorem 1); the microstates are named — $U(1)_{Q_i}$ holonomy configurations of the hidden fiber, one register per Planck-area pixel (§31.6); and the area-scaling is the bandwidth of the port — entropy lives on the boundary because the boundary is where the projection truncates the bulk (§25). The species problem becomes the program’s one declared open computation

(Appendix R).

Problems and worked solutions. G.2.1. Prove the dimensional-analysis uniqueness claim behind (96).

Solution. Step 1: seek $S = k_B \Xi$ with Ξ a dimensionless function of A , G_N , $\hbar$, c (mass does not appear separately: for Schwarzschild, A carries it). Step 2: from G_N , $\hbar$, c exactly one length can be built, $\ell_P = \sqrt{G_N \hbar / c^3}$, so the only dimensionless combination involving A is A/ℓ_P^2 (and functions thereof). Step 3: additivity fixes the function: for two distant holes, entropies add and areas add, so $\Xi(A_1 + A_2) = \Xi(A_1) + \Xi(A_2)$, forcing linearity, $\Xi = \eta A/\ell_P^2$. Physics intuition: the quantum enters through ℓ_P before any quantum *dynamics* is known — dimensional analysis alone says that if a horizon has entropy at all, $\hbar$ is in it. That is why Bekenstein could be right in 1973 with no mechanism.

G.3.1. Show that the Schwarzschild hole saturates the Bekenstein bound (99) exactly.

Solution. Step 1: insert $E = Mc^2$ and $R = r_s = 2G_N M/c^2$ into the bound: $S_{\max} = 2\pi k_B Mc^2 \cdot 2G_N M/(\hbar c^2) = 4\pi k_B G_N M^2/(\hbar c)$. Step 2: evaluate $S_{\text{BH}} = k_B A c^3/(4G_N \hbar)$ with $A = 16\pi G_N^2 M^2/c^4$: $S_{\text{BH}} = 4\pi k_B G_N M^2/(\hbar c)$. Step 3: the ratio is exactly 1 — no factors of 2 survive. Physics intuition: the bound is a property of causality and thermodynamics; the black hole is the object that converts the inequality into an identity. Saturation is the information-theoretic face of the no-hair property: nothing about the hole is wasted on internal structure a boundary observer could have resolved.

G.3.2. How far below the bound does a real storage device sit?

Solution. Step 1: for a 1-kg, 10-cm device the bound allows

$$\frac{S_{\max}}{k_B \ln 2} = \frac{2\pi (1 \text{ kg}) c^2 (0.1 \text{ m})}{\hbar c \ln 2} \approx 2.6 \times 10^{42} \text{ bits.}$$

Step 2: a contemporary solid-state drive of that size stores $\sim 10^{13}$ bits. Step 3: the gap is twenty-nine orders of magnitude — and even the device’s fundamental entropy content ($\sim 10^{25}$ atomic degrees of freedom) sits seventeen orders below the ceiling. Physics intuition: ordinary matter stores information in atoms; the bound prices information in *gravity’s* currency, and the exchange rate ($\sim 10^{19}$ GeV vs. eV chemistry) is the whole gap. Only when self-gravity takes over does storage approach the ledger’s limit — and the object that achieves it is, by G.3.1, a black hole.

G.4.1. Verify the GSL for pure Hawking emission: show the radiation gains more entropy than the hole loses.

Solution. Step 1: the hole loses $dS_{\text{BH}} = dE/T_{\text{H}}$ (first law, $dE < 0$). Step 2: the radiation is (nearly) blackbody at T_{H} : free-streaming photons carry entropy flux $\frac{4}{3}(aT^3)$ against energy flux aT^4 , so $dS_{\text{rad}} = \frac{4}{3}|dE|/T_{\text{H}}$. Step 3: the balance is $dS_{\text{gen}} = (\frac{4}{3} - 1)|dE|/T_{\text{H}} = \frac{1}{3}|dE|/T_{\text{H}} > 0$: the generalized second law (100) holds with a strict margin (graybody filtering shifts the 4/3 at the percent level but never below 1). Physics intuition: evaporation is not the hole “leaking order” — it is the conversion of one large, perfectly packed ledger page into many small thermal ones, at a fixed markup of $\frac{1}{3}$; the exact exchange-class version of this bookkeeping is Appendix O, Step 4.

Take-away. *Standard result:* $S_{\text{BH}} = \eta k_B A / \ell_P^2$ on consistency grounds, the Bekenstein bound (99), and the GSL (100). *What it assumed:* the second law is worth saving, and dimensional analysis. *What it could not supply:* the coefficient, the microstates, and a mechanism for area-scaling. *The GG-Theory reading:* coefficient from the conical replica, microstates from the fiber holonomies, area-scaling from port bandwidth — with the species sum left open, honestly, as Appendix R.

H Explicit derivation of Hawking temperature in GDOS

What this appendix teaches. The Hawking temperature derived from first principles with the full Bogoliubov machinery: quantum field theory on a fixed curved background set up from scratch, the Rindler warm-up solved exactly, the two routes to the surface gravity, the restoration of units, and the GDOS reading in which the continuation to the Euclidean section stops being a trick and becomes a licensed theorem.

Read alongside: Birrell and Davies (1982) chs. 3 and 8 (the canonical treatment this Step follows most closely); Hawking (1975) and Unruh (1976) for the originals; Parker and Toms (2009) at depth.

This chapter is the course’s first worked calculation: the Hawking temperature, derived from first principles with every intermediate manipulation displayed. The main text asserted the result three times over — the Bogoliubov ratio of §26.4, the clock-face theorem T-BH-5 (§28), and the Euclidean period (Appendix I); here a graduate student meets the machinery those assertions run on. The plan: Step 1 sets up quantum field theory on a fixed curved background from scratch; Step 2 builds the Bogoliubov machinery; Step 3 works the flat-space warm-up with an exact answer (the Unruh effect); Step 4 transfers it to the black hole by two independent routes; Step 5 restores units and evaluates the numbers; Step 6 states what GDOS adds. Nothing here is GG-Theory-specific until Step 6: Steps 1–5 are the standard material (Hawking, 1975; Unruh, 1976; Wald, 1984), worked in the conventions of this paper.

H.1 Step 1 — the stage: a scalar field on a fixed curved background

Take a real scalar field on a globally hyperbolic spacetime $(\mathcal{M}, g_{\mu\nu})$ with signature $(-, +, +, +)$, action

$$S[\phi] = -\frac{1}{2} \int d^4x \sqrt{-g} (g^{\mu\nu} \partial_\mu \phi \partial_\nu \phi + m^2 \phi^2 + \xi R \phi^2), \quad (101)$$

whose variation gives the curved-space Klein–Gordon equation

$$(\square - m^2 - \xi R)\phi = 0, \quad \square \phi \equiv \frac{1}{\sqrt{-g}} \partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \phi). \quad (102)$$

On any Cauchy surface Σ with future-directed unit normal n^μ and induced volume element $\sqrt{h} d^3x$, define the Klein–Gordon inner product

$$(u, v) = -i \int_\Sigma d^3x \sqrt{h} n^\mu (u \partial_\mu v^* - v^* \partial_\mu u). \quad (103)$$

Two properties carry the whole subject and are worth verifying once by hand: (a) for solutions of the field equation, (103) is independent of the choice of Σ (the current $j^\mu = -i(u \partial^\mu v^* - v^* \partial^\mu u)$ is

conserved on shell, $\nabla_\mu j^\mu = 0$, by the field equation applied twice); (b) it is not positive definite: $(u^*, u^*) = -(u, u)$. Property (b) is the origin of everything that follows — the space of solutions splits into a positive-norm half and its conjugate, and *which* split is chosen is precisely the choice of vacuum.

Pick any complete set of positive-norm solutions $\{u_i\}$ with $(u_i, u_j) = \delta_{ij}$, $(u_i^*, u_j^*) = -\delta_{ij}$, $(u_i, u_j^*) = 0$, and expand

$$\phi(x) = \sum_i (a_i u_i(x) + a_i^\dagger u_i^*(x)). \quad (104)$$

The canonical commutation relations for ϕ are then equivalent to $[a_i, a_j^\dagger] = \delta_{ij}$, and the vacuum belonging to this basis is $a_i|0_u\rangle = 0$ for all i . In Minkowski space the modes $u_{\mathbf{k}} \propto e^{i\mathbf{k}\cdot\mathbf{x} - i\omega_{\mathbf{k}}t}$, positive frequency with respect to inertial time, are selected by Poincaré invariance and every inertial observer agrees. On a general curved background there is no preferred time, hence no preferred split, hence no preferred particle notion — only preferred *observers* can select one.

Physical picture. A vacuum is not an empty box; it is a bookkeeping convention tied to a clock. “Positive frequency” means “rotating one way with respect to *this* time parameter.” Flat space owns a global inertial clock, so all inertial bookkeepers agree on what counts as a particle. A curved spacetime in general owns several inequivalent clocks — and the whole content of horizon thermodynamics is what happens when two lawful bookkeepers compare ledgers. In the cohort’s language the split is the derived complex structure of the causal clock (Arisaka, 2026, P4-Real); T-BH-5 records when that selection is unique.

H.2 Step 2 — the Bogoliubov machinery

Let $\{v_j\}$ be a second complete positive-norm basis — another lawful bookkeeper. Completeness of $\{u_i, u_i^*\}$ lets us expand

$$v_j = \sum_i (\alpha_{ji} u_i + \beta_{ji} u_i^*), \quad \alpha_{ji} = (v_j, u_i), \quad \beta_{ji} = -(v_j, u_i^*), \quad (105)$$

where the coefficient formulas follow by taking the inner product of the expansion with u_i and u_i^* and using the orthonormality relations of Step 1. Substituting (105) into the field expansion and matching coefficients of u_i and u_i^* gives the operator transformation

$$b_j = \sum_i (\alpha_{ji}^* a_i - \beta_{ji}^* a_i^\dagger), \quad (106)$$

and demanding $[b_j, b_k^\dagger] = \delta_{jk}$ yields the consistency conditions

$$\sum_i (\alpha_{ji} \alpha_{ki}^* - \beta_{ji} \beta_{ki}^*) = \delta_{jk}, \quad \sum_i (\alpha_{ji} \beta_{ki} - \beta_{ji} \alpha_{ki}) = 0. \quad (107)$$

The payoff line of the whole subject is one trace: the number of v -quanta in the u -vacuum is

$$\langle 0_u | b_j^\dagger b_j | 0_u \rangle = \sum_i |\beta_{ji}|^2. \quad (108)$$

Derivation: insert (106); the aa and $a^\dagger a^\dagger$ terms annihilate the vacuum on one side or the other; the surviving term is $\sum_{i,i'} \beta_{ji} \beta_{ji'}^* \langle 0_u | a_i a_{i'}^\dagger | 0_u \rangle = \sum_i |\beta_{ji}|^2$.

Physical picture. $\beta \neq 0$ means one bookkeeper's annihilation operator contains the other's creation operator: the two ledgers do not just relabel the particles, they disagree about how many there are. Particle number is dictionary data, not bulk data — the same custody division the cohort enforces everywhere (the dictionary owns the sectors), met here in its oldest habitat.

H.3 Step 3 — the warm-up with an exact answer: the accelerated observer

Everything essential happens already in flat spacetime. Work in 1+1 dimensions (the transverse directions are spectators) with null coordinates $U = t - x$, $V = t + x$, and consider the right Rindler wedge $x > |t|$, the region accessible to a uniformly accelerated observer with proper acceleration a . Rindler coordinates (η, χ) ,

$$U = -\frac{1}{a} e^{-a(\eta-\chi)}, \quad V = \frac{1}{a} e^{a(\eta+\chi)}, \quad (109)$$

cover the wedge, and the boost Killing vector ∂_η is the accelerated observer's time. Outgoing right-wedge modes, positive frequency in η , depend on the single combination $u_R \equiv \eta - \chi = -\frac{1}{a} \ln(-aU)$ (valid for $U < 0$):

$$u_\omega^R \propto e^{-i\omega u_R} = (-aU)^{i\omega/a}, \quad U < 0, \quad (110)$$

with a mirror family $u_\omega^L \propto (aU)^{-i\omega/a}$ supported at $U > 0$ (the left wedge). Each family vanishes on half the U axis, so neither is analytic at $U = 0$.

Now the decisive criterion. A mode is positive frequency with respect to *inertial* time iff it is a superposition of $e^{-i\Omega U}$ with $\Omega > 0$ only — equivalently, iff it is analytic and bounded in the *lower half* of the complex U plane (close the Fourier contour there). Continue (110) from $U < 0$ to $U > 0$ through the lower half-plane: along that path $\arg(-U)$ runs from 0 to $-\pi$, so

$$(-aU)^{i\omega/a} \longrightarrow (e^{-i\pi} aU)^{i\omega/a} = e^{\pi\omega/a} (aU)^{i\omega/a} = e^{\pi\omega/a} (u_\omega^L)^*. \quad (111)$$

The factor $e^{\pi\omega/a}$ is the entire effect: continuation through the horizon point $U = 0$ costs one exponential. Therefore the two combinations

$$w_\omega^{(1)} = \frac{e^{\pi\omega/2a} u_\omega^R + e^{-\pi\omega/2a} u_\omega^{L*}}{\sqrt{2 \sinh(\pi\omega/a)}}, \quad w_\omega^{(2)} = \frac{e^{\pi\omega/2a} u_\omega^L + e^{-\pi\omega/2a} u_\omega^{R*}}{\sqrt{2 \sinh(\pi\omega/a)}} \quad (112)$$

are analytic in the lower half-plane — hence purely positive frequency in the inertial sense — and correctly normalized (the denominators follow from $e^{\pi\omega/a} - e^{-\pi\omega/a} = 2 \sinh(\pi\omega/a)$ and the opposite norms of a mode and its conjugate). Reading (112) against the template (105): the Bogoliubov

coefficients connecting the Rindler bookkeeper to the inertial one are diagonal in ω with

$$\alpha_\omega = \frac{e^{\pi\omega/2a}}{\sqrt{2\sinh(\pi\omega/a)}}, \quad \beta_\omega = \frac{e^{-\pi\omega/2a}}{\sqrt{2\sinh(\pi\omega/a)}}, \quad \left| \frac{\beta_\omega}{\alpha_\omega} \right| = e^{-\pi\omega/a}, \quad (113)$$

and $|\alpha_\omega|^2 - |\beta_\omega|^2 = 1$ checks the consistency condition (107). The occupation seen by the accelerated bookkeeper in the inertial vacuum is then, by (108),

$$\langle N_\omega \rangle = |\beta_\omega|^2 = \frac{e^{-\pi\omega/a}}{2\sinh(\pi\omega/a)} = \frac{1}{e^{2\pi\omega/a} - 1} = \frac{1}{e^{\hbar\omega/k_B T_U} - 1}, \quad k_B T_U = \frac{\hbar a}{2\pi c}, \quad (114)$$

an exact Planck distribution at the Unruh temperature (Davies and Fulling, 1976; Unruh, 1976). No approximation was made: the thermality is an identity of the analytic continuation.

Physical picture. The accelerated observer's horizon at $U = 0$ cuts the Minkowski vacuum in half, and the vacuum's correlations do not respect the cut: each right-wedge mode is entangled with a left-wedge partner (the pairing exhibited by (112)). An observer confined to one side sees her half of an entangled pure state — which is a mixed state, and the analytic structure forces the mixture to be exactly thermal. Acceleration does not *heat* the vacuum; it *hides half of it*, and the hidden half's ledger entry is a temperature. This is the OGGEp partial trace in its simplest costume.

H.4 Step 4 — from Rindler to the black hole: two routes to κ

In GG-Theory the object carrying the horizon is the 6D bulk; near the horizon it looks like

$$ds_6^2 \approx -f(r_4)dt^2 + dr_4^2/f(r_4) + r_4^2 d\Omega^2 + dr^2 + d\theta^2, \quad (115)$$

for a Schwarzschild 4D metric with $f(r_4) = 1 - r_s/r_4$: the fiber directions (r, θ) ride along as spectators, exactly like the transverse directions of Step 3, so the 1+1 analysis applies verbatim to each 4D mode. Two independent routes now deliver the surface gravity as the acceleration parameter.

Route (a): the near-horizon Rindler frame. In geometrized units ($G = c = 1$), expand f about the horizon radius r_s :

$$f(r_4) = f'(r_s)(r_4 - r_s) + O((r_4 - r_s)^2) = 2\kappa(r_4 - r_s), \quad \kappa \equiv \frac{1}{2}f'(r_s) = \frac{1}{2r_s}, \quad (116)$$

using $f'(r_4) = r_s/r_4^2$. Trade r_4 for the proper distance from the horizon,

$$x = \int_{r_s}^{r_4} \frac{dr'}{\sqrt{f(r')}} = \int_{r_s}^{r_4} \frac{dr'}{\sqrt{2\kappa(r' - r_s)}} = \sqrt{\frac{2(r_4 - r_s)}{\kappa}} \implies f = \kappa^2 x^2, \quad (117)$$

so the (t, x) block of the metric is $ds^2 = -\kappa^2 x^2 dt^2 + dx^2$: precisely Rindler, with the static observer at proper distance x accelerating at $a(x) = 1/x$. By Step 3 she measures the local temperature $T(x) = \hbar/(2\pi x)$ — divergent at the horizon, as an infinite acceleration demands — and the Tolman

redshift carries it out: the asymptotic observer sees

$$T_\infty = \sqrt{-g_{tt}} T(x) = (\kappa x) \cdot \frac{\hbar}{2\pi x} = \frac{\hbar\kappa}{2\pi}, \quad (118)$$

with the x -dependence canceling identically. The horizon temperature is the redshift-invariant content of the diverging local Unruh effect.

Route (b): the collapse ray-tracing. Transforming to Kruskal-like coordinates $U = -e^{-\kappa u}$, $V = e^{\kappa v}$ with $\kappa = 1/(2r_s)$:

$$UV = -e^{r_4/r_s}(r_4/r_s - 1), \quad ds^2|_{t,r_4} = -\frac{4r_s^3}{r_4} e^{-r_4/r_s} dU dV, \quad (119)$$

the second relation the regular Kruskal form of the (t, r_4) block, its prefactor finite and nonzero at $r_4 = r_s$. The near-horizon metric in UV coordinates is *regular* at the horizon (the Finkelstein observation), so the state regular there — the one a collapse actually prepares at late times — is the one positive-frequency in U . The exterior bookkeeper uses $u = -\frac{1}{\kappa} \ln(-U)$, so her outgoing modes are $e^{-i\omega u} = (-U)^{i\omega/\kappa}$ — literally (110) with $a \rightarrow \kappa$ (the constant rescale of U between the two conventions shifts a phase, not the continuation factor). The continuation argument (111) then transfers word for word and yields the ratio of the main text, $|\beta_\omega/\alpha_\omega| = e^{-\pi\omega/\kappa}$ (Eq. (16)), hence the Planck occupation at

$$k_B T_H = \frac{\hbar\kappa}{2\pi}. \quad (120)$$

This is Hawking’s derivation (Hawking, 1975) in its analytic kernel: the collapse’s role is only to prepare regularity at the horizon, and the exponential map $u = -\frac{1}{\kappa} \ln(-U)$ between the two bookkeepers’ clocks does the rest. What the exterior curvature adds is the graybody filter $\Gamma_\ell(\omega)$ of Eq. (19); what rotation adds is the anchor selection of Theorem 3; what collapse transients add is quantified by Theorem 4.

Physical picture. A horizon is an exponential magnifier. Outgoing rays that escape at late exterior time u hugged the horizon at Kruskal parameter $U \propto -e^{-\kappa u}$: equal steps in the observer’s clock correspond to exponentially shrinking steps of the regular clock. A logarithmic lens applied to a smooth state produces a scale-invariant exponential ratio between absorption and emission — one number, κ — and a spectrum that can only be Planck. Thermality is not a property of what fell in; it is the signature of the lens.

H.5 Step 5 — restoring units, and the numbers

Restoring $\hbar$, c , k_B , and G_N : the geometric surface gravity $\kappa = 1/(2r_s)$ (an inverse length, $r_s = 2G_N M/c^2$) gives

$$T_H = \frac{\hbar\kappa}{2\pi k_B c}. \quad (121)$$

in the acceleration convention $\kappa_{\text{acc}} = c^2 \kappa_{\text{geom}}$ used throughout this paper (equivalently $T_{\text{H}} = \hbar c \kappa_{\text{geom}} / (2\pi k_B)$). Explicitly,

$$T_{\text{H}} = \frac{\hbar c^3}{8\pi G_N M k_B} \approx 6.17 \times 10^{-8} \text{ K} \left(\frac{M_{\odot}}{M} \right), \quad (122)$$

the reference-card value (Appendix D; verified in `p2c2_verify`). The inverse scaling with mass is the signature fact of black-hole thermodynamics: M87* at $6.5 \times 10^9 M_{\odot}$ sits at $\sim 10^{-17}$ K, colder than anything else in physics, while a 10^{12} kg primordial hole would radiate at $\sim 10^{11}$ K. Big holes are cold, small holes are furnaces, and the negative specific heat this implies is derived honestly in Appendix I, Step 3.

H.6 Step 6 — what GDOS adds: the custody reading

The OGGEF projection to the exterior region preserves the thermal structure Steps 3–4 derived: the output $\rho_{\partial}^{\text{out}}$ is thermal at T_{H} , because the trace over the hidden partners is exactly the operation that turned the squeezed pair state into a Planck mixture (Eq. (18)). The derivation used only the bulk geometry and the OGGEF partial trace; it did not postulate Hawking’s specific mechanism. Nor is the analytic continuation of Step 3 — the single non-algebraic move in the whole chapter — a prescription: its existence is the Wick-rotation clause of the cohort’s Wick-quartet theorem (T-P5-B1, Arisaka, 2026, P5-CPT), consumed at the horizon as Theorem 5; the geometry contributes only the period $2\pi/\kappa$. The custody division is exact: the dictionary owns which modes count as particles and which partners are hidden (port data); the geometry owns the clock whose one-sided spectrum licenses the continuation and whose surface gravity sets its cost.

Problems and worked solutions. H.2.1. Prove the single-mode consistency condition

$|\alpha|^2 - |\beta|^2 = 1$ directly from the commutator.

Solution. Step 1: for one mode, $b = \alpha^* a - \beta^* a^\dagger$ (the single-mode case of (105)). Step 2: impose $[b, b^\dagger] = 1$: $[b, b^\dagger] = |\alpha|^2 [a, a^\dagger] + |\beta|^2 [a^\dagger, a] = |\alpha|^2 - |\beta|^2$. Step 3: hence $|\alpha|^2 - |\beta|^2 = 1$ — the one-mode face of (107), and the algebraic reason the Planck occupation (108) can be written $1/(e^{2\pi\omega/\kappa} - 1)$: the ratio $|\beta/\alpha| = e^{-\pi\omega/\kappa}$ plus this normalization fixes both moduli. Physics intuition: the condition is a hyperbolic identity, $\cosh^2 - \sinh^2 = 1$ — a Bogoliubov transformation is a boost in the $(a, a^\dagger)$ plane, and thermality is what a boost does to a vacuum.

H.3.1. Evaluate the Unruh temperature for one Earth gravity, and the acceleration required to feel room temperature.

Solution. Step 1: $k_B T_{\text{U}} = \hbar a / (2\pi c)$, so $T_{\text{U}}(g) = \hbar g / (2\pi c k_B) \approx 4.0 \times 10^{-20}$ K. Step 2: inverting for $T = 300$ K gives $a = 2\pi c k_B T / \hbar \approx 7.4 \times 10^{22} \text{ m s}^{-2}$ — about 10^{21} Earth gravities, the surface gravity of nothing tamer than a horizon. Step 3: this is why the Unruh effect has never been seen directly, and why the black hole is its one cosmic amplifier: at a horizon the *geometry* supplies the acceleration budget no laboratory can. Physics intuition: $\hbar / (2\pi c k_B)$ is nature’s exchange rate between acceleration and temperature, and it is minuscule — thermality from motion is real but priced beyond any engine except gravity’s.

H.4.1. Show that the constant rescaling of the Kruskal coordinate between the two routes shifts no physics.

Solution. Step 1: Route (b) uses $U = -e^{-\kappa u}$; any alternative affine choice is $U' = \lambda U$ with $\lambda > 0$ constant. Step 2: the continuation argument of (111) tracks $(-U)^{i\omega/\kappa} \rightarrow e^{\pi\omega/\kappa}(U)^{i\omega/\kappa}$; under $U \rightarrow \lambda U$ both sides acquire the same factor $\lambda^{i\omega/\kappa}$ — a constant *phase* on the mode, which cancels in $|\beta_\omega/\alpha_\omega|$. Step 3: hence the ratio $e^{-\pi\omega/\kappa}$, and with it T_H , is invariant — only the branch-crossing exponential is physical. Physics intuition: the temperature is not stored in any coordinate convention but in the *one* analytic fact — crossing the horizon zero of U costs $e^{\pi\omega/\kappa}$ — which is why three superficially different derivations (Appendices H and I, and §28) keep landing on the same number.

H.5.1. Find the mass at which a black hole is in equilibrium with today’s cosmic microwave background, and classify the population.

Solution. Step 1: set $T_H = \hbar c^3/(8\pi G_N M k_B) = 2.725$ K. Step 2: solve: $M = 4.5 \times 10^{22}$ kg ≈ 0.61 lunar masses ($2.3 \times 10^{-8} M_\odot$). Step 3: every astrophysical hole ($M \gtrsim M_\odot$) is therefore *colder* than the sky it sits in and absorbs net energy today; only primordial holes below lunar mass — none yet observed — would be net emitters. The crossover drifts upward as the universe cools. Physics intuition: evaporation is real but *scheduled*: the universe must first cool below each hole’s temperature before its long shrinkage can begin — thermodynamics queues the black holes behind the expansion.

H.6.1. Show that the Planck occupation and the branch-cut ratio are the same statement (detailed balance).

Solution. Step 1: from $|\beta_\omega/\alpha_\omega| = e^{-\pi\omega/\kappa}$ and H.2.1, $|\beta|^2 = 1/(e^{2\pi\omega/\kappa} - 1)$: the ratio implies Planck. Step 2: conversely, a KMS (thermal) state at temperature T satisfies detailed balance — emission and absorption amplitudes related by $e^{-\hbar\omega/k_B T}$; reading $|\beta/\alpha|^2 = e^{-2\pi\omega/\kappa}$ as that Boltzmann factor identifies $k_B T = \hbar\kappa/2\pi c$: Planck implies the ratio. Step 3: the equivalence is the content of T-BH-4/T-BH-5 — the branch-cut asymmetry *is* the thermal face, one datum wearing two names. Physics intuition: nothing “heats up” at a horizon; a one-sided exponential map re-weights outgoing modes exactly as a Boltzmann factor would, and a detector cannot tell the difference — which is the paper’s whole thermality thesis in one mode.

Take-away. *Standard result:* $k_B T_H = \hbar\kappa/(2\pi c)$, from Bogoliubov coefficients on the collapse background. *What it assumed:* free fields on a fixed classical geometry, and a positive-frequency continuation prescription. *The GG-Theory reading:* the prescription is licensed by the one-sided spectrum of the port clock (T-P5-B1, Theorem 5); the geometry contributes only the period $2\pi/\kappa$; and the thermality is the KMS face of an honest CPTP channel, not evidence of information loss.

I Euclidean derivation of temperature and entropy

What this appendix teaches. The Euclidean route to the same temperature and to the entropy: the cigar geometry and its regularity period, the on-shell action computed term by term including the Gibbons–Hawking boundary term, the saddle-point thermodynamics — and the GG-Theory reading in which both formal moves of 1977, the continuation and the arena, carry licenses.

Read alongside: Gibbons and Hawking (1977) and Gibbons and Hawking (1977, action) for the originals; Wald (1994) for the careful modern statement.

The Euclidean route computes the same temperature from global regularity, and adds the entropy from the saddle-point free energy. Historically both moves were formal prescriptions; in this paper each is licensed — the continuation by T-P5-B1 (consumed as Theorem 5), the arena by the analytic-shore clause of §31.5. This chapter works the computation a graduate student should reproduce line by line: Step 1 the cigar and the period; Step 2 the on-shell action, term by term; Step 3 the thermodynamics of the saddle; Step 4 the GG-Theory reading. Geometrized units $G = c = 1$ (with G restored where it earns its keep); $\hbar = k_B = 1$ inside the saddle algebra.

I.1 Step 1 — the cigar and the period

From the Euclidean Schwarzschild metric ($t \rightarrow -i\tau$)

$$ds_E^2 = \left(1 - \frac{2GM}{r}\right) d\tau^2 + \left(1 - \frac{2GM}{r}\right)^{-1} dr^2 + r^2 d\Omega_2^2, \quad (123)$$

zoom in on the horizon $r_s = 2GM$. Write $r = r_s + \rho$ with $\rho \ll r_s$, so $f = 1 - r_s/r = \rho/(r_s + \rho) \simeq \rho/r_s$. Substitute $\rho = x^2/(4r_s)$, i.e. $x = 2\sqrt{r_s\rho}$, and work out both metric blocks:

$$\frac{dr^2}{f} = \frac{d\rho^2}{\rho/r_s} = \frac{r_s}{\rho} \left(\frac{x dx}{2r_s}\right)^2 = \frac{4r_s^2}{x^2} \cdot \frac{x^2 dx^2}{4r_s^2} = dx^2, \quad f d\tau^2 = \frac{x^2}{4r_s^2} d\tau^2, \quad (124)$$

so that near the horizon

$$ds_E^2 \simeq dx^2 + x^2 d\left(\frac{\tau}{2r_s}\right)^2 + r_s^2 d\Omega_2^2 : \quad (125)$$

a plane in polar coordinates $(x, \tau/2r_s)$ times the horizon sphere. The plane is smooth at $x = 0$ if and only if the angle $\tau/2r_s$ has period 2π — otherwise the tip carries a conical singularity, which the vacuum saddle is not allowed to have. Hence

$$\tau \sim \tau + \beta, \quad \beta = 4\pi r_s = 8\pi GM \implies T_H = \frac{1}{8\pi GM} = \frac{\kappa}{2\pi}, \quad (126)$$

in exact agreement with Appendix H, Step 4 (with $\kappa = 1/2r_s = 1/4GM$). The same tip, *deliberately* given a deficit angle, is the replica geometry of the main text's entropy derivation (§31.5) — regularity is the $n = 1$ member of the replica family.

Physical picture. Periodic imaginary time is the field-theory encoding of a temperature ($\text{Tr } e^{-\beta H}$), and the geometry votes on the period: only one β makes the Euclidean section a smooth cigar rather than a cone. The black hole does not sit *in* a heat bath at T_H ; its own geometry *is* the statement that T_H is the only temperature it can be in equilibrium at.

I.2 Step 2 — the on-shell action, term by term

The saddle's free energy needs the Euclidean action

$$I_E = -\frac{1}{16\pi G} \int_{\mathcal{M}} d^4x \sqrt{g} R - \frac{1}{8\pi G} \oint_{\partial\mathcal{M}} d^3x \sqrt{h} (K - K_0), \quad (127)$$

the second term the Gibbons–Hawking–York boundary term with the flat-space subtraction K_0 (Gibbons and Hawking, 1977, action). For Schwarzschild $R = 0$: the entire on-shell action is boundary. Evaluate on the surface $r = R$ (to be sent to infinity), with induced metric $h = \text{diag}(f, r^2, r^2 \sin^2 \theta)|_{r=R}$ and outward unit normal $n = \sqrt{f} \partial_r$.

The trace of the extrinsic curvature. For a metric of this form, $K = \nabla_\mu n^\mu = \frac{1}{\sqrt{g}} \partial_r (\sqrt{g} \sqrt{f})$ with $\sqrt{g} = r^2 \sin \theta$, so

$$K = \frac{1}{r^2} \partial_r (r^2 \sqrt{f}) = \frac{f'}{2\sqrt{f}} + \frac{2\sqrt{f}}{r}. \quad (128)$$

The integral. The boundary volume is $\sqrt{h} = \sqrt{f} R^2 \sin \theta$ with τ -length β , so

$$\oint \sqrt{h} K = \beta 4\pi R^2 \sqrt{f} \left(\frac{f'}{2\sqrt{f}} + \frac{2\sqrt{f}}{R} \right) = 4\pi\beta \left(\frac{R^2 f'}{2} + 2Rf \right) = 4\pi\beta (2R - 3GM), \quad (129)$$

using $f = 1 - 2GM/R$, $f' = 2GM/R^2$, so $R^2 f'/2 = GM$ and $2Rf = 2R - 4GM$. *The subtraction.* Flat space at the same boundary has $K_0 = 2/R$, but its time circle must be matched to the same proper length, $\beta_0 = \beta \sqrt{f(R)}$:

$$\oint \sqrt{h} K_0 = \beta \sqrt{f} \cdot 4\pi R^2 \cdot \frac{2}{R} = 8\pi\beta R \sqrt{f} = 4\pi\beta (2R - 2GM) + O(1/R), \quad (130)$$

expanding $\sqrt{f} = 1 - GM/R + O(R^{-2})$. Subtracting and taking $R \rightarrow \infty$:

$$I_E = -\frac{4\pi\beta}{8\pi G} \left[(2R - 3GM) - (2R - 2GM) \right] = \frac{\beta M}{2} = \frac{\beta^2}{16\pi G}, \quad (131)$$

the last form using $M = \beta/8\pi G$ from (126). The divergent $2R$ pieces canceled exactly — that is what the subtraction is for — and the finite remainder is the whole answer.

I.3 Step 3 — the thermodynamics of the saddle

With $Z(\beta) \simeq e^{-I_E}$ at the saddle, every thermodynamic quantity follows by differentiation, and each line is a consistency check:

$$\begin{aligned} E &= \partial_\beta I_E = \frac{\beta}{8\pi G} = M && \text{(the saddle's energy is the mass: check),} \\ F &= \frac{I_E}{\beta} = \frac{M}{2}, && S = \beta E - I_E = \frac{\beta^2}{16\pi G} = 4\pi G M^2. \end{aligned} \quad (132)$$

Since $A = 4\pi r_s^2 = 16\pi G^2 M^2$, the entropy line reads

$$S_{\text{BH}} = 4\pi G M^2 = \frac{16\pi G^2 M^2}{4G} = \frac{A}{4G_N}, \quad (133)$$

and the first law closes: $dS/dM = 8\pi G M = \beta = 1/T$. One more derivative is physics: $\partial T/\partial M = -1/(8\pi G M^2) < 0$, so the specific heat is

$$C = \frac{\partial M}{\partial T} = -8\pi G M^2 < 0: \quad (134)$$

the asymptotically flat black hole heats up as it loses mass. A negative-specific-heat system cannot equilibrate with an infinite bath — the runaway that makes evaporation accelerate (Appendix K, Step 3) and that the AdS box of Appendix J exists to tame.

Physical picture. Self-gravitating systems generically carry negative specific heat (a satellite dropped to a lower orbit speeds up), and the saddle knows it: the free energy is dominated by a configuration whose temperature falls as its energy grows. The Euclidean computation is thermodynamics done *by geometry* — the partition function evaluated not over microstates but over the one smooth classical shape the boundary conditions admit.

I.4 Step 4 — the GG-Theory reading

In the standard reading, (132) is a thermodynamic consequence of the saddle. In the GG-Theory reading the same computation supplies the geometric piece F_{GEP} of the operational ledger (43): the Euclidean trick does not compete with GG-Theory, it feeds it. Three points, stated plainly. (i) The periodicity argument fixes the temperature and the entropy *functional* but does not say where the microstates live — that is the content of the induced-gravity Part. (ii) The derivation is inherently coarse-grained: it computes a saddle free energy seen by an observer who has already projected to a regular Euclidean boundary problem — exactly the observer-facing object GDOS tracks. (iii) The “trick” itself is a theorem — the continuation to the Euclidean section (the analytic shore among time’s four endowments; Arisaka, 2026, P5-CPT) exists because the exterior clock’s spectrum is one-sided (T-P5-B1, consumed as Theorem 5), the saddle’s periodicity being the geometry’s regularity datum.

Problems and worked solutions. I.1.1. Repeat the cigar argument for the de Sitter horizon and extract its temperature.

Solution. Step 1: the static patch (Appendix Q) has $f(r) = 1 - H^2 r^2$, with horizon at $r_H = 1/H$ and $f'(r_H) = -2H$. Step 2: near the horizon set $\rho = r_H - r$: $f \approx 2H\rho$, and the proper distance is $x = \int dr/\sqrt{f} = \sqrt{2\rho/H}$, so $f = H^2 x^2$ and the Euclidean (t, x) block is a plane in polar coordinates with angle $H\tau$. Step 3: smoothness demands period $\beta = 2\pi/H$, i.e. $k_B T_{\text{dS}} = \hbar H/2\pi$ — the first entry of (174), by exactly the argument of (126). Physics intuition: the cigar does not care which side of the horizon the observer is on — enclosure is enclosure — which is why one theorem prices both the black-hole sky and the cosmological one.

I.2.1. Verify that $I_E = \beta M/2$ yields exactly $S = Ac^3 k_B/(4G_N \hbar)$.

Solution. Step 1 (geometrized units): from (131), $I_E = \beta M/2$ with $\beta = 8\pi G_N M$. Step 2: the saddle-point free energy is $F = I_E/\beta = M/2$, and $S = \beta(M - F) = \beta M/2 = 4\pi G_N M^2$. Step 3: express through the area $A = 16\pi G_N^2 M^2$: $S = A/4G_N$ — and restoring $\hbar, c, k_B$ gives the Bekenstein–Hawking value with coefficient exactly $\frac{1}{4}$. The factor-of-two anatomy: half of βM is subtracted by the boundary term ($F = M/2$), and the *other* half is the entropy. Physics intuition: in gravity the on-shell bulk action vanishes ($R = 0$); the entire thermodynamics lives on the boundary term — an early, exact hint that the physics of horizons is boundary bookkeeping, which the main text’s port reading takes literally.

I.3.1. Compute the specific heat of the Schwarzschild saddle and interpret its sign.

Solution. Step 1: from (137)-adjacent scaling $T = 1/(8\pi G_N M)$ (geometrized), $M = 1/(8\pi G_N T)$. Step 2: $C = dM/dT = -1/(8\pi G_N T^2) < 0$ — restoring units, $C = -\hbar c^5/(8\pi G_N k_B T^2) \cdot k_B^2$, negative always. Step 3: negative C means no stable canonical equilibrium in flat space: a hole warmer than its bath radiates, loses mass, gets *warmer* — runaway; a hole colder than its bath grows without bound. The AdS box of Appendix J is the standard cure, and the branch structure there ($C < 0$ small branch, $C > 0$ large branch, (138)) is this computation stabilized by curvature. Physics intuition: self-gravitating systems heat up by losing energy (virial logic); the black hole is the extreme case, and its thermodynamic instability is why evaporation, once begun, accelerates to the end.

I.4.1. Derive the conical deficit for a period $\beta' \neq \beta$, and connect it to the replica method.

Solution. Step 1: with the near-horizon plane $ds^2 = dx^2 + \kappa^2 x^2 d\tau^2$ and $\tau \sim \tau + \beta'$, the angle variable $\phi = \kappa\tau$ runs over $[0, \kappa\beta']$; smoothness needs $\kappa\beta' = 2\pi$. Step 2: for $\beta' = n\beta$ (the n -fold replica) the tip carries deficit $\delta = 2\pi(1 - n)$, and the curvature acquires a distributional piece $R = 2\delta\delta^{(2)}(x) + R_{\text{smooth}}$. Step 3: inserting this delta into the Einstein–Hilbert action gives the replica free energy its area term — ∂_n at $n = 1$ yields $A/4G_N$ — which is precisely the engine of the main text’s exact-1/4 derivation (§31.5) and of the holographic entanglement entropy (Appendix L, Step 3). Physics intuition: entropy is the response of the action to a deficit angle at the horizon — geometry’s way of counting what the tip hides; every modern route to $A/4G_N \hbar$, including this paper’s, passes through this one computation.

Take-away. *Standard result:* T_H from cigar regularity and $S_{\text{BH}} = A/4G_N \hbar$ from the on-shell action. *What it assumed (in 1977):* that the Euclidean continuation and the saddle-point ensemble mean something — two formal moves. *The GG-Theory reading:* both moves carry licenses — the continuation from the one-sided clock spectrum, the arena from the analytic-shore clause — so the “trick” is a theorem consumed, not a prescription trusted.

J Hawking–Page free-energy notebook

What this appendix teaches. The Hawking–Page transition worked as a free-energy notebook: the two black-hole branches in the AdS box, the free-energy ledger that selects between the gas and the large hole, and the stability analysis — the cleanest classical template for a branch switch selected by a ledger, which is how the main text reads the Page turnover.

Read alongside: [Hawking and Page \(1983\)](#) for the original; [Harlow \(2016\)](#) for the transition’s modern role in holography.

Anti-de Sitter space is the box that tames the negative specific heat of Appendix I: its curvature acts as a reflecting wall, so a black hole can genuinely equilibrate with its own radiation, and the canonical ensemble becomes honest. The price of honesty is a phase transition ([Hawking and Page, 1983](#)), and this notebook works it in three Steps — the cleanest template in gravity for a branch switch selected by a free-energy ledger, which is exactly how this paper reads the Page turnover (Theorem 2). Units $G \equiv G_N$, $c = \hbar = k_B = 1$.

J.1 Step 1 — temperature from regularity, and the two branches

For the 4D AdS–Schwarzschild metric function

$$f(r) = 1 - \frac{2GM}{r} + \frac{r^2}{L^2}, \quad (135)$$

the horizon r_+ is the largest root of f , so $2GM = r_+ + r_+^3/L^2$. The cigar argument of Appendix I, Step 1 fixes the period as $\beta = 4\pi/f'(r_+)$; differentiate and eliminate M :

$$f'(r) = \frac{2GM}{r^2} + \frac{2r}{L^2} \implies f'(r_+) = \frac{r_+ + r_+^3/L^2}{r_+^2} + \frac{2r_+}{L^2} = \frac{1}{r_+} + \frac{3r_+}{L^2}, \quad (136)$$

$$T(r_+) = \frac{1}{4\pi r_+} \left(1 + \frac{3r_+^2}{L^2} \right). \quad (137)$$

The two terms compete: $T \sim 1/4\pi r_+$ for small holes (the flat-space behavior) and $T \sim 3r_+/4\pi L^2$ for large ones (the box behavior). Minimizing, $dT/dr_+ = 0$ gives $-1/r_+^2 + 3/L^2 = 0$:

$$r_+^{\min} = \frac{L}{\sqrt{3}}, \quad T_{\min} = \frac{\sqrt{3}}{2\pi L}. \quad (138)$$

Below $T_{\min}$ no black hole exists at all; above it there are *two*, a small branch ($r_+ < L/\sqrt{3}$) and a large branch ($r_+ > L/\sqrt{3}$), and the ensemble must choose.

J.2 Step 2 — the free-energy ledger

The entropy, mass, and free energy on a branch are

$$S = \frac{\pi r_+^2}{G_N}, \quad M = \frac{r_+}{2G_N} \left(1 + \frac{r_+^2}{L^2} \right), \quad F = M - TS = \frac{r_+}{4G_N} \left(1 - \frac{r_+^2}{L^2} \right), \quad (139)$$

where the F line is worth one displayed check:

$$F = \frac{r_+}{2G} \left(1 + \frac{r_+^2}{L^2}\right) - \frac{1}{4\pi r_+} \left(1 + \frac{3r_+^2}{L^2}\right) \frac{\pi r_+^2}{G} = \frac{r_+}{4G} \left[2 + \frac{2r_+^2}{L^2} - 1 - \frac{3r_+^2}{L^2}\right] = \frac{r_+}{4G} \left(1 - \frac{r_+^2}{L^2}\right). \quad (140)$$

Equivalently, the on-shell Euclidean action (same subtraction discipline as Appendix I, Step 2, with thermal AdS as the reference background) is

$$I_E = \beta F = \frac{\pi r_+^2 (L^2 - r_+^2)}{G (L^2 + 3r_+^2)}. \quad (141)$$

F changes sign at $r_+ = L$, i.e. at the Hawking–Page temperature

$$T_{\text{HP}} = T(L) = \frac{1}{\pi L} : \quad (142)$$

below it thermal AdS dominates the ensemble (the black-hole saddle has $F > 0$), above it the large black hole dominates ($F < 0$). In the dual gauge theory this is the confinement/deconfinement transition (Maldacena, 1998).

J.3 Step 3 — stability, and the branch-switch template

Which branch is even locally stable? The specific heat along the branch is $C = dM/dT = (dM/dr_+)/ (dT/dr_+)$, and $dM/dr_+ = (1 + 3r_+^2/L^2)/2G > 0$ always, so the sign of C is the sign of dT/dr_+ :

$$C < 0 \quad (r_+ < L/\sqrt{3}, \text{ small branch}), \quad C > 0 \quad (r_+ > L/\sqrt{3}, \text{ large branch}). \quad (143)$$

The small branch is the flat-space runaway of Appendix I, Step 3, embedded in AdS — a local maximum of the free energy, the instanton mediating the transition — while the large branch is a genuine equilibrium.

Physical picture, and the GG-Theory close. Three saddles sit in one ensemble — thermal gas, small hole, large hole — and the physical state is whichever minimizes the ledger; as T rises through T_{HP} the minimum jumps branches discontinuously while everything on each branch stays smooth. That is the entire grammar of a first-order transition, realized in pure geometry. In GG-Theory language this is the cleanest template for a branch switch selected by the generalized free-energy ledger: phase grammar, stability, and information bookkeeping fused before any GG-Theory ingredient is added — and Theorem 2 reads the Page turnover as precisely such a switch, with (43) the ledger and the island branch the analogue of the large hole.

Problems and worked solutions. J.1.1. Fill in the minimization that yields the minimum temperature (138).

Solution. Step 1: from (137), $4\pi T = 1/r_+ + 3r_+/L^2$. Step 2: $dT/dr_+ = 0$ gives $-1/r_+^2 + 3/L^2 = 0$, i.e. $r_+^{\min} = L/\sqrt{3}$; the second derivative $2/r_+^3 > 0$ confirms a minimum. Step 3: substituting back,

$4\pi T_{\min} = \sqrt{3}/L + \sqrt{3}/L = 2\sqrt{3}/L$, so $T_{\min} = \sqrt{3}/(2\pi L)$ (138). Below $T_{\min}$ the ensemble contains no black-hole saddle at all. Physics intuition: the two terms are the two regimes — the flat-space $1/r_+$ (small holes, negative specific heat) and the box's r_+/L^2 (large holes, positive) — and $T_{\min}$ is where the box first wins; everything the transition does is a negotiation between these two terms.

J.2.1. Show the transition is first order, with entropy jump of order the horizon area.

Solution. Step 1: at $T_{\text{HP}} = 1/\pi L$ (142), the dominant saddle switches between thermal AdS ($F \approx 0$, entropy of a gas, $O(1)$ in G_N) and the large hole ($r_+ = L$). Step 2: the large hole's entropy is $S = \pi r_+^2/G_N = \pi L^2/G_N$ — in Planck units, a number of order $(L/\ell_P)^2$. Step 3: since F is continuous at the crossing but $S = -\partial F/\partial T$ jumps by $\sim L^2/G_N$, the transition is first order with a macroscopic latent heat $T_{\text{HP}} \Delta S$. Physics intuition: the ensemble reorganizes from $O(1)$ to $O(1/G_N)$ degrees of freedom in one step — in the dual gauge theory this is deconfinement, and in the main text's language it is the cleanest classical instance of a ledger-selected branch switch, the grammar of Theorem 2.

J.3.1. Confirm the stability assignment of the two branches from the specific heat.

Solution. Step 1: $C = dM/dT = (dM/dr_+)/(dT/dr_+)$, and $dM/dr_+ = (1 + 3r_+^2/L^2)/2G_N > 0$ always. Step 2: from J.1.1, $dT/dr_+ < 0$ for $r_+ < L/\sqrt{3}$ and > 0 for $r_+ > L/\sqrt{3}$: hence $C < 0$ (small branch), $C > 0$ (large branch). Step 3: the small branch is the flat-space runaway of I.3.1 embedded in AdS — a local *maximum* of free energy acting as the nucleation barrier — while the large branch is a genuine equilibrium. Physics intuition: the box turns the runaway into a phase diagram: the unstable saddle does not disappear, it becomes the mountain pass between gas and hole, exactly the role the sub-dominant branch plays in the Page-turnover ledger race (Appendix K, Step 4).

Take-away. *Standard result:* the first-order Hawking–Page transition between thermal AdS and the large black hole, selected by the free-energy ledger. *What it assumed:* the AdS box and the Euclidean ensemble. *The GG-Theory reading:* the cleanest pre-existing template for a ledger-selected branch switch — the grammar in which the main text reads the Page turnover (Theorem 2).

K The Page curve entropy calculation

What this appendix teaches. The Page curve as a counting theorem: the exact random-state average, the curve and its corner, the evaporation clock with the ratio $t_{\text{Page}} \approx 0.646 t_{\text{evap}}$ derived rather than quoted, and the statement of exactly which premise (typicality of the port ensemble) GDOS keeps and flags.

Read alongside: Page (1993) for the original; Harlow (2016) for the modern quantum-information formulation; Hayden and Preskill (2007) for the mirror refinement.

The Page curve is a counting statement before it is a gravity statement, and this chapter works the counting in full: Step 1 the exact random-state average (Page, 1993); Step 2 the curve and its corner; Step 3 the evaporation clock, including the reference-card ratio $t_{\text{Page}} \approx 0.646 t_{\text{evap}}$ derived rather than quoted; Step 4 what GDOS keeps and what it flags.

K.1 Step 1 — the exact random-state average

For a unitary evaporation process of a black hole of initial entropy $S_{\text{BH}}^{(0)} = N$ coarse-grained qubit units, consider the system when m units have been emitted: the radiation factor has dimension $d_R = 2^m$, the black-hole factor $d_B = 2^{N-m}$. The global state is bipartite pure, so the two reduced entropies agree,

$$S_{\text{vN}}(\rho_R) = S_{\text{vN}}(\rho_B). \quad (144)$$

— whatever the dynamics, the two subsystems' ledgers are mirror images. The dynamical input is Page's postulate: after scrambling, the global pure state is well modeled as Haar-random in the $d_R d_B$ -dimensional Hilbert space. For such a state Page conjectured and subsequent work proved the exact average (for $d_R \leq d_B$)

$$\langle S_{\text{vN}}(\rho_R) \rangle = \sum_{k=d_B+1}^{d_R d_B} \frac{1}{k} - \frac{d_R - 1}{2d_B} = H_{d_R d_B} - H_{d_B} - \frac{d_R - 1}{2d_B}, \quad (145)$$

with $H_n = \sum_{k \leq n} 1/k$ the harmonic numbers. The asymptotics follow from $H_n = \ln n + \gamma_E + 1/2n + O(1/n^2)$:

$$\langle S_{\text{vN}}(\rho_R) \rangle = \ln d_R - \frac{d_R}{2d_B} + O\left(\frac{1}{d_R d_B}\right), \quad (146)$$

i.e. the radiation is *maximally* entangled with the hole up to an exponentially small deficit $d_R/2d_B = 2^{2m-N}/2$. For $d_R \geq d_B$ exchange the labels (the spectrum of the two reduced states is identical by (144)).

Physical picture. A random pure state hides its information democratically: any subsystem smaller than half the total looks *exactly* thermal to within $2^{-|\text{rest}|}$ corrections. Early radiation therefore carries essentially zero information — not because information is lost but because it is stored in correlations *with* the remainder, unreadable locally. The deficit term is the whole story of the information return, waiting in the second-order bookkeeping.

K.2 Step 2 — the curve and its corner

Reading (146) across the evaporation, with m emitted of N total:

- For $m < N/2$: $S_R \approx m \ln 2$ (linear growth, matching Hawking's calculation exactly — this half of the curve is not modified).
- For $m = N/2$ (the Page point): the exact average sits $\frac{1}{2}$ nat below the maximum, $\langle S_R \rangle = \frac{N}{2} \ln 2 - \frac{1}{2} + O(2^{-N})$ — the corner is rounded only at the third decimal for astrophysical N .
- For $m > N/2$: $S_R \approx (N - m) \ln 2$: the radiation entropy now *tracks the shrinking hole*, decreasing to zero at complete evaporation.

The triangle $S_R(m) \simeq \min(m, N-m) \ln 2$ is the Page curve. Its rising edge is thermodynamics; its falling edge is unitarity; the corner is where the smaller-half rule changes sides.

K.3 Step 3 — the evaporation clock

To place the corner in time rather than in emitted qubits, run the clock. The Hawking flux is Stefan–Boltzmann at $T_H \propto 1/M$ through an effective area $\propto M^2$, so dimensional analysis fixes $dM/dt = -\alpha \hbar c^4/(G_N^2 M^2)$ with α a graybody-weighted constant; the geometric-optics photon-graviton estimate underlying the reference card corresponds to $\alpha = 1/15360\pi$. Integrating from M_0 ,

$$M(t) = M_0 \left(1 - \frac{t}{t_{\text{evap}}}\right)^{1/3}, \quad t_{\text{evap}} = \frac{5120 \pi G_N^2 M_0^3}{\hbar c^4} \approx 2 \times 10^{67} \text{ yr} \left(\frac{M_0}{M_\odot}\right)^3, \quad (147)$$

the card’s value (verified in `p2c2_verify`). The Page point is reached when half the *entropy* has been emitted, and $S_{\text{BH}} \propto M^2$, so $M_{\text{Page}} = M_0/\sqrt{2}$; inverting (147),

$$\frac{t_{\text{Page}}}{t_{\text{evap}}} = 1 - \left(\frac{M_{\text{Page}}}{M_0}\right)^3 = 1 - 2^{-3/2} \approx 0.6464, \quad (148)$$

which is the reference card’s $t_{\text{Page}} \approx 0.646 t_{\text{evap}}$, now derived: it is a pure number, independent of M_0 and of α , fixed entirely by $S \propto M^2$ and $\dot{M} \propto -M^{-2}$. (Species-weighted emission shifts α , not the ratio; frequency-dependent graybody corrections shift the ratio at the percent level.)

Physical picture. Two-thirds of the lifetime passes before the turnover because the hole spends its entropy slowly at first — big, cold, dim — and the halfway point of the *entropy* ledger sits far past the halfway point of the *mass* ledger. After the corner the hole is hotter, brighter, and increasingly transparent: the last emissions are the most information-rich, exactly as the smaller-half rule demands.

K.4 Step 4 — what GDOS keeps, and what it flags

In GDOS, the bulk evolution is unitary (AX-1), and the random-state hypothesis is supplied at long times by the standard eigenstate-thermalization-hypothesis arguments applied to the bulk — an imported dynamical premise, flagged as such rather than proved here; the lemma that consumes it (Lemma 2.1, L-BH-Fresh) carries it as an explicit condition, and the freshness meter $I(q:R)$ gives the premise its operational form: the turnover is the scheduled expiry of the Port-Freshness Clause, not a new mechanism. Given the premise, the Page curve follows, and the branch-switch reading of the corner is Theorem 2 (T-BH-2). The qubit-model numerics behind these statements ($N = 12$, Haar ensembles, growth-ratio and saturation checks) are logged in the benchmark audit (Appendix U).

Problems and worked solutions. K.1.1. Prove that the two reduced entropies of a bipartite pure state are equal.

Solution. Step 1: Schmidt-decompose $|\psi\rangle = \sum_k \sqrt{p_k} |k\rangle_R |k\rangle_B$ — possible for any bipartite pure state (singular-value decomposition of the coefficient matrix). Step 2: the reduced states are $\rho_R = \sum_k p_k |k\rangle\langle k|_R$ and $\rho_B = \sum_k p_k |k\rangle\langle k|_B$: *identical spectra* $\{p_k\}$. Step 3: von Neumann entropy depends only on the spectrum, so $S(\rho_R) = S(\rho_B)$ — the mirror-ledger identity (144) on which the

whole curve rests. Physics intuition: entanglement is a shared property, not a possession — the radiation cannot “know” more about the hole than the hole matches about the radiation, which is why one curve serves both subsystems.

K.1.2. Evaluate the exact Page average for the smallest nontrivial case, $N = 2$ qubits with one emitted.

Solution. Step 1: $d_R = d_B = 2$; the exact average (145) reads $\langle S \rangle = \sum_{k=3}^4 \frac{1}{k} - \frac{d_R-1}{2d_B}$. Step 2: evaluate: $\frac{1}{3} + \frac{1}{4} - \frac{1}{4} = \frac{1}{3}$ nat = 0.48 bit. Step 3: the maximum possible is $\ln 2 = 0.693$ nat = 1 bit, so even at the Page point of the smallest system the typical state sits 0.36 nat below maximal — the deficit term is visible at $N = 2$ and shrinks exponentially with N (146). Physics intuition: “maximally entangled up to a deficit” is not an asymptotic slogan but an exact, checkable number already for two qubits — and the deficit, tiny as it is, is where all the information lives (Step 1’s physical picture).

K.3.1. Derive the Page-time ratio $t_{\text{Page}}/t_{\text{evap}} = 1 - 2^{-3/2}$ from the two scaling laws alone.

Solution. Step 1: $S_{\text{BH}} \propto M^2$, so half the entropy is emitted when $M = M_0/\sqrt{2}$. Step 2: Stefan–Boltzmann at $T \propto 1/M$ through area $\propto M^2$ gives $dM/dt \propto -1/M^2$, hence $M(t)^3 = M_0^3(1 - t/t_{\text{evap}})$ (147). Step 3: solve $M_0^3 2^{-3/2} = M_0^3(1 - t_{\text{Page}}/t_{\text{evap}})$: $t_{\text{Page}}/t_{\text{evap}} = 1 - 2^{-3/2} = 0.6464$ (148) — a pure number, independent of M_0 and of the graybody constant. Physics intuition: two-thirds of the lifetime passes before the turnover because the hole spends its entropy slowly while big and cold; the information era is the loud final third.

K.4.1. At what average rate does a solar-mass black hole return its information?

Solution. Step 1: the archive is $S_{\text{BH}}/k_B \ln 2 \approx 1.5 \times 10^{77}$ bits (reference card, Appendix D). Step 2: it is returned over $t_{\text{evap}} \approx 2 \times 10^{67}$ yr = 6.6×10^{74} s. Step 3: the average rate is ~ 230 bits per second — a slow teletype running for 10^{57} ages of the universe. (The instantaneous rate is far below this before the Page time and above it after; the average is fixed by the two totals alone.) Physics intuition: unitarity is patient — “the information comes back” is a statement about correlations accumulated across the whole evaporation, not a broadcast; no observer who samples less than the late-time majority of the radiation ever decodes a word (the Hayden–Preskill refinement, Hayden and Preskill, 2007).

Take-away. *Standard result:* for unitary evaporation of a typical state, $S_R(m) \approx \min(m, N - m)$ — rise, corner at the Page point, fall — with $t_{\text{Page}} \approx 0.646 t_{\text{evap}}$. *What it assumed:* global unitarity and typicality. *The GG-Theory reading:* unitarity is the bulk’s (AX-1); typicality is a premise made operational by the freshness meter, and its scheduled expiry is the turnover (Theorem 2).

L From state counting to holography: the modern era

What this appendix teaches. The three modern answers to “what does black-hole entropy count?”, each derived to its punchline: the Strominger–Vafa D-brane count that first matched $A/4G_N\hbar$ microscopically; the holographic principle and AdS/CFT, where evaporation is manifestly unitary; and the quantum-extremal-surface / islands program that recovered the Page curve from the gravitational path integral — followed by a precise statement of what each achievement still leaves open, and the map of where the main text picks up each thread.

Read alongside: [Susskind and Lindesay \(2005\)](#) for the arc; [Harlow \(2016\)](#) for the RT/QES technology; [Almheiri et al. \(2021\)](#) for the islands synthesis this appendix compresses.

The Page curve of Appendix K set a target that stood for a quarter century: exhibit, from first principles, the microstates behind (96) and the unitarity behind the curve. This appendix works the three modern campaigns in the order they arrived, at the level of detail a student needs to see *what exactly was proven* — because the precise domain of each proof is where the remaining problems live, and Appendix M then compares the winning formalism with this paper’s reading line by line.

L.1 Step 1 — Strominger–Vafa: the first microscopic count

The 1996 breakthrough ([Strominger and Vafa, 1996](#)) counts the states of a bound system of Q_1 D1-branes and Q_5 D5-branes carrying n units of momentum, in five dimensions, at weak string coupling where the system is an ordinary (if elaborate) quantum mechanics. The count is a Cardy-type asymptotic of the two-dimensional world-volume theory: the central charge is $c_{\text{eff}} = 6 Q_1 Q_5$, and the degeneracy of states at level n obeys

$$S_{\text{micro}} = k_B \ln d(n) \simeq 2\pi k_B \sqrt{\frac{c_{\text{eff}} n}{6}} = 2\pi k_B \sqrt{Q_1 Q_5 n}. \quad (149)$$

Because the states are BPS (supersymmetry-protected), the count is rigid as the coupling is dialed up — the same system gravitationally collapsed is an extremal five-dimensional black hole, and evaluating the horizon area of the corresponding supergravity solution gives, with no adjustment,

$$\frac{k_B A c^3}{4G_5 \hbar} = 2\pi k_B \sqrt{Q_1 Q_5 n} = S_{\text{micro}}. \quad (150)$$

The coefficient $1/4$ — Hawking’s η — emerges from an honest count of named states, exactly. (*Physics intuition:* supersymmetry is the bridge that lets one count the building at blueprint stage, weak coupling, and trust the count after construction, strong coupling; the BPS property is the guarantee that no states move in or out during construction.) What was proven: for a special class — supersymmetric, extremal or near-extremal, in special charge sectors — the Bekenstein–Hawking entropy is state counting. What was not: anything about the Schwarzschild hole in our sky, whose temperature is nonzero, whose charges are astrophysically neutralized, and whose states are protected by nothing.

L.2 Step 2 — the holographic principle and AdS/CFT

Bekenstein’s area-scaling (Appendix G) was promoted by ‘t Hooft (1993) and Susskind (1995) to a principle: the maximum entropy in any region is its *boundary* area in Planck units, so the fundamental description of any gravitating region lives on its boundary. [Maldacena \(1998\)](#) made the principle exact in a specific universe: string theory in anti-de Sitter space is *equal* — not approximately dual — to a conformal field theory on the boundary. For the information paradox the consequence is structural: a black hole in AdS is some high-energy state of the boundary CFT; the CFT is an ordinary quantum system with manifestly unitary evolution; therefore black-hole formation and evaporation, whatever they look like geometrically, cannot lose information. From

1997 on, the question was no longer *whether* unitarity wins but *how the bulk geometry implements it* — and by what mechanism the Hawking calculation’s apparent thermality is corrected. (*Physics intuition*: AdS/CFT settles the verdict without showing the crime scene. Knowing the boundary books balance does not tell you what entry, at the horizon, does the balancing — that gap is exactly where the next twenty years of work, and Steps 3–4, live.)

L.3 Step 3 — geometrizing entanglement: Ryu–Takayanagi and the quantum extremal surface

The instrument that eventually cracked the Page curve was a formula for *entanglement* entropy, not black-hole entropy. [Ryu and Takayanagi \(2006\)](#) proposed, and later derivations confirmed, that the entropy of a boundary region B is the area of the smallest bulk surface γ_B anchored on ∂B :

$$S(B) = \min_{\gamma_B} \frac{A(\gamma_B) k_B c^3}{4G_N \hbar}, \quad (151)$$

a strict generalization of Bekenstein–Hawking (take B = the whole boundary of a black-hole spacetime: γ hugs the horizon). Quantum corrections promote the minimized functional to the *generalized entropy* $S_{\text{gen}} = A/4G_N \hbar + S_{\text{bulk}}$ — area plus the entropy of bulk fields outside γ — and the surface to a *quantum extremal surface* (QES): extremize S_{gen} , then take the minimum over extrema. Two features deserve the student’s attention. First, (151) is derived, in the Euclidean formulation, by the *same conical-deficit replica method* this paper uses for Theorem 1 — the technology of §31.5 and Appendix I is the common engine of the modern era. Second, the functional being minimized is exactly the object this paper calls the GFF ledger (43): area term plus bulk-entropy term, in competition.

L.4 Step 4 — islands and replica wormholes: the Page curve from gravity

Apply the QES prescription to the radiation R of an evaporating hole ([Almheiri et al., 2019](#); [Penington, 2020](#)). At early times the minimal QES is the empty surface: $S(R)$ = the semiclassical radiation entropy, rising exactly as Hawking computed. But there is a competitor: a QES just inside the horizon, whose generalized entropy is (horizon area)/ $4G_N \hbar$ plus a small correction, and which — crucially — assigns to R an *island* I : the bulk region inside the surface, whose state is counted with the radiation,

$$S(R) = \min \left\{ S_{\text{no-island}}(R), \text{ext}_I \left[\frac{A(\partial I) k_B c^3}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I) \right] \right\}. \quad (152)$$

The island branch is flat-to-decreasing (the shrinking horizon area), so the minimum switches branches exactly at the Page time: the curve of Appendix K emerges from geometry, with no unitarity assumption inserted by hand. The 2020 replica computations ([Almheiri et al., 2020](#)) exhibited the mechanism inside the gravitational path integral: connected “replica wormhole” saddles, negligible early, dominate after the Page time and enforce the island answer. What was proven: in two-dimensional dilaton-gravity models coupled to matter (and, with more assumptions,

in higher-dimensional setups), the gravitational path integral itself computes a unitary Page curve. This is the achievement the present paper’s Part IX weighs and the main text engages theorem by theorem.

L.5 Step 5 — the remaining problems, stated precisely

A course honest about triumphs must be precise about residues; these five are the live ones, and the reader will recognize them as the openings the main text walks through. *(i) Dimensional rigor.* The island derivations are theorems in $d = 2$ toy models; in four asymptotically flat dimensions the prescription is motivated, not derived (Almheiri et al., 2021). *(ii) Mechanism vs. prescription.* The QES rule says *which* entropy to compute, not *what physical process* the island represents — what does it mean, operationally, for interior degrees of freedom to “belong” to the radiation? *(iii) The ensemble/factorization puzzle.* Replica wormholes compute averages that behave like an *ensemble* of theories; a single unitary quantum system should have no such average (§33 engages this directly). *(iv) Genericity.* The microscopic counts (Step 1) live in supersymmetric corners; the island results live in low dimensions; the astrophysical Schwarzschild hole — generic, non-extremal, non-supersymmetric — has still never had its microstates named by these methods. *(v) The species/induced- G_N closure.* The $A/4G_N\hbar$ in every formula above uses the observed G_N ; none of the derivations computes G_N from the field content it sums over — the species problem of Appendix G, still open in the modern formalism.

L.6 Step 6 — where GG-Theory enters

The main text’s program is positioned against exactly these residues, and the third-step convention of this course sends the reader there: the microstates for the *generic* hole are named and counted in §31.6 (with the coefficient fixed independently of the count, §31.5); the island functional is derived as the branch-minimized GFF ledger in §33, with the dictionary tabulated in Appendix N; the Page curve arrives from bulk unitarity plus projection in §36 (Theorem 2); the factorization puzzle is engaged as a named commitment in §33; and the species closure is the program’s own declared open computation (Appendix R). Part IX weighs the whole against the frameworks of this appendix in both table and argument.

Problems and worked solutions. L.1.1. Evaluate the Strominger–Vafa degeneracy for the modest charges $Q_1 = Q_5 = n = 10$.

Solution. Step 1: from (149), $S/k_B = 2\pi\sqrt{Q_1Q_5n} = 2\pi\sqrt{1000} \approx 198.7$. Step 2: the number of microstates is $d = e^{S/k_B} \approx e^{198.7} \approx 10^{86}$. Step 3: for scale, that is more states than there are protons in the observable universe ($\sim 10^{80}$) — carried by a single bound state of thirty units of charge, and counted *exactly* by a weak-coupling quantum mechanics. Physics intuition: black-hole entropies are not merely large, they are exponentially large in the charges — which is why no census of states could ever proceed by listing, and why the Cardy asymptotic (a saddle point, not a tally) is the right tool. The astrophysical analogue of this number is the $10^{77} k_B$ of the reference card (Appendix D), counted in this paper’s reading by §31.6.

L.2.1. Show that gravity itself protects the holographic bound: a thermal gas collapses before its entropy can exceed its boundary area.

Solution. Step 1 (natural units, $\hbar = c = k_B = 1$): a photon gas of temperature T in a box of size R has $S_{\text{gas}} \sim (RT)^3$ and $E \sim R^3 T^4$; the area bound allows $S_{\text{max}} \sim R^2/\ell_P^2$. Step 2: the gas would beat the bound when $(RT)^3 > R^2/\ell_P^2$, i.e. $R > 1/(T^3 \ell_P^2)$. Step 3: but the gas’s Schwarzschild radius $r_s \sim G_N E \sim \ell_P^2 R^3 T^4$ exceeds R already when $R > 1/(T^2 \ell_P)$ — and $1/(T^2 \ell_P) < 1/(T^3 \ell_P^2)$ whenever $T \ell_P < 1$, i.e. always below the Planck temperature. The box collapses into a black hole — whose entropy saturates the bound by G.3.1 — before the bound is ever threatened. Physics intuition: the holographic principle needs no enforcement agency; gravity is the enforcement. Try to overfill a region with entropy and you build the very object that defines the limit.

L.3.1. Verify the Ryu–Takayanagi formula against the known entanglement entropy of a two-dimensional critical system.

Solution. Step 1: for an interval of length ℓ in a 2D conformal field theory, the entanglement entropy is $S = \frac{c_{\text{CFT}}}{3} \ln(\ell/\epsilon)$ (ϵ a short-distance cutoff) — a standard, independently derived result of condensed-matter and field theory. Step 2: the bulk dual is AdS_3 with radius L ; the minimal surface anchored on the interval’s endpoints is a geodesic, and a direct integration in the Poincaré metric gives its regulated length $\mathcal{L} = 2L \ln(\ell/\epsilon)$. Step 3: apply (151) with the three-dimensional Newton constant G_3 (natural units): $S = \mathcal{L}/(4G_3) = \frac{2L}{4G_3} \ln(\ell/\epsilon)$, and with the boundary central charge $c_{\text{CFT}} = 3L/2G_3$ (the Brown–Henneaux value) this is exactly $\frac{c_{\text{CFT}}}{3} \ln(\ell/\epsilon)$. Physics intuition: the RT surface converts entanglement — a notoriously hard quantum quantity — into geometry’s easiest one, a length; the match above, exact including coefficient, is why the formula was believed long before it was derived, and the derivation, when it came, used the conical-replica engine of I.4.1.

L.4.1. Show that the replica-wormhole saddle overtakes the disconnected one precisely when the naive radiation entropy exceeds the horizon’s.

Solution. Step 1: for a state of entropy S , $\text{Tr} \rho_R^n \approx e^{-(n-1)S}$; the disconnected saddle realizes this with the semiclassical radiation entropy, $\sim e^{-(n-1)S_{\text{rad}}}$, while the fully connected wormhole pays one horizon area across the whole braid, $\sim e^{-(n-1)S_{\text{BH}}}$ (the area enters as $(n-1)A/4G_N$ through the conical deficit of I.4.1). Step 2: the path integral is dominated by the *larger* contribution, so the wormhole wins exactly when $S_{\text{BH}} < S_{\text{rad}}$. Step 3: the entropy, computed by whichever saddle dominates, is therefore $\min(S_{\text{rad}}, S_{\text{BH}})$: the Page curve of Appendix K, now spoken in saddle-point language, with the crossover at exactly the Page time. Physics intuition: the “exotic” wormhole is doing what sub-dominant saddles always do in statistical mechanics — waiting for its free energy to win; the Page turnover is a first-order phase transition in the replica ensemble, the same grammar as Appendix J (J.2.1) and the same branch switch as (44).

Take-away. *Standard results:* exact microscopic entropy in protected sectors (149); unitarity settled in principle by AdS/CFT ; the Page curve recovered from the gravitational path integral via QES/islands (152). *What each assumed:* supersymmetry and extremality; an AdS box and a dual; two dimensions or motivated prescriptions. *What remains open:* 4D rigor, mechanism, factorization, genericity, and the species closure. *The GG-Theory reading:* each residue is an address in the main text — §31.6, §33, §36, Appendix R — with Part IX as the scorecard.

M Comparison with the islands formula

What this appendix teaches. The islands formula and the GDOS statement set side by side: the QES condition derived, the two-branch race that produces the Page turnover in each formalism, and the precise statement of where the two prescriptions provably agree and where they part — the technical companion to the comparison of Part IX and the history of Appendix L.

Read alongside: [Penington \(2020\)](#) and [Almheiri et al. \(2019\)](#) for the originals; [Almheiri et al. \(2021\)](#) for the review this comparison tracks.

The islands formula is

$$S(R) = \min \text{ext} \left\{ \frac{A(\partial I)}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I) \right\}. \quad (153)$$

In GDOS, the corresponding statement — registered as a dictionary-grade proposition in the named-statement deck (Appendix V) — is

$$S(R) = S_{\text{vN}}(\text{Tr}_{\text{in}}[\rho_{\text{bulk}}]_R) = S_{\text{vN}}(\rho_{\partial}^R), \quad (154)$$

where ρ_{∂}^R is the 4D reduced state of the radiation. This chapter works the mechanism in two Steps and then records the equivalence and its limits.

M.1 Step 1 — the QES condition, derived, and the two-branch race

Write the generalized entropy of a candidate island with boundary at radial position x as

$$S_{\text{gen}}(x; t) = \frac{A(x)}{4G_N \hbar} + S_{\text{semi-cl}}(R \cup I(x); t). \quad (155)$$

The extremum condition defining the quantum extremal surface is

$$\frac{\partial S_{\text{gen}}}{\partial x} = 0 \quad \Longleftrightarrow \quad \frac{A'(x)}{4G_N \hbar} = - \frac{\partial S_{\text{semi-cl}}}{\partial x} : \quad (156)$$

the surface settles where the marginal area cost of enlarging the island exactly cancels the marginal entanglement saved by including more interior partners. Since A' is set by the horizon scale while $\partial_x S_{\text{semi-cl}}$ is set by the near-horizon mode density, the solution pins ∂I to just inside the horizon — which is why the island’s area term is $\approx 2S_{\text{BH}}$ for the two-sided problem and $\approx S_{\text{BH}}$ (the instantaneous horizon) for evaporation.

Now race the two branches of the min:

- *No-island branch* ($I = \emptyset$): $S(R) = S_{\text{semi-cl}}(R; t)$, the naive radiation entropy, growing linearly — each emitted quantum adds its thermal share, $\dot{S} \approx \text{const}$ (the rising Page edge).
- *Island branch*: $S(R) \approx S_{\text{BH}}(t) + S_{\text{matter}}$, essentially time-independent on emission timescales (and slowly *decreasing* as the hole shrinks), because once the interior partners are inside I they are “kept,” not traced, and the only cost is the area at ∂I .

A linearly growing branch racing a flat one guarantees a crossing; the min switches branches at exactly the Page point of Appendix K, and the late-time entropy rides $S_{\text{BH}}(t)$ down the falling edge. This is the replica-wormhole result (Almheiri et al., 2019, 2020; Penington, 2020), stated as a ledger race.

M.2 Step 2 — why the two formulations agree

At sub-Page time, the island $I = \emptyset$, the bulk is traced in the normal way, and the radiation entropy is $S_R \sim m \ln 2$ (Hawking result). At super-Page time, the island I picks up bulk volume deep inside the horizon, and $S(R)$ starts decreasing.

The equivalence of the two formulations:

1. The island boundary ∂I is the boundary in the bulk between “traced-out” and “kept” regions.
2. The $A(\partial I)/(4G_N \hbar)$ term is the gravitational entropy associated with the partial-trace structure.
3. The $S_{\text{semi-cl}}$ term is the remaining entropy of the semi-classical state on $R \cup I$.
4. Extremizing over ∂I is the same as optimizing the partial-trace structure to match the minimal von Neumann entropy.

In GG-Theory, the 6D bulk is explicitly unitary, and the OGGEp partial trace is a specific operation. The islands formula is the semi-classical approximation of the OGGEp output, and the QES (∂I) is the boundary of the traced-over region. The precise proposition and its grade — where the equivalence is derived, where it is a dictionary entry — are registered in the named-statement deck (Appendix V); the reading is a translation, not a competing mechanism, and Appendix N tabulates it term by term.

Physical picture. The island prescription looks exotic — interior regions suddenly counted as part of the radiation — until it is read as bookkeeping: the min-ext is deciding, ledger in hand, which purchase is cheaper: keep paying thermal rent on every emitted quantum’s hidden partner, or buy the interior outright at the fixed price $A/4G_N \hbar$. At the Page point the mortgage beats the rent, and the books — which were always unitary in the 6D bulk — start showing it in 4D.

Problems and worked solutions. M.1.1. Work the two-branch race with constant rates, find the corner, and explain why the true answer (0.646) differs.

Solution. Step 1: model the no-island branch as $S_{\text{no-isl}}(t) = \lambda t$ and the island branch as $S_{\text{isl}}(t) = S_0 - \mu t$, with constant rates; the min switches at $t_* = S_0/(\lambda + \mu)$. Step 2: for thermal emission, G.4.1 gives $\lambda = \frac{4}{3}\mu$ (radiation gains 4/3 of what the hole loses), and the hole spends S_0 over t_{evap} , so $\mu = S_0/t_{\text{evap}}$: hence $t_* = \frac{3}{7} t_{\text{evap}} \approx 0.43 t_{\text{evap}}$. Step 3: the true ratio is $1 - 2^{-3/2} = 0.646$ (148) because the rates are *not* constant: $T \propto 1/M$ makes the early hole cold and slow, back-loading the entropy spend; the constant-rate toy front-loads it. The corner’s existence is rate-independent; its location is a clock question. Physics intuition: every version of the Page curve is the same race between a rising rent and a falling mortgage (the ledger picture of the text); the schedules differ, the crossing does not.

M.1.2. Explain why the island’s area term is $\approx 2S_{\text{BH}}$ for the eternal two-sided problem but $\approx S_{\text{BH}}$ for evaporation.

Solution. Step 1: for the eternal (two-sided) black hole with both radiation baths collected, the island is a spacelike slab through the interior whose boundary ∂I has *two* components — one surface just inside each horizon — so the area term in (153) pays twice: $2 \times A/4G_N\hbar \approx 2S_{\text{BH}}$. Step 2: for a one-sided evaporating hole there is a single horizon, a single QES component, and a single payment $\approx S_{\text{BH}}(t)$ at the *instantaneous* area. Step 3: correspondingly the eternal case’s entropy saturates at $2S_{\text{BH}}$ (the two-sided thermofield double’s total) while the evaporating case’s island branch tracks the shrinking horizon down to zero — which is what makes the curve fall. Physics intuition: the area term counts the *doors* between the kept and traced regions; the eternal hole has two doors, the dying hole one, and the ledger charges per door.

M.2.1. Derive the equivalence claim: extremizing S_{gen} is branch-minimizing the GFF ledger, row by row.

Solution. Step 1: the QES functional (155) is $A(x)/4G_N\hbar + S_{\text{semi-cl}}$: term one is the geometric free-energy piece F_{GEP} and term two the entropic piece F_{CEP} of the ledger (43) — the identification of Table 20, rows one and two. Step 2: extremization over the surface position x (156) is stationarity of the ledger over the admissible coarse-grained saddles, i.e. (44) restricted to the one-parameter family the island geometry exposes; the min over extrema is the ledger’s branch minimum. Step 3: hence wherever the island family exhausts the admissible saddles — the regime stated in the text — the two prescriptions compute the same number by construction; where the ledger admits saddles outside the island family (the DGI-constrained reconstructions of §33), the GG-Theory reading is the wider net. Physics intuition: the dictionary is not an analogy but an accounting identity — two bookkeeping systems posting the same two costs; they can only disagree where one system has ledger lines the other lacks, and that residue is precisely what Part VII prices.

Take-away. *Standard result:* the QES condition and the island rule (153). *The GDOS statement:* (154), the entropy of the projected state. *Where they provably agree:* whenever the branch-minimized GFF ledger is the dominant saddle of the island extremization — the regime tabulated in Appendix N. *Where they part:* provenance — prescription within the path integral vs. theorem of the projection — and the parting is testable only through the exterior corrections priced in Part VII.

N Island/QES and GG-Theory language

What this appendix teaches. The translation table: every term of the island/QES program — generalized entropy, extremization, islands, replica wormholes, bulk reconstruction — next to its GG-Theory/GDOS reading, one line each, with pointers to the equations that justify each row.

The island/quantum-extremal-surface program and the GG-Theory reading are translations of one another in the regime where both apply; the table records the dictionary.

Table 20: Island/QES terms and their GG-Theory/GDOS readings. Five rows, and the translation runs one way only: every island-program term is given a reading in this framework’s ledger, and no claim is made that the reverse holds. The purpose is to let a reader fluent in the island literature check this Part’s derivations against machinery they already trust, and to make plain where the two are doing the same arithmetic under different names.

Island / QES language	GG-Theory / GDOS reading
Generalized entropy $S_{\text{gen}} = A/4G + S_{\text{bulk}}$	Black-hole instance of the universal ledger $\text{GFF} = F_{\text{GEP}} + F_{\text{CEP}}$ (43).
Extremize/minimize over candidate QES	Branch-minimize the ledger over admissible coarse-grained saddles (44).
Appearance of an island in the radiation wedge	Late-time dominance of the area-limited bridge branch (§36).
Replica wormhole saddle	Sub-dominant-to-dominant branch exchange in the GFF ledger.
Quantum error correction of the bulk	Non-naive, DGI-constrained reconstruction of the accessible algebra (§33).

O Entropy balance, channels, and data processing

What this appendix teaches. The quantum-information engine room: reduction as a CPTP channel (Kraus and Stinespring, constructed), the data-processing inequality, Spohn’s entropy-production identity worked in full, and the single-mode exchange ledger in closed form — the machinery behind the paper’s exact entropy-production statements.

Read alongside: Harlow (2016) for the channel toolkit; Spohn (1978) for the production identity; Wall (2012) for the GSL this machinery feeds.

The black-hole story requires distinguishing fine-grained from operational entropy. Under closed unitary evolution the global von Neumann entropy is constant, $S(\rho_{\text{global}}(t)) = S(\rho_{\text{global}}(0))$, yet a reduced state changes because reduction is a channel, and in GDOS language $\rho_{\partial}(t) = \Pi_t(\rho_6(0))$ with Π_t CPTP. The operational entropy at the boundary is therefore the entropy of the channel *image*, not of “reality itself.” This is why data-processing inequalities are physical, not technical: for any channel $\mathcal{N}$, relative entropy contracts, $S(\mathcal{N}\rho \parallel \mathcal{N}\sigma) \leq S(\rho \parallel \sigma)$, so accessible mutual information is bounded by the parent theory’s. The boundary description may lose information without the bulk destroying it — the precise content of “information is preserved, but not by the 4D observer alone.”

That paragraph is the chapter’s thesis; the four Steps below equip it. This is the mathematical spine of Lemma 2.2 (L-BH-Arrow) and Corollary 2.1 (C-BH-1), worked at the graduate level.

O.1 Step 1 — reduction is a channel: Kraus and Stinespring

Let the parent system evolve unitarily on $\mathcal{H}_S \otimes \mathcal{H}_E$ from a product start $\rho \otimes |0\rangle\langle 0|_E$, and let the observer keep only S . Then

$$\Pi(\rho) = \text{Tr}_E[U(\rho \otimes |0\rangle\langle 0|_E)U^\dagger] = \sum_k K_k \rho K_k^\dagger, \quad K_k \equiv \langle k|_E U |0\rangle_E, \quad (157)$$

by inserting a basis $\{|k\rangle_E\}$ in the trace; completeness of the basis and unitarity of U give $\sum_k K_k^\dagger K_k = \mathbb{K}$ (trace preservation), and the manifest form $\sum_k K_k \cdot K_k^\dagger$ is complete positivity. Stinespring’s theorem is the converse: *every* CPTP map arises this way from some dilation $(U, \mathcal{H}_E)$. So “the observer’s description” and “a channel” are the same mathematical object — exactly the GDOS statement that $\rho_\partial(t) = \Pi_t(\rho_6(0))$, with the horizon making the environment factor causally compulsory (L-BH-Port, §25.2).

O.2 Step 2 — the contraction: data processing

The data-processing inequality (DPI) states that for every channel $\mathcal{N}$ and every pair of states,

$$S(\mathcal{N}\rho \parallel \mathcal{N}\sigma) \leq S(\rho \parallel \sigma), \quad S(\rho \parallel \sigma) \equiv \text{Tr } \rho (\ln \rho - \ln \sigma). \quad (158)$$

Its proof pedigree is worth recording once: monotonicity under partial trace is equivalent to strong subadditivity of the von Neumann entropy, which rests on Lieb’s concavity theorem; the general-channel case follows by Stinespring — adjoin the environment (relative entropy is additive under $\otimes$ with a fixed ancilla), act unitarily (invariance), and trace out (partial trace case). Every arrow statement in this paper descends from (158); its QFT-grade formulation, needed when the subsystems are relativistic wedges rather than tensor factors, is the algebraic monotonicity used by Casini (2008) and by Wall (2012) — the ceiling declared in Lemma 2.2’s Wall-class clause.

Physical picture. Relative entropy is the ledger distance between “what the state is” and “what the observer’s best hypothesis says”: operationally, the exponential rate at which the hypothesis can be rejected. A channel can only blur the evidence — no post-processing manufactures distinguishability — so the distance contracts. Everything irreversible in this paper is this one sentence wearing different clothes.

O.3 Step 3 — Spohn’s production identity, worked

Let $\Phi_t = e^{t\mathcal{L}}$ be a quantum dynamical semigroup with stationary state σ ($\Phi_t\sigma = \sigma$) — here, the horizon’s KMS face at T_H (Theorem 5). Two lines give the theorem. First, since $\text{Tr } \dot{\rho}_t = 0$,

$$\frac{d}{dt} S(\rho_t \parallel \sigma) = \text{Tr} [\dot{\rho}_t (\ln \rho_t - \ln \sigma)] \quad (159)$$

(the term $\text{Tr} [\rho_t \frac{d}{dt} \ln \rho_t] = \text{Tr } \dot{\rho}_t = 0$ drops). Second, apply the DPI (158) to the channel Φ_ϵ :

$$S(\rho_{t+\epsilon} \parallel \sigma) = S(\Phi_\epsilon \rho_t \parallel \Phi_\epsilon \sigma) \leq S(\rho_t \parallel \sigma) \implies \varsigma(t) \equiv -\frac{d}{dt} S(\rho_t \parallel \sigma) \geq 0. \quad (160)$$

This is Spohn’s theorem (Spohn, 1978): the entropy production ς of a semigroup with a stationary face is pointwise non-negative, vanishing iff the state sits on the face. Lemma 2.2’s three named links are visible here in miniature: the semigroup direction is *consumed* (it is the arrow custody of AX-3/P5-IX/A5, not derived from thermodynamics), the DPI supplies the contraction, and the horizon translation identifies σ with the KMS face and ς with GSL production.

O.4 Step 4 — the single-mode exchange ledger, in closed form

The exchange-class identity of Lemma 2.2 — ΔS_{gen} equals the relative entropy spent toward the KMS face — can be displayed mode by mode. For one bosonic mode, a thermal state of occupation n is $\rho_n = (1+n)^{-1} \left(\frac{n}{1+n}\right)^{a^\dagger a}$, so

$$\ln \rho_n = a^\dagger a \ln \frac{n}{1+n} - \ln(1+n), \quad S_{\text{vN}}(\rho_n) = (1+n) \ln(1+n) - n \ln n. \quad (161)$$

Let the mode arrive in a thermal state at the wrong temperature — occupation $\bar{n}_b$ — and let the face sit at $\bar{n}_H = (e^{\hbar\omega/k_B T_H} - 1)^{-1}$. Assemble the relative entropy term by term:

$$\begin{aligned} S(\rho_{\bar{n}_b} \parallel \rho_{\bar{n}_H}) &= -S_{\text{vN}}(\rho_{\bar{n}_b}) - \text{Tr}[\rho_{\bar{n}_b} \ln \rho_{\bar{n}_H}] \\ &= -(1+\bar{n}_b) \ln(1+\bar{n}_b) + \bar{n}_b \ln \bar{n}_b - \bar{n}_b \ln \frac{\bar{n}_H}{1+\bar{n}_H} + \ln(1+\bar{n}_H) \\ &= \bar{n}_b \ln \frac{\bar{n}_b}{\bar{n}_H} - (1+\bar{n}_b) \ln \frac{1+\bar{n}_b}{1+\bar{n}_H}. \end{aligned} \quad (162)$$

Positivity is one derivative: $\partial_{\bar{n}_b} S = \ln[\bar{n}_b(1+\bar{n}_H)/\bar{n}_H(1+\bar{n}_b)]$ vanishes iff $\bar{n}_b = \bar{n}_H$, and $\partial_{\bar{n}_b}^2 S = 1/\bar{n}_b - 1/(1+\bar{n}_b) > 0$, so (162) is strictly convex with minimum value 0 exactly on the face. Read through the identity of Lemma 2.2, this is the per-mode GSL: the generalized entropy gained in an exchange is (162), non-negative always, and *zero* iff $T_b = T_H$ — which is Corollary 2.1’s statement that the conditional-GGA crossover is the GSL’s equality point. (The vacuum limit $\bar{n}_b \rightarrow 0$ gives $S \rightarrow \ln(1+\bar{n}_H)$, the per-emitted-mode closed form quoted in the main text.) The identity, the positivity grid, and the closed forms are verified to machine precision in the `s4_arrow` suite (Appendix U); the formula (162) itself is re-checked in `p2c2_verify`.

Physical picture. The horizon face at T_H is the sea level of the entropy ledger. A mode exchanged with the hole at any other temperature carries a measurable “height” above sea level — the relative entropy (162) — and the GSL is the statement that exchanges only ever spend height, never mint it. Equilibrium at exactly T_H is sea level itself: the unique exchange that costs nothing, and the crossover point where the conditional-GGA bound is saturated.

Problems and worked solutions. O.1.1. Prove that the Kraus form is trace-preserving and completely positive.

Solution. Step 1 (TP): from (157), $\text{Tr} \Pi(\rho) = \text{Tr}[\sum_k K_k^\dagger K_k \rho]$, and $\sum_k K_k^\dagger K_k = \sum_k \langle 0|U^\dagger|k\rangle \langle k|U|0\rangle_E = \langle 0|U^\dagger U|0\rangle_E = \mathbf{1}$ by completeness and unitarity: traces are preserved. Step 2 (CP): extend by an ancilla A of any dimension; $(\Pi \otimes \text{id}_A)$ acting on a positive ρ_{SA} gives $\sum_k (K_k \otimes \mathbf{1}) \rho_{SA} (K_k \otimes \mathbf{1})^\dagger$, a sum of terms of the form $X \rho X^\dagger \geq 0$: positivity survives the extension, which is the definition of complete positivity. Step 3: Stinespring’s converse (quoted in the text) closes the loop: *every* CP, trace-preserving map has this dilated form — observer descriptions and channels are one class. Physics intuition: complete positivity is the statement that a physical evolution must remain physical even when the system is secretly entangled with a bystander — exactly the situation of every black-hole interior register, which is why CPTP is the non-negotiable grammar of the whole

paper.

O.2.1. Prove the two easy legs of the data-processing pedigree: unitary invariance and stability under fixed ancillas.

Solution. Step 1 (unitaries): relative entropy (158) is a trace of operator functions of its two arguments, and $\text{Tr } f(U\rho U^\dagger) = \text{Tr } f(\rho)$: hence $S(U\rho U^\dagger \| U\sigma U^\dagger) = S(\rho \| \sigma)$ — reversible processing costs nothing. Step 2 (ancillas): for product extensions, $\ln(\rho \otimes \tau) = \ln \rho \otimes \mathbf{1} + \mathbf{1} \otimes \ln \tau$, so $S(\rho \otimes \tau \| \sigma \otimes \tau) = S(\rho \| \sigma) + \text{Tr } \tau(\ln \tau - \ln \tau) = S(\rho \| \sigma)$: attaching an uncorrelated bystander adds nothing. Step 3: the third leg — monotonicity under partial trace, the hard one resting on Lieb’s concavity — is exactly what the text imports; combining the three gives contraction under every Stinespring dilation, i.e. under every channel. Physics intuition: the two easy legs say distinguishability is a property of *information*, not of representation; the hard leg says forgetting can only blur it. Everything irreversible in this paper is the hard leg wearing different clothes.

O.4.1. Check the three limits of the closed-form exchange ledger (162).

Solution. Step 1 ($\bar{n}_b \rightarrow \bar{n}_H$): both terms of (162) vanish — $\bar{n}_b \ln(\bar{n}_b/\bar{n}_H) \rightarrow 0$ and the second term’s argument $\rightarrow 1$: zero cost on the face, as Spohn’s theorem requires (equilibrium is sea level). Step 2 ($\bar{n}_b \rightarrow 0$, the vacuum limit): the first term dies and the second gives $-\ln[1/(1 + \bar{n}_H)] = \ln(1 + \bar{n}_H)$ — the per-emitted-mode closed form quoted in the main text. Step 3 ($\bar{n}_b \rightarrow \infty$): the two logarithms combine and the ledger grows as $\bar{n}_b \ln[(1 + \bar{n}_H)/\bar{n}_H]$ — linearly in the occupation: overdriving a mode far above the face costs proportionally to how much is sent. Physics intuition: the ledger is a distance-from-equilibrium meter with the right three readings — zero on the face, finite for the vacuum, unbounded for a blowtorch — and the GSL is the statement that the meter never runs backwards.

O.4.2. Price a typical Hawking quantum: the ledger cost of one vacuum-limit mode at $\hbar\omega = k_B T_H$.

Solution. Step 1: the face occupation at $\hbar\omega = k_B T_H$ is $\bar{n}_H = 1/(e - 1) \approx 0.582$. Step 2: the vacuum-limit cost (O.4.1, Step 2) is $\ln(1 + \bar{n}_H) = \ln[e/(e - 1)] \approx 0.459 \text{ nat} \approx 0.66 \text{ bit per mode}$. Step 3: this is the per-mode microscopic face of the global $\frac{1}{3}$ margin of G.4.1 — summed over the emitted spectrum, the per-mode ledger reproduces the $\frac{4}{3}$ radiation-to-hole entropy ratio (the integral is logged in the benchmark audit’s `s4_arrow` suite). Physics intuition: each Hawking quantum leaves carrying about two-thirds of a bit of irreversibility — the price the exterior pays for reading the bulk through a thermal port, mode by mode, exactly accounted.

Take-away. *Standard results:* Stinespring dilation, data processing, Spohn’s identity, and the exchange-class entropy production in closed form. *What they assume:* only that reduction is a channel — no gravity anywhere. *The GG-Theory reading:* at a horizon these theorems stop being effective descriptions and become exact, because the channel is geometrically forced; the GSL’s direction is inherited, not assumed.

P Extremal black holes and the third law

What this appendix teaches. The charged, rotating sector where black-hole thermodynamics earns its Legendre structure: the Kerr–Newman horizon data, the area and the Christodoulou mass split, the first law and the Smarr formula derived, the third law in its GDOS reading, and superradiance worked as the rotating anchor’s instance.

Read alongside: Chandrasekhar (1983) chs. 6–7 and Frolov and Novikov (1998) chs. 6–8 for the full classical theory; Poisson (2004) ch. 5 for the mechanics.

Rotation and charge are where black-hole thermodynamics earns its Legendre structure: new charges, new potentials, a Smarr formula, and a third law. This chapter works the Kerr–Newman horizon data (Step 1), the area and the Christodoulou mass (Step 2), the first law and Smarr formula (Step 3), the third law in its GDOS reading (Step 4), and superradiance as the worked instance of the rotating anchor (Step 5). Geometrized units $G_N = c = 1$ except where displayed.

P.1 Step 1 — the Kerr–Newman horizon data

The Kerr–Newman family (mass M , spin $a = J/M$, charge Q) has horizon structure governed by

$$\Delta(r) = r^2 - 2Mr + a^2 + Q^2 = (r - r_+)(r - r_-), \quad r_{\pm} = M \pm \sqrt{M^2 - a^2 - Q^2}, \quad (163)$$

so horizons exist iff $a^2 + Q^2 \leq M^2$, with equality the *extremal* case $r_+ = r_-$. The surface gravity and horizon angular velocity are

$$\kappa = \frac{r_+ - r_-}{2(r_+^2 + a^2)} = \frac{\sqrt{M^2 - a^2 - Q^2}}{r_+^2 + a^2}, \quad \Omega_H = \frac{a}{r_+^2 + a^2}, \quad (164)$$

the first equality the standard evaluation of $\kappa^2 = -\frac{1}{2}(\nabla^\mu \xi^\nu)(\nabla_\mu \xi_\nu)|_{r_+}$ for the horizon generator $\xi = \partial_t + \Omega_H \partial_\phi$ (equivalently $\kappa = \Delta'(r_+)/2(r_+^2 + a^2)$). At extremality ($a^2 + Q^2 = M^2$) the surface gravity vanishes: $\kappa = 0$. The Hawking temperature correspondingly vanishes: $T_H = 0$ — a thermodynamic system with zero temperature but, as Step 2 shows, non-zero entropy.

P.2 Step 2 — area, spin, and the Christodoulou mass

For Kerr ($Q = 0$), the horizon area is

$$A_{\text{Kerr}} = 4\pi(r_+^2 + a^2) = \frac{8\pi G_N^2 M^2}{c^4} \left[1 + \sqrt{1 - a^{*2}}\right], \quad (165)$$

where $r_+ = G_N M/c^2 + \sqrt{(G_N M/c^2)^2 - a^2}$ is the outer horizon radius and $a^* = ac^2/G_N M$ is the dimensionless spin. The first equality reduces via one identity worth deriving: r_+ solves $r_+^2 - 2Mr_+ + a^2 = 0$, so

$$r_+^2 + a^2 = 2Mr_+ = 2M(M + \sqrt{M^2 - a^2}) = 2(M^2 + \sqrt{M^4 - J^2}), \quad (166)$$

using $a = J/M$. Extremal Kerr has $a^* = 1$ and $A^{\text{ext}} = 8\pi G_N^2 M^2/c^4$, one-half the Schwarzschild value. Inverting for M : from $A = 8\pi(M^2 + \sqrt{M^4 - J^2})$, isolate the root and square,

$$\left(\frac{A}{8\pi} - M^2\right)^2 = M^4 - J^2 \implies \boxed{M^2 = \frac{A}{16\pi} + \frac{4\pi J^2}{A}} \quad (167)$$

— the Christodoulou mass formula. Its first term defines the irreducible mass, $M_{\text{irr}}^2 = A/16\pi$: the part of the mass that the area theorem forbids extracting. The second term is rotational energy, extractable in principle (Penrose process, superradiance — Step 5) down to $M \rightarrow M_{\text{irr}}$; for extremal Kerr, up to $1 - 1/\sqrt{2} \approx 29\%$ of the mass.

P.3 Step 3 — the first law and the Smarr formula

Differentiate the Christodoulou formula at fixed J and fixed A respectively:

$$\left(\frac{\partial M}{\partial A}\right)_J = \frac{1}{2M} \left(\frac{1}{16\pi} - \frac{4\pi J^2}{A^2}\right) = \frac{\kappa}{8\pi}, \quad \left(\frac{\partial M}{\partial J}\right)_A = \frac{4\pi J}{AM} = \Omega_H, \quad (168)$$

where both right-hand equalities are checked against (164) by direct substitution (and to 10^{-12} numerically in `p2c2_verify`). Hence the first law

$$dM = \frac{\kappa}{8\pi} dA + \Omega_H dJ = T_H dS_{\text{BH}} + \Omega_H dJ, \quad (169)$$

identifying $S_{\text{BH}} = A/4$ once $T_H = \kappa/2\pi$ is fixed by Appendix H. The Smarr formula follows from scaling alone: under $M \rightarrow \lambda M$ every length doubles with λ , so $A \rightarrow \lambda^2 A$ and $J \rightarrow \lambda^2 J$, i.e. $M(A, J)$ is homogeneous of degree $\frac{1}{2}$; Euler's theorem gives

$$M = 2 \frac{\partial M}{\partial A} A + 2 \frac{\partial M}{\partial J} J = 2 T_H S_{\text{BH}} + 2 \Omega_H J, \quad (170)$$

the algebraic shadow of the first law, with the factor 2 the homogeneity degree's inverse. (With charge, $+\Phi_H Q$ joins the first law and $M = 2T_H S_{\text{BH}} + 2\Omega_H J + \Phi_H Q$: the Q -term enters with weight one because $Q \rightarrow \lambda Q$.)

Physical picture. The first law is local bookkeeping at the horizon; the Smarr formula is the same books audited globally, and the audit works because gravity has no intrinsic scale — homogeneity is the entire input. That thermodynamic identities follow from dimensional analysis is the first hint that the entropy is geometric.

P.4 Step 4 — the third law, GDOS-read

In GDOS, the third law of thermodynamics (Nernst-unattainability) is automatic: reaching the extremal state requires the scale profile at the horizon to vanish $e^{A(r_h)} \rightarrow 0$, which corresponds to infinite scale depth and hence cannot be achieved in finite external time. The classical statement is the same limit read as geometry: as $a^2 + Q^2 \rightarrow M^2$, the two roots $r_{\pm}$ merge, $\kappa \rightarrow 0$ by (164), and the Euclidean cigar of Appendix I degenerates — its tip recedes to infinite proper distance, so

there is no smooth cap and no regularity argument: extremal holes sit outside the thermal ensemble, reachable only asymptotically.

Extremal black holes are interesting for microstate-counting arguments because their entropy survives at $T = 0$: $S_{\text{BH}}^{\text{ext}} = A^{\text{ext}} c^3 / (4G_N \hbar)$, giving e^{S_{BH}/k_B} microstates in the ground state. Strominger–Vafa counted these microstates explicitly for certain 5D extremal black holes using D-brane configurations in string theory (Strominger and Vafa, 1996). In GG-Theory, the counting is via $U(1)_{Qi}$ holonomies on S_θ^1 at each horizon pixel.

P.5 Step 5 — superradiance, the worked instance

The Bogoliubov transformation near a rotating horizon involves both the surface gravity κ and the horizon angular velocity Ω_H ; the resulting spectrum is thermal with chemical potential:

$$\langle N_{\omega \ell m} \rangle = \frac{1}{e^{(\hbar\omega - m\hbar\Omega_H)/k_B T_H} - 1}, \quad (171)$$

and one gets superradiant emission for $\omega < m\Omega_H$. The mechanism is one boundary condition: a mode $e^{-i\omega t + im\phi}$ near the horizon depends on the generator’s combination $\tilde{\omega} \equiv \omega - m\Omega_H$ (the frequency in the co-rotating frame), entering the horizon as $e^{-i\tilde{\omega} r^*}$. The conserved energy flux through the horizon is proportional to $\omega \tilde{\omega} |T_{\omega \ell m}|^2$, so for $0 < \omega < m\Omega_H$ the “absorbed” flux is *negative*: the hole pays energy into the mode, the reflected wave exceeds the incident one, and the occupation formula above encodes the same fact as a pole crossing ($\langle N \rangle < 0$ formally flags the amplified channel). The structural theorem behind this reading — the rotating-frame selection failure and superradiance as a symplectic pumping channel on the asymptotic anchor — is T-BH-3, §27. In GDOS, the superradiance is the 4D observer seeing energy flow *out* of the black hole via modes whose bulk phase-space volume extends outside the rotating ergosphere.

Physical picture. A spinning hole is a flywheel wearing a one-way membrane: modes that rotate slower than the horizon drag against the flywheel and are paid for the friction. The Christodoulou split of Step 2 says exactly how much is in the till — everything above M_{irr} — and the area theorem is the statement that the till never absorbs from the irreducible account.

Problems and worked solutions. P.1.1. Check the Kerr–Newman surface gravity (164) in its three instructive limits.

Solution. Step 1 (Schwarzschild, $a = Q = 0$): $r_\pm = M \pm M$, so $r_+ = 2M$, $r_- = 0$, and $\kappa = (r_+ - r_-)/2(r_+^2 + a^2) = 2M/8M^2 = 1/4M$ — the geometric form of (91). Step 2 (extremal, $a^2 + Q^2 = M^2$): $r_+ = r_-$ and $\kappa = 0$ identically, whatever the split between spin and charge: extremality is a property of the double root, not of the route to it. Step 3 (slow spin, $Q = 0$, $a \ll M$): with $x = a^2/M^2$, the numerator of (164) is $M(1 - x/2) + O(x^2)$ and the denominator $4M^2(1 - x/4) + O(x^2)$, so $\kappa \approx \frac{1}{4M}(1 - a^2/4M^2)$: rotation *cools* the hole at second order, the first hint of the third law’s asymptote. Physics intuition: κ measures the gap between the two roots of Δ — the horizon and its inner partner — and temperature dies exactly as the two merge; extremality is the thermodynamic zero written as a coincidence of radii.

P.2.1. How much of an extremal Kerr hole’s mass is extractable in principle?

Solution. Step 1: for extremal Kerr ($a = M$, $Q = 0$) the horizon area (165) is $A = 8\pi M^2$ (half the Schwarzschild value for the same mass). Step 2: the irreducible mass is $M_{\text{irr}} = \sqrt{A/16\pi} = M/\sqrt{2}$: the area theorem forbids touching it. Step 3: the extractable fraction is $1 - 1/\sqrt{2} \approx 29.3\%$ of the total mass-energy — by Penrose processes or superradiance, in principle down to a Schwarzschild remnant of mass $M/\sqrt{2}$. For comparison: nuclear fusion frees 0.7%, ISCO accretion 5.7% (E.3.3). Physics intuition: a maximally spinning hole is the universe’s best battery — nearly a third of its rest energy sits in the till as rotational energy above the irreducible account, and the area theorem is the banking regulation that says which withdrawals are lawful.

P.3.1. Extend the Smarr formula to include charge by the same scaling argument.

Solution. Step 1: under $M \rightarrow \lambda M$ the scaling weights are $A \rightarrow \lambda^2 A$, $J \rightarrow \lambda^2 J$ (as in the text), and $Q \rightarrow \lambda Q$ — charge carries length dimension one in geometrized units. Step 2: Euler’s theorem for a function of weight 1 in arguments of weights (2, 2, 1) reads $M = 2 \frac{\partial M}{\partial A} A + 2 \frac{\partial M}{\partial J} J + \frac{\partial M}{\partial Q} Q$. Step 3: inserting the first law (168) identifications gives $M = 2T_H S_{\text{BH}} + 2\Omega_H J + \Phi_H Q$ (170): the charge term enters with weight one, exactly as the text’s parenthetical states. Physics intuition: the Smarr formula is dimensional analysis promoted to a theorem — gravity has no intrinsic scale, so the first law’s local bookkeeping must integrate to a global identity, with each charge weighted by how it scales.

P.5.1. Compute the superradiant window for a $10 M_\odot$ hole at $a_* = 0.9$.

Solution. Step 1: the horizon radius is $r_+ = M(1 + \sqrt{1 - a_*^2}) = 1.436 M$ (geometrized), and the horizon angular velocity $\Omega_H = a/(r_+^2 + a^2) = 0.313/M$. Step 2: restore units with $1/M \rightarrow c^3/(G_N M)$: for $M = 10 M_\odot$, $\Omega_H \approx 6.4 \times 10^3 \text{ rad s}^{-1}$. Step 3: the $m = 1$ superradiant condition $\omega < m \Omega_H$ (171) amplifies waves below $f = \Omega_H/2\pi \approx 1.0 \text{ kHz}$ — squarely the ground-based gravitational-wave band, which is why superradiance physics is not exotica but part of the waveform modeling for the ringdowns Part VIII scores. Physics intuition: the hole is a flywheel spinning a thousand times a second; any wave patient enough to rotate more slowly gets dragged, and paid — the one-way membrane runs a toll in reverse for slow traffic.

Take-away. *Standard results:* the Kerr–Newman horizon data, M_{irr} , the first law with work terms, the Smarr formula, the third law, and superradiance. *What they assume:* classical stationarity and cosmic censorship. *The GG-Theory reading:* the work terms are the ledger’s chemical potentials; extremality is the port never quite closing; superradiance is the rotating anchor paying friction through the one-way membrane.

Q De Sitter horizon and cosmology

What this appendix teaches. The same theorem with the observer on the inside: the static patch and its surface gravity, the temperature and entropy of the cosmological horizon with units restored, and the GDOS reading that makes the de Sitter entropy the black-hole theorem evaluated on a different real section — the cohort’s bridge to the cosmology of GG-A10-Cosmos and GG-C2-Hubble.

Read alongside: Gibbons and Hawking (1977) for the original; Arizaka (2026, A10-Cosmos) Part VII for the cosmological consumer.

The de Sitter horizon is the same theorem with the observer on the inside, and it is the cohort's bridge to cosmology. Step 1 works the static patch; Step 2 restores units and evaluates the two famous numbers; Step 3 gives the GDOS reading.

Q.1 Step 1 — the static patch and its surface gravity

A de Sitter space (Hubble expansion with $H = \text{const}$, positive cosmological constant $\Lambda = 3H^2$) admits, around any observer, static coordinates

$$ds^2 = -(1 - H^2 r^2) dt^2 + (1 - H^2 r^2)^{-1} dr^2 + r^2 d\Omega_2^2, \quad (172)$$

with $f(r) = 1 - H^2 r^2$ vanishing at the cosmological horizon $r_H = 1/H$ ($= c/H$ with units). The observer sits at the *center*; the horizon surrounds her. The surface gravity follows exactly as in Appendix H, Step 4:

$$\kappa = \frac{1}{2} |f'(r_H)| = \frac{1}{2} \cdot 2H^2 r_H = H, \quad (173)$$

and the near-horizon metric is again Rindler ($f = \kappa^2 x^2$ in the proper distance x from the horizon — the same two-line computation), so the whole Bogoliubov/continuation apparatus of that chapter, and the cigar regularity of Appendix I (the Euclidean section is the round four-sphere, whose smoothness fixes the period $2\pi/H$), apply verbatim.

Q.2 Step 2 — temperature and entropy, with units

The Gibbons–Hawking temperature and entropy are

$$k_B T_{\text{GH}} = \frac{\hbar H}{2\pi}, \quad \frac{S_{\text{dS}}}{k_B} = \frac{\pi r_H^2}{G_N \hbar / c^3} = \frac{\pi c^5}{G_N \hbar H^2} \quad (174)$$

(Gibbons and Hawking, 1977) — the first line is $T_H = \hbar c \kappa_{\text{geom}} / 2\pi k_B$ with $\kappa_{\text{geom}} = H/c$, the second is $A_H / 4\ell_P^2$ with $A_H = 4\pi r_H^2$ and $\ell_P^2 = G_N \hbar / c^3$. For our universe — taking the derived background rate $H_0 \simeq 66.6 \text{ km/s/Mpc}$ (the two-road reading the companions carry; the distance ladder's 73.5 is read as a calibration systematic, not a second rate, and the order-of-magnitude values below are insensitive to the choice) — $r_H = c/H_0 \sim 1.4 \times 10^{26} \text{ m}$, $T_{\text{GH}} \sim 3 \times 10^{-30} \text{ K}$, and $S_{\text{dS}} \sim 2 \times 10^{122}$ (the famous cosmological entropy bound; all three evaluated in `p2c2_verify`). The temperature is thirty orders of magnitude below the CMB — unobservable directly, which is why the de Sitter rows of the falsifier apparatus run through structural consistency rather than direct detection.

Physical picture. The black-hole observer stands outside a region she cannot see into; the de Sitter observer stands *inside* a region she cannot see out of. Both horizons hide partners of local vacuum modes, both therefore carry a temperature set by their surface gravity, and both charge entropy by the same quarter-area rule — one theorem, two signs of the enclosure.

Q.3 Step 3 — the GDOS reading

In GDOS, the de Sitter horizon is a cosmological OGGEF port, exactly analogous to the black-hole horizon but with the observer *inside* a causally disconnected region rather than outside. The information content observable to an inertial observer is bounded by S_{dS} from above; this is Banks’s idea of “holographic cosmology.” The same-arena clause of the analytic-shore license (§31.2) covers this case explicitly: the continuation that builds the Euclidean four-sphere is licensed by the identical one-sided spectrum argument, so the de Sitter entropy is the same theorem as the black-hole one, evaluated on a different real section.

The relation to the Hubble paper (Arisaka, 2026, C2-Hubble) is direct: there is one Friedmann background and one present rate, every pipeline is audited for whether it reads that background or a calibrated rung, and the calibrated ladder alone is displaced. The de Sitter horizon is the port through which the background is read at all.

Problems and worked solutions. Q.1.1. Verify that the static patch solves the Einstein equations with $\Lambda = 3H^2$.

Solution. Step 1: the metric is (74)-form with $e^{2\alpha} = e^{-2\beta} = f = 1 - H^2 r^2$; the vacuum-with- Λ equations read $G_{\mu\nu} + \Lambda g_{\mu\nu} = 0$, equivalently $R_{\mu\nu} = \Lambda g_{\mu\nu}$ in four dimensions. Step 2: for this metric family a two-line computation (the $\theta\theta$ equation of Appendix E, Step 1, with the new f) gives $(rf)' = (r - H^2 r^3)' = 1 - 3H^2 r^2$, so the $\theta\theta$ equation $1 - (rf)' = \Lambda r^2$ holds iff $\Lambda = 3H^2$. Step 3: the remaining components are then satisfied identically ($R_{\mu\nu} = 3H^2 g_{\mu\nu}$: de Sitter is maximally symmetric). Physics intuition: $\Lambda = 3H^2$ is the dictionary between the cosmologist’s expansion rate and the relativist’s constant — one number wearing two habits, and the horizon at $r_H = 1/H$ is where the two descriptions meet.

Q.2.1. Show that the de Sitter thermal wavelength is exactly the horizon scale.

Solution. Step 1: from (174), $k_B T_{\text{dS}} = \hbar H / 2\pi$. Step 2: the corresponding thermal wavelength is $\lambda = \hbar c / (k_B T_{\text{dS}}) = 2\pi c / H$. Step 3: the horizon radius is $r_H = c/H$, so $\lambda = 2\pi r_H$ — the circumference of the horizon, exactly. A “detector” smaller than the horizon can never resolve even one full oscillation of the typical quantum. Physics intuition: this is why the de Sitter temperature is unobservable in any local sense and why the text’s falsifier rows for the de Sitter sector run through structural consistency rather than direct detection: the radiation is not faint so much as *larger than any laboratory*, including the causal patch itself.

Q.2.2. Derive the famous 10^{122} and give it its holographic reading.

Solution. Step 1: from (174), $S_{\text{dS}}/k_B = \pi c^5 / (G_N \hbar H^2)$; with the present-epoch $H \approx 2.2 \times 10^{-18} \text{ s}^{-1}$, $S_{\text{dS}} \approx 2 \times 10^{122} k_B$. Step 2: rewrite using $r_H = c/H$ and $\ell_P^2 = G_N \hbar / c^3$: $S_{\text{dS}}/k_B = \pi r_H^2 / \ell_P^2$ — one quarter of the horizon area in Planck pixels, the same quarter-area law as the black hole (C-BH-2). Step 3: the number is the largest entropy in the observable universe by some seventeen orders of magnitude (the supermassive-hole inventory totals $\sim 10^{105} k_B$): the cosmological ledger is dominated not by anything *in* the universe but by the horizon that bounds the view. Physics intuition: 10^{122} is the information capacity of our port — in the paper’s reading, the total bandwidth of the projection every 4D observer lives behind; cosmology’s famous coincidences (the cohort’s Λ -sector results,

GG-A10-Cosmos) are statements about this one number.

Take-away. *Standard result:* $T_{\text{dS}} = \hbar c H / (2\pi k_B)$ and $S_{\text{dS}} = \pi k_B c^5 / (G_N \hbar H^2) \sim 10^{122} k_B$ today. *What it assumed (in 1977):* the same two formal moves as the black-hole case. *The GG-Theory reading:* the same theorem on a different real section, with the observer inside — the module’s bridge to the cosmological boundary stratum (GG-A10-Cosmos, GG-C2-Hubble).

R The induced-gravity species sum

What this appendix teaches. The program’s one declared open computation, set up honestly: the heat-kernel toolkit and where the sum rule comes from, the closure condition as a sum rule on the field content, why the Kaluza–Klein tower makes it self-consistent, the role of GU-5, the Standard-Model sector — and the honest verdict on what is and is not established.

Read alongside: [Parker and Toms \(2009\)](#) for the heat-kernel technology; [Sakharov \(1968\)](#) and [Jacobson \(1994\)](#) for induced gravity; [Susskind and Uglum \(1994\)](#) for the species problem this appendix closes in on.

Section 31.5 established $S_{\text{BH}} = A/4G_N$ with the exact coefficient, for *whatever* value the induced Newton constant takes. This appendix asks the remaining quantitative question: does the induced G_N , summed over the actual GG-Theory spectrum, equal the G_N that the scale reduction (§31.3) and GU-5 assign? We set up the computation, identify what makes it close, and state plainly what is and is not established.

R.1 Step 1 — the heat-kernel toolkit: where the sum rule comes from

The sum rule that this appendix runs on is a one-loop statement, and its derivation is short enough to display in full. For a field with quadratic fluctuation operator $\Delta = -\square + m^2 + \xi R$ on a Euclidean background, the one-loop effective action is

$$W = \frac{1}{2} \ln \det \Delta = -\frac{1}{2} \int_0^\infty \frac{ds}{s} \text{Tr} e^{-s\Delta}, \quad (175)$$

the proper-time representation (s has dimension length²). The trace has the Seeley–DeWitt small- s expansion

$$\text{Tr} e^{-s\Delta} = \frac{e^{-sm^2}}{(4\pi s)^2} \int d^4x \sqrt{g} \left[1 + s \left(\frac{1}{6} - \xi \right) R + O(s^2) \right], \quad (176)$$

whose first two coefficients are the only ones this appendix needs: $a_0 = 1$ (the flat-space volume term, feeding the induced cosmological constant) and $a_1 = (\frac{1}{6} - \xi)R$ (the curvature term, feeding the induced Einstein–Hilbert action) ([Visser, 2002](#); [Wald, 1984](#)). Insert (176) into (175) and read off the coefficient of $\int \sqrt{g} R$:

$$\frac{1}{16\pi G_{\text{ind}}} = \frac{1}{2} \frac{\frac{1}{6} - \xi}{16\pi^2} \int_{1/\Lambda^2}^\infty \frac{ds}{s^2} e^{-sm^2} = \frac{(\frac{1}{6} - \xi)\Lambda^2}{32\pi^2} + O(m^2 \ln \Lambda), \quad (177)$$

the quadratically divergent leading term regulated by a proper-time cutoff at $s = 1/\Lambda^2$. Three remarks make this graduate-safe. (i) The $O(1)$ normalization of the divergence is regulator con-

vention — proper-time, dimensional subtraction at a matching scale, and Pauli–Villars distribute the factor between Λ^2 and the coefficient differently; the physical statement is the *structure* $1/G_{\text{ind}} \propto (\text{weighted count}) \times \Lambda^2$, and the convention is absorbed into the definition of the physical cutoff (this is half of the “cutoff identification” entry in the honest-gap list below, and Visser (2002) catalogs the freedom). Arithmetically: the sum rule of Step 2 quotes the χ table in the convention $\Lambda_{\text{Step 2}}^2 \equiv \frac{1}{2} \Lambda_{\text{proper-time}}^2$, which absorbs the explicit factor $\frac{1}{2}$ of (177); the two displays agree identically under that identification. (ii) A conformally coupled scalar, $\xi = \frac{1}{6}$, contributes zero at this order — the first entry of the χ table in the sum rule, now derived rather than quoted. (iii) Higher spins repeat the computation with the appropriate fluctuation operators; their a_1 coefficients carry the spin-dependent signs that the honest-gap list flags (Callan and Wilczek, 1994; Demers et al., 1995; Kabat, 1995). With the toolkit displayed, the sum rule of Step 2 is a definition-chase: sum (177) over every propagating degree of freedom and name each field’s coefficient χ_i .

Physical picture. Sakharov’s point (Sakharov, 1968) survives every refinement: gravity’s stiffness $1/G$ is (at least in part) vacuum polarization — the elasticity a medium acquires because bending it costs its resident fields zero-point energy. The heat kernel is the bookkeeping device: each field, each mass, each spin deposits its a_1 coefficient, and the ledger’s total is the Newton constant the low-energy observer measures.

R.2 Step 2 — the closure condition as a sum rule

The one-loop induced Einstein–Hilbert coefficient is, from the a_1 Seeley–DeWitt coefficient of each field’s fluctuation operator,

$$\frac{1}{16\pi G_{\text{ind}}} = \frac{\Lambda^2}{16\pi^2} \sum_i \chi_i, \quad \chi_{\text{scalar}} = \frac{1}{6} - \xi, \quad \chi_{\text{conf. scalar}} = 0, \quad \chi_{1/2}, \chi_1 \text{ (spin-dependent)}, \quad (178)$$

with Λ the physical ultraviolet cutoff and the sum over every propagating degree of freedom. Closure is the statement

$$\boxed{\frac{\Lambda^2}{16\pi^2} \sum_i \chi_i = \frac{1}{16\pi G_N}} \quad \Longleftrightarrow \quad \sum_i \chi_i = \pi \left(\frac{M_{\text{Pl}}}{\Lambda} \right)^2, \quad (179)$$

using $G_N = 1/M_{\text{Pl}}^2$. This is a *sum rule* relating the weighted species count $\sum_i \chi_i$ to the Planck-to-cutoff ratio — the modern form of Sakharov’s self-consistency condition (Adler, 1982; Sakharov, 1968; Visser, 2002).

R.3 Step 3 — why the Kaluza–Klein tower makes it self-consistent

The sum in (178) is not over a finite list: every bulk field contributes an entire Kaluza–Klein tower of four-dimensional modes along $I_r \times S_\theta^1$. The decisive structural fact is that *the tower sum reconstructs the dimensional reduction* — summing the four-dimensional induced-gravity contributions over the KK modes of a bulk field is identical to reducing the six-dimensional

induced-gravity term over the scale profile:

$$\sum_{n \in \text{KK}} \left[\frac{1}{16\pi G_N} \right]_n^{(4\text{D})} = \left[\frac{1}{16\pi G_6} \right]^{(6\text{D})} \int_0^{2\pi} d\theta \int_0^K dr e^{-2r} = \frac{1}{16\pi G_N} \Big|_{\text{geom}}, \quad (180)$$

the right-hand side being exactly Step 1, Eq. (26). The species sum and the geometric reduction are therefore *the same object computed two ways*; they cannot disagree without the six-dimensional theory being inconsistent. In this sense the scale profile does not merely regulate the species sum — it *performs* it.

Lemma 2.3 (The tower identity (L-BH-Tower)). *For each bulk field, the Kaluza–Klein tower sum of its four-dimensional induced-gravity contributions equals the six-dimensional reduction of its induced term over the canonical scale profile, Eq. (180): the species sum and the geometric reduction are one object computed two ways. Grade: Derived (an identity of the reduction); the closure argument that consumes it is graded separately in the honest verdict below, and its failure mode is the registered internal kill row F-BH-9 (Table 9).*

R.4 Step 4 — the role of GU-5

Whether (179) is a *derivation* of G_N or a *consistency check* on it depends on the status of G_N . In GG-Theory G_N is a member of GU-5, the geometric conversion factor of the scale integration and the cohort’s only dimensionful anchor. It is therefore an *input*, and (180) guarantees that the entanglement entropy of the spectrum reproduces $A/4G_N$ *by construction*, provided the bulk theory is the consistent ultraviolet completion that the constitution’s admissibility filter (the AX-4 lineage) requires. Closure is then not an extra postulate but a self-consistency already enforced, at the level of the leading quadratically divergent term, by the tower identity (180).

R.5 Step 5 — the Standard-Model sector, and the honest gap

What (180) does *not* settle is the precise numerical coefficient $\sum_i \chi_i$ at the leading (zero-mode, Standard-Model) level, for two reasons we state plainly.

1. **Spin-dependent signs.** The coefficients $\chi_{1/2}$ and χ_1 for fermions and gauge fields enter with signs and curvature couplings that are notoriously delicate; for some spectra the naive induced G_N even emerges with the wrong sign (Visser, 2002). A reliable $\sum_i \chi_i$ for the GG-A9-SM field content requires the full sign-resolved bookkeeping, which we have not carried out.
2. **Cutoff identification.** The physical Λ is the bulk scale at which the KK description takes over; pinning $(M_{\text{Pl}}/\Lambda)^2$ to the scale depth $K = 35$ of OGN-5 is plausible — the scale profile generates precisely such a hierarchy — but we do not derive the exact relation.

An order-of-magnitude orientation: for Λ at the bulk scale the rule (179) requires $\sum_i \chi_i \sim \pi(M_{\text{Pl}}/\Lambda)^2$, which the KK tower supplies through (180); the zero-mode Standard-Model piece is an $O(10^{-10^2})$ contribution on top. These figures are consistent with closure but do not constitute a proof.

R.6 Step 6 — honest verdict

The species sum *closes structurally*: by (180) the weighted tower sum is identical to the scale reduction that defines G_N , so the entanglement entropy reproduces $A/4G_N$ given the GU-5 anchor and a consistent bulk completion. What is *not* established here is the precise leading numerical coefficient $\sum_i \chi_i$ for the GG-Theory spectrum, which is sign-sensitive and awaits a dedicated computation with the full bulk field content. We record this as the principal open calculation of the entropy program. None of it affects the coefficient $1/4$, which §31.5 fixed independently of the species content.

Problems and worked solutions. R.1.1. Show that a conformally coupled scalar contributes zero to the induced Newton constant at leading order, and say why.

Solution. Step 1: from (178), the scalar’s coefficient is $\chi_{\text{scalar}} = \frac{1}{6} - \xi$; the conformal value is $\xi = \frac{1}{6}$. Step 2: hence $\chi = 0$: the field’s heat kernel has no $a_1 R$ term, and its vacuum polarization stiffens spacetime not at all at this order. Step 3 (why): conformal coupling makes the massless scalar locally scale-invariant, and a scale-invariant field cannot contribute to a dimensionful stiffness $1/G_{\text{ind}} \propto \Lambda^2$ at leading order — there is no scale in its dynamics for the cutoff to dress. Physics intuition: induced gravity is bought with broken scale invariance; each field pays in proportion to how badly its fluctuations know about scale, which is why the χ table is sign- and coupling-sensitive and why the honest verdict of Step 6 calls the full sum a real computation rather than a formality.

R.2.1. Evaluate the closure requirement at the ledger’s scale depth.

Solution. Step 1: the sum rule (179) demands $\sum_i \chi_i = \pi(M_{\text{Pl}}/\Lambda)^2$ for physical cutoff Λ . Step 2: pinning Λ at the ledger’s depth- K scale gives $M_{\text{Pl}}/\Lambda \sim e^K$; with $K = 35$, $(M_{\text{Pl}}/\Lambda)^2 \sim e^{70} \approx 2.5 \times 10^{30}$, so the requirement is $\sum_i \chi_i \sim \pi e^{70} \approx 8 \times 10^{30}$. Step 3: a Standard-Model-sized field content ($O(10^2)$ zero-mode degrees of freedom) is thirty orders short — and that is precisely the gap the Kaluza–Klein tower fills via the identity (180): the tower sum over $\sim e^{70}$ effective modes below the cutoff *is* the dimensional reduction of the 6D term. Physics intuition: the Planck scale’s enormity relative to particle physics is, in this reading, a counting statement — gravity is weak because $1/G_N$ is the pooled contribution of an exponentially deep tower, and the depth is the ledger’s K , not a new dial.

R.3.1. Verify the tower identity in its simplest instance: one bulk scalar on the flat interval.

Solution. Step 1: take the 6D scale profile flat ($A = 0$) over $r \in [0, K]$ with unit compact circle; the KK modes of a bulk scalar are labeled n , each a 4D field contributing one equal quantum of induced stiffness per mode below the cutoff (the $\Lambda^2/16\pi^2$ of (178)). Step 2: the number of modes below Λ is the phase-space volume of the two compact directions, $N \propto \int_0^{2\pi} d\theta \int_0^K dr = 2\pi K$ per unit cutoff cell: the mode *sum* is the compact *integral*. Step 3: the right side of (180) is the 6D induced term reduced over the same integral — identical object, computed once in 4D language (sum over modes) and once in 6D language (integral over profile): for the flat profile the identity is arithmetic; the canonical profile e^{-2r} weights the integral without changing the logic (Lemma 2.3). Physics intuition: “the tower performs the reduction” is not a metaphor — summing a Kaluza–Klein tower *is* integrating over the extra dimensions, mode by mode, which is why the species sum and the geometric $1/G_N$ cannot disagree without the 6D theory itself being inconsistent.

Take-away. *Standard technology:* heat kernels and the induced Newton constant. *The closure condition:* the species sum over the actual GG-Theory spectrum must reproduce the G_N that the scale reduction assigns. *Status:* the sum rule and its self-consistency mechanism are established; the leading coefficient over the full bulk field content is the program’s declared open computation — and the coefficient $1/4$ does not wait on it.

S Internal source ledger

Every load-bearing element of the module, with its source status: *derived here* (this paper’s own result), *consumed* (cohort result used under its key, never re-derived), or *imported* (literature result used with attribution). The ledger is the anti-arbitrage instrument: nothing below is claimed at two grades in two places.

Table 21: The internal source ledger. “Consumed” items are cited under permanent cohort keys and carry their home grades; “imported” items are the module’s declared literature inputs. The status column is the one that matters. Derived here means the argument is in this paper and can be attacked here; consumed means the result belongs to a companion at that companion’s grade and this paper only uses it; imported means it is somebody else’s published result. Mixing the three is how a module comes to look better supported than it is, so they are separated by column rather than by prose.

Element	Status	Source / key
Exact $1/4$ coefficient (T-BH-1)	derived here	Steps 1–3 (§31.5); premises GGEP, canonical metric, OGGEF trace
Induced $1/G_N$ from the profile	derived here	Step 1; same object as GG-A5-GDOS’s boundary reduction
Thermal clock face (T-BH-5)	derived here	composition of T-P5-B1 (consumed) with Kay and Wald, 1991 (imported)
Euclidean arena license (L-BH-Shore)	derived here	shore from T-P5-B1 (consumed); replica saddle (imported, Solodukhin, 2011)
Freshness expiry at t_{Page} (L-BH-Fresh)	derived here	Page-average arithmetic + typicality premise (imported, flagged)
GSL direction and price (L-BH-Arrow)	derived here	DPI (identity) + consumed arrow + induced first law
Rotating anchor / collapse channel (T-BH-3/4)	derived here	model classes; Kerr Hadamard question imported (Kay and Wald, 1991)
Constitutional arrow	consumed	AX-3 (GG-A1-Constitution); law-level in Arisaka, 2026 , P5-CPT Part IX; dynamical in GG-A5-GDOS
Port-Freshness Clause	consumed	L-GDOS-Fresh, GG-A5-GDOS Definition 5 (its grade lives there)
Wick quartet	consumed	T-P5-B1 (Arisaka, 2026 , P5-CPT)
Real-category grading	consumed	L-Constitution-1, candidate grade (GG-A1-Constitution)
Born-weight grounding	consumed	GG-A8-QM: closed at the relaxation level; the field-grade representation step is its declared residue
Canonical bulk metric; arena	consumed	GG-A2-GG6 Theorems 0 and 2
Observer-vacuum reading	consumed	Arisaka, 2026 , A10-Cosmos (its observer-vacuum section: R14–R15 closed at model level in GG-A8-QM/GG-P4-Real); the curved-space existence residue it imports adopted open here
Replica saddle $\ln Z_n = -I_E$	imported	Solodukhin, 2011 (declared at §31.5)
Bifurcate-horizon uniqueness/thermality	imported	Kay and Wald, 1991
Page averages; scrambling	imported	Page, 1993 ; Hayden and Preskill, 2007
Wall-class GSL machinery	imported (ceiling)	Wall, 2012 ; registered open

T Cross-paper inheritance table

Table 22: The compact inheritance table: one row per companion, consumption only — this module re-derives none of it. One row per companion, and the second column is a list of takings rather than of agreements. This module re-derives nothing it inherits, which means an error anywhere in the first column propagates here unchecked. That is the cost of working in a cohort, and this table is where it is made visible; the last column says which step would fail.

Companion	What this paper takes	Where it is used
GG-A1-Constitution	AX-1–AX-4; L-Constitution-1 (candidate); the Reality Register	premises of T-BH-1; the funnel; O-BH-4
GG-A2-GG6	Theorem 0 (forced arena); Theorem 2 (canonical metric, no free parameters)	Part II; Step 1
GG-A5-GDOS	T2–T4 open-system law; L-GDOS-Fresh; boundary-gravity reduction	the port; Lemma 2.1; Step 1 consistency
GG-A6-GRIP, GG-A7-Qi4	the Gi-4/Qi-4 pillars forcing GDOS; the ledger structure theorems	the funnel; the four laws; Lemma 2.2
GG-A8-QM	hidden-fiber language; real-bulk series home	Step 4; §29
GG-A9-SM	the boundary field content	the species sum (F-BH-9); $c_{\log}$
GG-A10-Cosmos	observer-vacuum reading; de Sitter GGA; PBH subdominance	§27; way-station close; Part VII
GG-A11-Origin	GRIP regimes; No-Exit theorem	the way-station Part VI sections
GG-P1-KK	M_{lock} and ℓ_{lock}	every Kerr-null suppression
GG-P2-N3	$N_g = 3$, not a ledger entry	field content of the sum
GG-P4-Real	the derived clock (origin of i); real-bulk treatment	T-BH-5(i); §29
GG-P5-CPT	T-P5-B1; the four endowments; the law-level arrow	T-BH-5; the licensed arena; Lemma 2.2
GG-C1-Kaxion, GG-C2-Hubble	Kaxion DM saturation; the one background rate	Part VII; de Sitter evaluation

U Consolidated benchmark audit

Every number quoted in a proof or verification clause of this paper, with the script that reproduces it (all scripts in the program’s verification suite, www.arisaka-gg.org).

Table 23: The consolidated benchmark audit. A quoted number with no row here is a defect; report it. The middle column prints the tolerance each claim was verified to, not the claim itself, because a verified statement with no number beside it is an assertion. Where a tolerance is loose — the collapse-channel quadrature at the 10^{-3} class — that is the honest limit of the numerics and not a rounded result. The last column names the script, so every row here can be re-run.

Result	Verified statement (value)	Script
T-BH-3	exact symplecticity (10^{-9}); $\sinh^2(gt)$ law (4×10^{-14}); form indefiniteness; polar uniqueness	rotating-anchor suite
T-BH-4	branch-cut ratio $e^{-2\pi\nu}$ by quadrature ($< 10^{-3}$ rel. at $\nu = \frac{1}{2}, 1, 2$); late-time factorization ($0.13 \rightarrow < 10^{-2}$)	collapse-channel suite
T-BH-5	polar split exact (0); KMS ratio and Planck occupation ($\leq 3 \times 10^{-16}$); one-sidedness obstruction displayed; Γ -identity ($\leq 2 \times 10^{-16}$)	s1_kms
L-BH-Fresh	$I(q:R) = 0.002 \ln 2$ deep-fresh; growth ratios 4.05, 4.13, 4.12, 4.03 (predicted 4); Page-point crossing; saturation $1.997 \ln 2$ ($N = 12$, Haar, 5 realizations)	s3_fresh
L-BH-Arrow	DPI contraction in 300/300 random channels; exchange identity to 3.6×10^{-15} ; positivity over a 1.2×10^5 -point grid (min $+1.2 \times 10^{-8}$, at $\bar{n}_b \rightarrow \bar{n}_H$); closed forms to 10^{-12}	s4_arrow
Block II worked chapters	$T_H(M_\odot) = 6.17 \times 10^{-8}$ K; $t_{\text{evap}}(M_\odot) \approx 2 \times 10^{67}$ yr; $t_{\text{Page}}/t_{\text{evap}} = 1 - 2^{-3/2}$; Page average vs. asymptotics; $T_{\text{min}} = \sqrt{3}/2\pi L$ and F -flip at $r_+ = L$; thermal relative entropy (162) (positivity + face zero); Kerr first-law partials (168) (10^{-12}); $\kappa_{\text{dS}} = H$; $T_{\text{GH}}, S_{\text{dS}}$ values	p2c2_verify

V Named-statement deck

The deck (Table 24) is the paper’s claims registry, printed in full: every named statement, its one-line content, its exact grade, its home, and its falsifier exposure. Nothing in the paper is claimed above this table’s register, and the standing challenge of the objections Part — check the deck, not the prose — points here.

W Promotion roadmap

What would raise each conditional or open item, stated so the promotions are auditable in advance. (i) *The species sum* (F-BH-9): a sign-resolved $\sum_i \chi_i$ over the GG-A9-SM content plus the tower — closes the module’s principal open computation and converts O-BH-5’s concession from open to settled; chartered as its own project. (ii) *The typicality premise* (Lemma 2.1): replace the imported ETH/Haar assumption by bulk dynamics — promotes L-BH-Fresh to Derived. (iii) *The exchange class* (Lemma 2.2): extend the identity toward interacting fields and arbitrary cuts

Table 24: The named-statement deck. Grades are exact; nothing in this paper is claimed above this table’s register. The Exposure column keys each claim to the falsifier rows of Table 9 (F-BH-1..10) that would retire it; the internal row F-BH-9 exposes only the closure reading, exactly the partition O-BH-52 audits.

Key	One-line statement	Grade	Home	Exposure
T-BH-1	horizon entropy is $A/4G_N$ with $1/4 = 4\pi/16\pi$, regulator-independently, for the generic black hole	Derived-conditional (three premises)	Thm 1	F-BH-7; 1–3
T-BH-2	the Page turnover is the GFF branch switch	named mechanism; see L-BH-Fresh	Thm 2	F-BH-6
T-BH-3	rotating anchor: selection fails in the ergoregion; anchor at infinity; superradiance is symplectic pumping	model level, verified	Thm 3	F-BH-2
T-BH-4	collapse channel: Planck factor = branch-cut asymmetry; universal late-time channel	model level, verified	Thm 4	F-BH-5
T-BH-5	the thermal clock face is proven: continuation and imaginary time by T-P5-B1; KMS at $2\pi/\kappa$	Derived (quasi-free)	Thm 5	F-BH-5
L-BH-Fresh	freshness expires at t_{Page} by counting; the information return is an assumption switching off	Derived-conditional	Lem 2.1	F-BH-6
L-BH-Arrow	GSL direction from three named links; $\Delta S_{\text{gen}} =$ spent relative entropy in the exchange class	Derived (class) / Derived-conditional; Wall-class open	Lem 2.2	F-BH-8
L-BH-Tower	the KK tower sum is the dimensional reduction, computed twice	Derived (identity)	Lem 2.3	F-BH-9
L-BH-Port	at a horizon the OGGEF trace is forced: the port is ontological	Derived (AX-1+AX-3)	§25.2	F-BH-8
L-BH-Shore	the Euclidean arena is the clock’s analytic shore: the continuation is a theorem (T-P5-B1); the saddle is a declared import	Derived-conditional (quasi-free; saddle imported)	§31.5	—
C-BH-1	the conditional-GGA crossover is the GSL’s equality point	Derived (first law + exchange identity; conditional on the quasi-static bath reading)	Cor 2.1	F-BH-5
C-BH-2	the de Sitter horizon entropy is the same theorem on the cosmological OGGEF port	Derived-conditional (T-BH-1’s premises; same-arena clause)	§52.2	F-BH-7
RB-13	horizon entropy data are real-section geometry	registered in GG-A8-QM (registrational; no new theorem)	§29	—
(prop.)	the islands formula is the semiclassical display of the OGGEF output; the QES is the traced-region boundary	dictionary-grade proposition	App. M	—

(Wall, 2012 machinery) — promotes the GSL statement from class-exact to Wall-class; registered open. (iv) *The quasi-free arena* (§31.5): an interacting-grade shore — promotes the arena license (L-BH-Shore) from Derived-conditional toward Derived. (v) *The real-bulk proposal* (§29): registration by the series’ home papers — converts the proposal to a numbered entry. (vi) *The curved-space existence residue* (imported at Arisaka, 2026, A10-Cosmos’s observer-vacuum section; the model-level closures R14–R15 are already in hand as T-BH-3/4): an existence theory at the interacting grade — discharges the module’s adopted open problem. (vii) *The PBH ladder* (F-BH-4): a derived exponent and an absolute anchor — promotes the conjecture to a prediction; needs cosmology-side input, chartered separately.

X Falsification-test catalog

The module’s kill apparatus lives in three displays, one address each: the claims-level kill table (F-BH-1.1.10, Table 9), the dated instrument timeline with pre-registered verdicts (Table 8), and the refresh discipline that keeps both honest (§59). The external hostages — the lock cluster, the Kaxion window, the neutrino ordering — are cohort rows, inherited by construction (§61.7); this module adds no dial to any of them. One-line rule of use: a row fires as written or not at all, and a fired row retires the module.

Y An annotated reading guide

Shelf 1 (the thermodynamic foundation): Bekenstein 1973; Bardeen–Carter–Hawking 1973; Hawking 1975; Page 1993 — the four papers this module’s Parts I and VI answer to. *Shelf 2 (the induced-gravity lineage)*: Sakharov 1968; Jacobson 1994; Frolov–Fursaev 1997; Susskind–Uglum 1994; Callan–Wilczek 1994; Visser 2002 — read in that order, the lineage whose missing premise (*no bare G_N*) Part V supplies. *Shelf 3 (the information decade)*: AMPS 2013; Hayden–Preskill 2007; Penington 2020; Almheiri et al. 2019; the replica-wormhole paper — the modern paradox literature Part VI converses with. *Shelf 4 (the channel mathematics)*: Spohn 1978; Casini 2008; Wall 2012 — the entropy-production spine of Lemma 2.2 and its declared ceiling. *Shelf 5 (the instruments)*: the LVK catalog papers and GW250114 companions; the EHT series — the record Part VIII scores. *Shelf 6 (the cohort)*: GG-A2-GG6 for the arena, GG-A5-GDOS for the law, GG-P4-Real/GG-P5-CPT for the clock and its mirrors, GG-A8-QM for the fiber, read in that order; then this paper re-read from Part IV, which is where it stops being a survey and starts being a derivation.

Part XIII

Common Objections and Responses

Thirty Questions from the General Reader, in Plain English, followed by Fifty-three Adversarial Readings — from the Classical Relativist, the Curved-Spacetime Field Theorist, the String Theorist, the Islands Theorist, the Canonical-Quantization Referees, the Emergent-Gravity Theorist, the Gravitational-Wave Observer, the Horizon-Imaging Astronomer, the Quantum-Information Theorist, the Quantum-Foundations Referee, the Historian and Philosopher, the Student, and the Cohort-Internal Referee — Each Answered with Pointers into the Paper

A claimed first-principles derivation of the Bekenstein–Hawking coefficient — for the generic black hole, outside string theory — will and should draw hard objections. This Part collects the fifty-three sharpest adversarial readings we could construct, grouped by the referee most likely to raise each, and answers each with explicit pointers to the theorem, section, register entry, or named open item that carries the load. Where the honest answer is “not yet,” “not ours to claim,” or “conceded,” it is said in those words. These are self-authored stress tests in referee personas, published so that real referees can start where they end. Seven of the objections carry the module’s permanent seed keys (O-BH-1 through O-BH-7), filed under their communities; the keys are cited elsewhere in the paper (the source ledger, the promotion roadmap) and never change.

Before the referees, however, come the readers. The questions a non-specialist asks about black holes are not easier than the technical ones — they are usually deeper, because they go straight to what the mathematics is *about*: what a horizon is, where the information goes, whether the future is fixed, what is real. The opening section therefore collects thirty such questions, distilled from the themes of Parts IX–XI, and answers each one in plain English at whatever length honesty requires, with a pointer to the place in the paper where the full treatment lives. A reader with no physics beyond curiosity can read that section alone and come away with the paper’s picture of the black hole; the fifty-three technical readings that follow are unchanged, and keep their numbering and their scorecard.

65 Questions from the general reader

Thirty questions, in the order a curious reader tends to ask them: first what a black hole is and why it matters, then the famous paradox and where the information really goes, then what this theory claims and how it could be wrong, and finally what it all means. The answers are in plain English; each ends with a pointer to the full treatment. Nothing here is claimed beyond what the main text establishes, and where a statement is conditional or open, the answer says so in those words.

What a black hole is, and why it matters (Q1–Q8)

Question Q1 (the commuter, on relevance)

“Black holes are thousands of light-years away. Why should anyone but an astronomer care?”

Answer. Because a black hole is not an exotic exception to ordinary physics — it is ordinary physics with the volume turned all the way up. The same bookkeeping that makes your coffee cool, your laptop shed heat, and a photograph fade in sunlight is, in this paper’s reading, the mechanism behind Hawking radiation: a system loses track of information that flows somewhere it cannot see, and thermodynamics is the shape of that loss. The black hole is simply the one place where geometry itself enforces the loss, perfectly and calculably, which is why it is the sharpest test bench nature offers for the three great theories — gravity, quantum mechanics, and thermodynamics — at once. And there is a second, quieter reason: the universe has a horizon of its own, and every one of us lives inside it. The mathematics developed for black holes is the mathematics of our own cosmic situation. (The full case: Part IX’s comparison and Part XI’s closing sections.)

Question Q2 (the dinner-table questioner, wanting it straight)

“In one paragraph, without equations: what does this theory say a black hole actually is?”

Answer. The theory says the world has six dimensions, of which we perceive four; the other two are not far away but everywhere, too fine and too orthogonal for our instruments to enter. Everything any observer measures is a *projection* of the six-dimensional world onto their four-dimensional view — like a shadow cast on a wall. A black hole is the place where that projection becomes most severe: near the horizon, more and more of the world’s degrees of freedom slip into the unseen directions, until at the horizon the view simply ends. The horizon is not a wall in space; it is the edge of a window. What falls in is not destroyed — it moves into registers the outside view cannot carry. What comes out as faint radiation is the lawful, thermal signature of a view that has been truncated. That is the whole picture; the rest of the paper is the mathematics that makes it precise and the measurements that could kill it. (Part IV builds this; Part XI says it slowly.)

Question Q3 (the careful reader, on frozen time)

“I have read that time freezes at the horizon — that from outside, nothing ever falls in. So does anything ever actually cross?”

Answer. Both things are true, and the trick is that they are true *for different observers* — which, after Einstein, is no scandal. Watching your friend fall toward the horizon, you see their signals climb out of an ever-deeper gravitational well: each pulse takes longer than the last, their image reddens, freezes, and fades. From your side, they never quite cross. But your friend’s own clock ticks normally; they cross the horizon in perfectly finite time and notice nothing at the crossing — the horizon is not marked locally in any way. The two accounts never contradict each other, because no experiment can compare them face-to-face afterwards. In this paper’s language,

your frozen, fading image of your friend is exactly what a port delivers: the last information the window can carry, stretched over your forever. (The relational reading of horizon quantities: §28 and Part XI.)

Question Q4 (the nervous voyager, on survival)

“Would falling into a black hole hurt?”

Answer. It depends entirely on the size, and the answer surprises everyone. Near a small black hole — a few times the Sun’s mass — the gravitational *gradient* is ferocious: your feet would be pulled so much harder than your head that you would be stretched apart well before reaching the horizon. But the bigger the hole, the gentler its horizon: for the giant at the center of galaxy M87, the tidal stress at the horizon is some ten thousand times *weaker* than what your body feels standing on Earth. You would cross the point of no return without a bump, a flash, or any sensation at all — and in this paper’s reading there is no “firewall” of radiation at the crossing either; the apparent quantum argument for one dissolves when the bookkeeping is done in the full six dimensions. The genuinely unpleasant part comes much later, deep inside, where the projection becomes extreme. (The numbers: Appendix E; the no-firewall argument: Part VI.)

Question Q5 (the reader of popular science, on the center)

“Is there really a singularity at the center — a point where density is infinite and physics breaks down?”

Answer. “Infinite density” has never been a physical claim; it is mathematics raising a flag that says *my description has ended here*. Classical general relativity does predict, by theorem, that paths into a black hole terminate — but history offers a precedent for how such flags get resolved. For forty years the horizon itself was believed to be a place where physics broke down, until 1958, when it was shown to be an artifact of a poorly chosen map: change coordinates and the horizon is smooth. This theory proposes the same lesson one level up: the center is singular *in the four-dimensional projection* but smooth in the six-dimensional world — the flag marks the edge of the shadow, not an edge of reality. That is a proposal, not a theorem we can test at the center of a black hole; its support comes from the framework’s testable consequences elsewhere. (The 1958 story: Appendix E; the 6D reading: Part VI.)

Question Q6 (the skeptic at the whiteboard, on hidden dimensions)

“What are these two extra dimensions? Where are they hiding?”

Answer. They are not hiding somewhere else; they are here, at every point, in the same way the third dimension is present at every point of a printed map without appearing on it. One is a *scale* direction — a depth along which fields stretch and compress — and one is a tiny closed circle around which the forces of nature wind. We do not perceive them for the same reason a figure drawn on paper would not perceive height: our instruments, made of four-dimensional projections themselves, measure only the projection. What makes them physics rather than poetry is that their geometry is *rigid*: one fixed shape has to account simultaneously for the masses of particles, the strength of gravity, the composition of the universe, and — in this paper — the thermodynamics of black holes, with no adjustable numbers. A wrong shape fails loudly. (The bulk geometry: Part II; what rigidity buys: Q20 below.)

Question Q7 (the philosopher of science, on unfalsifiable metaphysics)

“You admit the six-dimensional bulk cannot be observed directly. Isn’t this exactly the kind of untestable metaphysics science is supposed to avoid?”

Answer. The concern is the right one, and the defense is the standard one in science, applied strictly. No one has seen a quark, the interior of the Sun, or a single atom of the early universe; we believe in them because positing them explains diverse observations with fewer assumptions than any rival, and because the posit sticks its neck out — it forbids things, and would be abandoned if the forbidden things were seen. The bulk is held to that same standard, and the paper’s structure is built around it: the framework’s numbers are fixed before the comparisons are made, the register says plainly which claims are derived and which remain open, and Part VIII maintains a public list of observations that would kill the black-hole reading outright — from horizon-scale images to a laboratory neutrino measurement. Metaphysics does not volunteer its own death certificate. This does. (The kill list: Part VIII; the register: the named-statement deck.)

Question Q8 (the traveler, on borders)

“Is the horizon a real, physical thing — a membrane, a surface, a wall?”

Answer. There is nothing there — no material, no marker, nothing a passing observer could detect locally. And yet it has exact, measurable consequences: a temperature, an entropy, a size photographed by telescopes. The resolution is that the horizon is not a thing but a *relation*: it is defined by what a given observer can and cannot ever reach, the way a national border is defined by law rather than paint on the ground — invisible underfoot, decisive in consequence. A hill can block your view and not your neighbor’s without anyone doubting the landscape behind it is real; a surface can be a horizon for one observer and not for another in precisely the same way. This paper’s contribution is to make the relation quantitative: the horizon is where the projection that defines an observer’s view runs out of bandwidth. (The relation reading: Part XI; the port theorem: Part IV.)

The paradox, and where the information goes (Q9–Q17)**Question Q9 (the newcomer, wanting the paradox itself)**

“What exactly is the black-hole information paradox, in plain words?”

Answer. Quantum mechanics has one inviolable rule: the universe never forgets. Given the complete present, the past is in principle recoverable; information can be scrambled beyond any practical hope but never destroyed. Black holes seemed to break the rule twice. First, a hole swallows anything — an encyclopedia, a star — and settles into a state characterized by just its mass, spin, and charge, with no visible trace of what went in. Second, Hawking showed in 1974 that the hole slowly shines with perfectly *thermal* radiation — and thermal radiation is, by definition, radiation that carries no message. Run the process to the end: the hole evaporates away entirely, the radiation remembers nothing, and the encyclopedia is not scrambled but *gone*. Either quantum mechanics fails, or gravity’s account of the horizon is incomplete. That dilemma, sharpened over fifty years, is the paradox. (The precise statement: Part III.)

Question Q10 (the same newcomer, wanting the resolution)**“And in plain words, how does this theory resolve it?”**

Answer. By locating the error in the question’s hidden assumption: that the four-dimensional description is the whole story. Every working physicist knows a system can *look* like it is losing information when it is merely sharing it with something you are not tracking — a hot cup losing heat to a room is not deleting anything. Hawking’s calculation produced exactly that signature, and the fifty-year alarm came from the apparent absence of any “room” for the information to have gone into: the radiation seemed to be all there was. The six-dimensional bulk supplies the missing room. The interior registers behind the horizon are as physically real as the air in the room around the cup; the information flows there, stays in the universe’s books, and returns — slowly, scrambled into subtle correlations of the radiation — as the hole evaporates. Nothing is deleted; a view was mistaken for the world. (The mechanism: Parts IV–VI; the slow return: Q16.)

Question Q11 (the pragmatist, on recovery)**“So if I burned my diary and threw the ashes into a black hole, is it really still ‘there’? Could anyone ever read it again?”**

Answer. In principle yes; in practice never — and the distinction is exactly the one that already applies to the burning. When a diary burns, its text passes into correlations among smoke particles, photons, and ash; no law of physics erases it, and a demon with perfect instruments and the patience of ages could reconstruct every word. No such demon can be built, but physics cares about the principle, because *in principle recoverable* is what separates scrambling from destruction — and destruction would bring down the whole quantum edifice. The black hole, in this reading, adds one twist: the correlations pass through registers no exterior instrument can reach directly, and are returned over the hole’s evaporation — a timescale so long that for a stellar hole it makes the age of the universe look instantaneous. The books stay balanced; the reading room has a very long waiting list. (The bookkeeping theorems: Appendix O.)

Question Q12 (the puzzled reader, on temperature)**“How can a black hole — a hole — be hot? Hot to whom?”**

Answer. “Hot to whom” is precisely the right question, and its answer is the paper’s favorite lesson: temperature at a horizon is *relational*. An observer who hovers far away, watching the hole, finds their detector clicking at a definite temperature — for a solar-mass hole, about sixty billionths of a degree above absolute zero, unmeasurably faint but exactly calculable. An observer falling freely through the same region detects nothing at all: to them it is empty space. Neither is mistaken, because temperature is not a property of the hole alone — it is a property of the *relationship* between the hole and a particular observer’s view, of what that view keeps and what it truncates. Physics has absorbed such a lesson before: in 1905, simultaneity turned out to be observer-dependent without becoming unreal. Objective, it turns out, does not mean observer-free; it means lawfully related across observers. (The thermal clock: §28; the philosophy: Part XI.)

Question Q13 (the accountant, on entropy)

“Physicists say a black hole has more entropy than anything else in the universe. Entropy of *what*, exactly?”

Answer. Entropy always counts the same thing: the number of microscopic arrangements consistent with what you can see. A shuffled deck has entropy because fifty-two cards can hide their order behind one face-down stack. The scandal of black-hole entropy was that the classical theory said a hole has *no* hidden arrangements — mass, spin, charge, done — and yet thermodynamics insisted it carries more entropy than the star that made it. This paper’s answer names the arrangements: they are configurations of the hidden compact dimension at the horizon — one register per pixel of Planck-sized area — the degrees of freedom the four-dimensional view cannot resolve. The entropy is literally the measure of how entangled your view is with what sits behind the port. And that explains the old mystery of why the entropy grows with the horizon’s *area* rather than its volume: the entropy lives on the boundary because the boundary is where the window ends. (The count: Part V; the honest open piece of it: Q21.)

Question Q14 (the empiricist, on evidence)

“Has Hawking radiation ever actually been observed?”

Answer. No, and it almost certainly never will be directly. A stellar-mass black hole radiates at billionths of a degree, while the universe around it glows at nearly three degrees: every astrophysical hole is, today, absorbing far more than it emits, and the signal is hopelessly buried. What exists is indirect but serious: laboratory “analogue horizons” — in flowing fluids and cold-atom systems — exhibit the kinematic effect Hawking predicted, in systems where the mathematics is the same though the gravity is absent. This paper is candid about the situation: Hawking radiation functions in physics as a consistency requirement — three independent derivations, worked in Appendices [H](#) and [I](#), give one answer — rather than as an observed phenomenon, and the framework’s own testable content therefore rides on the *geometry* side: images, ringdowns, and the cross-domain predictions of Part VIII. (The honest evidence table: Part VIII.)

Question Q15 (the reader who wants the end of the story)

“What happens at the very end, when a black hole finishes evaporating?”

Answer. Here is the register speaking plainly: nobody knows, and this paper does not claim to. The final flash — when the hole shrinks to near the fundamental scale — is the one moment where the calculation tools of every framework, including this one, run out of warranty. What the theory does fix is the *bookkeeping* across the whole history up to that point: the information is not stored up desperately awaiting a final reveal; it has been leaving all along, in correlations, since the halfway point of the evaporation (Q16), so the ending — whatever its detailed physics — is not asked to perform a miracle. No information bomb, no paradox deferred to the last second, no need for an eternal remnant. The final chapter is open; the accounts are already balanced when it begins. (The turnover theorem: Part VI; the open end-state: the promotion roadmap.)

Question Q16 (the follower of physics news, on the Page curve)

“I keep seeing the phrase ‘Page curve.’ What is it, and why did its recovery in 2019 cause such excitement?”

Answer. Imagine tracking not the radiation’s energy but its *mystery* — its entropy, the amount of hidden information it carries. If a black hole truly deletes what it swallows, the mystery grows forever: more radiation, more featureless heat. But Don Page observed in 1993 that if quantum mechanics holds, the curve must turn: past the halfway point of evaporation, each new quantum of radiation, being secretly correlated with quanta emitted earlier, starts *resolving* mystery rather than adding to it, and the curve comes back down to zero. Rise, turn, fall — unitarity’s signature, drawn as a graph. For twenty-five years gravity’s own mathematics refused to draw it; in 2019 it finally did, through the island constructions — a landmark, achieved in simplified models. In this paper the curve arrives a different way: the bulk was never anything but unitary, and the turnover is a theorem about when the projection’s bookkeeping switches branches. Same curve, roots one level deeper. (Worked slowly: Appendix K; the theorem: Theorem 2.)

Question Q17 (the curious reader, on islands)

“What is an ‘island,’ in plain words? The term sounds mystical.”

Answer. It is an accounting rule with an evocative name. After the Page turnover (Q16), the modern gravity calculation says something strange-sounding: to compute what the *radiation* knows, you must include a region *inside* the black hole — the island — as if it already belonged to the radiation. Mystical as that sounds, it is bookkeeping: the calculation is deciding which of two ways of paying for information is cheaper — keep renting (count every radiated quantum’s hidden partner separately) or buy the interior outright (count it once, at a fixed price set by the horizon’s area). At the turnover, buying beats renting, and the books flip. In this paper’s reading the flip is no mystery at all: the interior registers were *always* shared with the radiation through the bulk — the island rule is the four-dimensional shadow of that six-dimensional fact. (The mortgage-versus-rent ledger: Appendix M; the dictionary between languages: Appendix N.)

The theory’s claims, and how it could be wrong (Q18–Q22)**Question Q18 (the fair-minded reader, on string theory)**

“How does this relate to string theory’s account of black holes? Are you saying string theory is wrong?”

Answer. No — and the paper is deliberate about this. String theory achieved something this framework openly admires and has not matched: for special, supersymmetric black holes, it counted the microscopic states one by one and got the entropy exactly right, a landmark of mathematical physics. What it has not yet done is name the states of the ordinary, spinning, electrically neutral black holes that actually populate the sky. This framework makes the complementary claim: it addresses precisely the generic astrophysical case, in a universe with our particle content, with no adjustable numbers — and it lacks string theory’s depth of mathematical control in the special corners. The two are answers to differently shaped questions,

and both are staked to their own falsifiable predictions. The paper’s comparison chapter sets them side by side in those terms — structural contrast, not verdict. (The full comparison: Part IX; the string count, worked with respect: Appendix L.)

Question Q19 (the historian, on hubris)

“Some of the best physicists of two generations worked on this problem. Why would this framework succeed where they did not?”

Answer. The claim is not superior cleverness, and the paper should be read as making a structural point instead. Most programs attacked the paradox from within the four-dimensional description, adding machinery — new surfaces, new saddle points, new principles — to repair its books. This framework changed the starting point: it takes *openness* as fundamental — every observation is a projection of a larger system — and had that principle fixed, with its machinery, long before black holes entered the program; the same engine was already accounting for particle masses and cosmology. When it reached the horizon, the paradox-generating assumption (that the 4D account is complete) was simply never made, so there was no paradox to repair — and, crucially, no freedom left to tune: the framework arrived with its numbers already committed elsewhere. That is either a deep advantage or a spectacular way to be wrong in public, and the paper maintains the kill list that decides which. (The inheritance: Part II; the kill list: Part VIII.)

Question Q20 (the buyer, on zero free parameters)

“The paper keeps saying ‘no free parameters.’ What does that mean, and why should it impress me?”

Answer. A free parameter is a dial: a number the theory cannot predict and adjusts to fit the data. Dials are not sins — the standard model of particle physics has about twenty, the standard cosmology six, and both are triumphs — but every dial weakens a test, because a theory with enough dials can absorb almost any observation. This framework’s structural constants are integers fixed by internal consistency: turn any of them and the whole interlocked structure fails, not just one prediction. Why it matters here: when the black-hole module derives the famous entropy coefficient, or prices the deviations from Einstein’s geometry at a definite (unobservably small) size, there is no dial anywhere to make an awkward measurement comfortable. Predictions without dials can *fail* — which is exactly what makes them worth something. Whether they do fail is a matter for instruments, and Part VIII lists which ones. (The ledger’s rigidity: Part II; the priced deviations: Part VII.)

Question Q21 (the numerate skeptic, on retrofitting)

“Couldn’t all this be sophisticated numerology — patterns found after the fact and dressed up as predictions?”

Answer. This is the most important skeptical question in the book, and the paper’s defense is procedural rather than rhetorical, because after-the-fact reasoning cannot be refuted by assertion — only by exposure. Three exposures are offered. First, pre-registration: the claims are stated with named kill conditions *before* the deciding measurements, including some — like the ordering of neutrino masses, measurable in a laboratory — that have nothing to do with

black holes yet would retire this paper’s central theorem if they go the wrong way. Numerology does not take hostages in other fields. Second, an audited register: every claimed result carries a grade — derived, conditional, or open — and the paper concedes in print the pieces it has not closed, including one computation at the heart of the entropy story (the species sum) that is declared open rather than papered over. Third, reproducibility: the derivation chains are laid out step by step in the appendices precisely so that a hostile reader can check them. The paper’s own words: it asks to be checked, not believed. (The register: the named-statement deck; the open computation: Appendix R.)

Question Q22 (the referee in all of us, on falsifiability)

“Name it plainly: what observations would kill this picture?”

Answer. In plain words, any one of these. A black-hole photograph or a gravitational-wave ringdown that deviates from Einstein’s geometry at the few-percent level — this framework predicts agreement to precisions no foreseeable instrument can challenge, so any confirmed deviation retires it. Echoes in gravitational waves after a merger’s ringdown — not predicted here, where the horizon carries no structure to reflect them. A demonstrated loss of information in any gravitational process. And — the unusual ones — failures far from black holes: the wrong neutrino mass ordering, or an empty dark-matter window in the haloscopes — or a second line inside it, where the cohort predicts exactly one — would bring down theorems this module depends on, because the same six-dimensional geometry underwrites all of it. A framework that can be shot from a neighboring field is staking more, not less. Each kill condition is registered, with its current experimental status, in the tables of Part VIII. (The dictionary of kills: Part VIII and the falsification catalog.)

What it means (Q23–Q30)

Question Q23 (the determinist’s child, on the future)

“Does this theory say the future is already fixed?”

Answer. At the deepest level the theory describes, yes: the six-dimensional bulk evolves deterministically, and nothing in it is ever random or lost. But no observer lives at that level. Every observer’s view is a projection, and the theory proves that the projection *necessarily* loses information — so from inside any view, the future is irreducibly probabilistic, and no instrument, however perfect, can do better. This is a determinism that cannot be exploited even in principle: the registers that would complete the prediction sit behind causal boundaries, not behind practical ones. So the honest one-sentence answer is: fixed as seen from nowhere, open as seen from anywhere — and “anywhere” is the only seat there is. (The full discussion, including what kind of probability survives: Part XI.)

Question Q24 (the reader of biographies, on the old argument)

“So did Einstein win after all? Does God play dice?”

Answer. The theory’s answer is that the two great antagonists were describing different levels of one world, and each was right about his level. Einstein’s instinct — that beneath the statistics

there is a complete, lawful reality in which nothing is truly random — is honored in the bulk, where evolution is deterministic and information is conserved. Bohr’s insistence — that physics must speak only of what observation can deliver, and that at that level probability is irreducible — is honored in every four-dimensional view, where the projection theorem makes uncertainty lawful and permanent. The horizon is where their argument is staged most starkly, and what this framework offers is not a winner but a dissolution: probability is what determinism looks like through a port. (The staging: Part XI; the quantitative version of “through a port” is the business of the companion paper on quantum mechanics, and is registered there as conditional — said here so the slogan is not mistaken for a settled theorem.)

Question Q25 (the human being, on free will)

“Where does free will fit, if the bulk is deterministic?”

Answer. Honestly labeled: physics narrows this question less than either side of the old debate hopes, and this paper adds one structural observation rather than a resolution. The observation is that you — as a describer of yourself — are also a projection: no view available to you, or to any neighbor, contains the registers that would make your next choice predictable. The unpredictability of persons is therefore not a temporary gap awaiting better brain scanners; in this framework it is lawful and permanent, the same theorem that protects the quantum dice. Whether lawful unpredictability *is* freedom, or merely feels like it, is a question philosophy has debated under determinisms far older than this one, and the paper takes no side — it only relocates the mystery to sharper ground: the boundary between a world that is whole and views that never are. (The projection theorem: Part IV; the philosophical setting: Part XI.)

Question Q26 (the unsettled reader, on reality)

“If everything I see is a projection of something deeper, is my world an illusion?”

Answer. No — and the distinction matters more than any other in this list. An illusion misinforms; a projection informs truly but partially. Your shadow on a wall is not a lie about you — its shape tracks your shape lawfully; it merely does not carry all of you. Every measurement ever made in four dimensions is correct, exactly as correct as it always was; the theory changes not one reading of one instrument. What it denies is only the silent assumption that the four-dimensional account is *complete* — and incompleteness is not falsehood. Physics has taken this step before: the solid table was not an illusion when atoms arrived; heat was not an illusion when it became molecular motion. The world you see is real. The claim is that it is a true view of something larger — and that the seam between the view and the whole is not hidden but measurable: it is the horizon, and it has a temperature. (The long version: Part XI’s closing sections.)

Question Q27 (the reader of quantum mysteries, on observers)

“Does ‘observation’ in this theory require a conscious observer? Does a falling tree make a sound?”

Answer. No consciousness is required, anywhere in the framework. An “observer” here is any physical system whose causal situation truncates its access to the whole — a photographic plate, a rock warmed by the sun, a second black hole. The projection is defined by geometry

and causal structure, not by awareness: the tree’s fall is registered in the correlations of the air whether or not anyone is listening, and a horizon traces out its interior for a detector exactly as it would for an astronaut. What consciousness does have is a stake in the *consequences*: beings like us only ever occupy views, so the theory explains why the world as experienced — by anything — is probabilistic, thermodynamic, and directional even though the world as a whole is not. The measurement problem in full is the business of the cohort’s quantum-mechanics companion, and this paper leans on it only where stated. (The port definition: Part IV.)

Question Q28 (the cosmic claustrophobe, on our own horizon)

“Are we living inside a black hole?”

Answer. Not in the technical sense — our universe’s expansion history does not match a black-hole interior, and the question as usually posed confuses two different geometries. But it gestures at something the paper does establish: we live *inside a horizon*. An accelerating universe has a boundary beyond which light can never reach us, and that cosmological horizon obeys, by theorem here, the same thermodynamics as a black hole’s — a temperature, and the largest entropy in nature — derived by the same argument with the observer on the inside rather than the outside. The mathematics of this paper is, in that precise sense, autobiographical: every observer, including you, reads the universe through a port, and the black hole is simply the case where the port belongs to someone else. (The inside-out theorem: Appendix Q; the cosmological consumer: the cohort’s cosmology paper.)

Question Q29 (the reader of Wheeler, on ‘it from bit’)

“Wheeler asked whether the world is, at bottom, made of information. Does this theory answer him?”

Answer. It answers with a friendly amendment. In this framework the bottom layer is not information but *geometry with a law*: a six-dimensional state evolving deterministically. Information is not what the world is made of; it is the currency in which *views* of the world are denominated — the measure of what a projection carries and what it must leave behind. What makes Wheeler’s instinct profound rather than merely poetic is that this currency obeys exact laws, and at a horizon the exchange rate becomes visible: area in, entropy out, at a fixed price. The black hole is where “it” and “bit” meet across a counter and the books are settled in public — one statement written simultaneously in geometry, in information, and in heat. Which script is fundamental? On this paper’s reading: the geometry writes, the information is the readership. (The three scripts: Part XI’s Rosetta section; labeled as reading, not theorem.)

Question Q30 (every reader, at the door)

“If I remember one thing from this paper, what should it be?”

Answer. That the horizon destroys nothing. For fifty years the black hole was read as the place where physics loses — where information dies, where the laws contradict each other, where theory goes to fail. Read through this framework, it is the opposite: the place where physics is most honest about what any single view of the universe can carry, and where the accounting of the unseen is done exactly, in public, with a temperature and an area for ledger lines. What falls in is not erased; it is kept, and returned in a currency our instruments are only beginning

to learn to count. If that is right, then the black hole’s fifty-year crisis was never about black holes: it was the discovery, made the hard way, that our window is not the world — and that nature, seen whole, keeps perfect books. (Said at full length: Part XI, closing paragraph.)

66 The classical relativist

Objection O-BH-8 (the membrane-paradigm veteran, on old wine)

“The membrane paradigm already treats the horizon as a physical boundary with temperature, viscosity, and resistivity. Your ‘ontological port’ is that, renamed.”

Response. The membrane paradigm is an *effective* description for exterior dynamics — its membrane properties are bookkeeping for what falls through, valid because the interior cannot answer back. The port statement (L-BH-Port, §25.2) is an information-theoretic claim with theorem-grade consequences: the trace over the hidden factor is *forced* by the causal structure plus AX-1/AX-3, and thermality is that trace computed (Eq. (18)), not a boundary condition posited. The membrane’s transport coefficients are recoverable downstream; the direction of explanation is reversed.

Objection O-BH-9 (the global-structure purist, on teleology)

“The event horizon is teleological — defined by the entire future. No local measurement locates it, yet your port carries ontological weight. You have hung physics on an unobservable.”

Response. Fair, and the module’s operational statements are built to dodge it: the collapse theorem (T-BH-4, §27) runs on the *late-time universal channel*, not on horizon location; the GSL lemma is quasi-stationary; the freshness clause counts emitted quanta. Where the exact horizon matters — the stationary entropy statement — event and apparent horizons coincide. The genuinely teleological content is inherited from classical general relativity unchanged, and the honest statement is that GG-Theory neither worsens nor repairs it.

Objection O-BH-10 (the area-theorem custodian, on arrows by fiat)

“Wall proved the GSL in its class. Your lemma ‘consumes’ the arrow from an axiom (AX-3) rather than deriving it. Consuming an axiom is not progress.”

Response. It is custody discipline, and the lemma says so in its statement (§41.1): deriving an arrow from arrowless ingredients is exactly the trap the cohort’s anti-arbitrage rule exists to prevent. What Lemma 2.2 adds is not the direction — that is consumed, with its source named — but the decomposition into three auditable links and the *exact* exchange-class identity $\Delta S_{\text{gen}} =$ spent relative entropy (Appendix O, Step 4), with Wall’s theorem (free and superrenormalizable fields, arbitrary cuts) registered as the strongest partial result on the declared open ceiling (§41.3), not as a rival.

Objection O-BH-11 (the mathematical relativist, on the extremal sector)

“Extremal entropy is subtle — Sen’s entropy function, the AdS₂ throat, the discontinuity

between extremal and near-extremal counting. Your appendix breezes past all of it.”

Response. The appendix claims only the standard statements (entropy survives at $T = 0$; Nernst unattainability; the degenerating cigar, Appendix P, Step 4), cites the throat literature where it is relevant (Sen, 2013), and places the GG-Theory extremal counting — holonomies at pixels — at the same open grade as the species sum. No resolution of the extremal-vs-near-extremal discontinuity is claimed; the module’s load-bearing statements are all non-extremal.

Objection O-BH-12 (the numerical relativist, on fixed backgrounds)

“Every theorem in your Part V lives on a fixed background. Evaporation is back-reaction by definition. Your entropy law is proved exactly where it is not needed.”

Response. The regime declarations are in the statements: T-BH-5 is quasi-free on a static exterior; T-BH-3/T-BH-4 are model level, verified; the evaporation clock (Appendix K, Step 3) is adiabatic Stefan–Boltzmann. The honest chain is: exact where stationary, adiabatic where slow, and open where neither — the 4D-flat rigor item is pre-fenced in the promotion roadmap (Appendix W). We note the same limitation governs every framework’s evaporation statements, including the replica-wormhole computations, which are also saddle-point statements.

67 The curved-spacetime field theorist

Objection O-BH-13 (the Kay–Wald reader, on bifurcation surfaces)

“Kay–Wald uniqueness holds on bifurcate Killing horizons. A collapse spacetime has no bifurcation surface, so T-BH-5’s citation does not license your physical case.”

Response. Agreed, and that is why the collapse case is a separate theorem rather than a footnote: T-BH-4 proves the late-time factorization — universal thermal channel plus decaying transient — at the model level, importing Hawking’s gravitating computation with attribution for the full statement. T-BH-5 is scoped to the static exterior where Kay–Wald applies, and the composition of the two is stated, not blurred (§28). The division of labor is printed in both statements.

Objection O-BH-14 (the microlocal analyst, on the Hadamard class)

“Which of your states are Hadamard, and does the OGGEp projection preserve the class? If not, your ‘regular state’ talk is decorative.”

Response. The regular states invoked are precisely the Hadamard ones — T-BH-5(iii) consumes the Kay–Wald theorems, whose regularity hypothesis is Hadamard form — and the OGGEp operation on them is restriction plus partial trace, which cannot manufacture singular behavior the parent state lacks. What is *not* claimed: a proof that the full 6D-to-4D projection preserves the Hadamard class for arbitrary bulk states; that belongs to the 4D-flat rigor item (Appendix W) and is flagged there, not here.

Objection O-BH-15 (the algebraic field theorist, on new words)

“Your ‘derived complex structure’ is the positive-frequency choice every textbook makes.

Renaming a choice does not make it a theorem.”

Response. The renaming is not the claim; the theorem is. T-P5-B1 proves that for quasi-free fields on a static arena the four analytic devices — positive-frequency selection, continuation through the complexified horizon, imaginary-time semigroup, holomorphic tube — are one datum, so selecting any one fixes the others with no further input. T-BH-5 consumes that at the horizon: what the older literature carried as three independent prescriptions is one licensed move.

Objection O-BH-16 (the constructive field theorist, on interactions)

“Your thermality proofs are quasi-free. The world is interacting. What survives?”

Response. Declared, graded, and routed: the quasi-free grade is printed in T-BH-5’s statement; the GSL beyond the exchange class is a registered open ceiling, with Wall’s theorem (free and superrenormalizable fields, arbitrary cuts) its strongest partial result (§41.3); and the promotion path — an interacting-grade shore — is roadmap item (iv) (Appendix W). The coefficient $1/4$ is insulated: the conical argument fixes it at the level of the gravitational effective action, independent of the matter sector’s interaction strength (Theorem 1).

Objection O-BH-17 (the entanglement-entropy specialist, on edge modes)

“Gauge fields and gravitons have edge modes and the Kabat contact term; ‘entanglement entropy = geometry’ is notoriously subtle exactly there. Your identification glosses a live technical literature.”

Response. Cited and carried as a known complication, not glossed: the honest-accounting section (§31.10) flags the gauge/graviton bookkeeping explicitly, with the contact-term literature (Kabat, 1995; Solodukhin, 2011). The clean statement — minimally coupled scalars — is where the worked derivation lives; the gauge-sector subtlety is part of the species-sum computation, which is the module’s declared principal open item (F-BH-9). The coefficient claim survives because it is a statement about the gravitational normalization ratio, not about any one species’ modes.

Objection O-BH-18 (the trans-Planckian skeptic, on the exponential lens)

“Hawking quanta observed at late times blueshift past every scale near the horizon, including your lock scale. Your derivation inherits the trans-Planckian problem and adds a scale at which it must break.”

Response. Three honest layers. (i) The analytic kernel — the one-sided spectrum and the continuation cost $e^{-\pi\omega/\kappa}$ — is not a mode-by-mode ultraviolet statement; T-BH-4’s branch-cut form shows the Planck factor arising from infrared-robust asymmetry, not from trusting arbitrarily short wavelengths. (ii) The analog-horizon record (§55) shows the thermal spectrum surviving explicit ultraviolet modification — imported as laboratory structure evidence with attribution. (iii) What a lock-scale dispersion computation would add — the GG-Theory analogue of the dispersive-derivation literature — has not been carried out and is not claimed; it is a legitimate audit item. The scale the objection says we add is the scale the falsifier cluster tests at colliders.

68 The string theorist

Objection O-BH-3 (the holographer, on what is new)

“AdS/CFT and the islands program already resolve the paradox; what is new?”

Response. Those frameworks resolve it in anti-de Sitter toy models or as an abstract boundary duality whose real-universe dual is unknown; GG-Theory works in the Standard-Model-containing universe. The islands formula is not contradicted — it is recovered as the GFF ledger at the horizon stratum (§33). What is new is the constitutional *source* (why generalized entropy and non-naïve subsystem reconstruction had to appear at all) and the exact $1/4$ for the *generic*, *non-supersymmetric* black hole, which neither AdS/CFT nor the explicit string count delivers.

Objection O-BH-5 (the D-brane counter, on explicit microstates)

“Strominger–Vafa is an explicit state count; you have only the coefficient.”

Response. Granted, and stated plainly throughout. The asymmetry is one of *completeness versus generality*: the string count is more complete in its protected (supersymmetric, extremal) sector; GG-Theory is general (Schwarzschild, Kerr, Reissner–Nordström) and fixes the coefficient without tuning, but its explicit enumeration — the species sum to a definite number — is the one open computational item (Appendix R). The microstates are *named* (hidden-fiber S_θ^1 holonomies at horizon cells, §31.6); they are not yet summed.

Objection O-BH-19 (the ultraviolet enforcer, on the consistency premise)

“Without supersymmetry or a string completion, what controls your 6D bulk in the ultraviolet? ‘AX-4 requires a consistent completion’ is a promissory note, not a theory.”

Response. It is a premise, and Theorem 1 carries it as one of its three named premises rather than hiding it. The difference from an unpayable note: the completion’s low-energy fingerprints are the cohort’s falsifier cluster — the lock scale and its collider criterion, the Kaxion window, the neutrino ordering — so the premise is exposed to instruments (§58) rather than sheltered. A conditional theorem with a testable condition is the honest form of the claim; the unconditional form would be the dishonest one.

Objection O-BH-20 (the fuzzball theorist, on horizon structure)

“The fuzzball program replaces the horizon with microstructure; information never enters a causally sealed region. You keep the classical horizon and move the problem into bookkeeping.”

Response. The two programs make opposite wagers, and the module prints the discriminating instrument: fuzzball-class microstructure at the horizon scale generically feeds post-merger echoes or QNM spectrum modification, and GG-Theory predicts their absence — clean general-relativity ringdown, with F-BH-2/F-BH-3 as permanent kill rows (§48, §58). A reproducible echo train retires this module and hands the sector to horizon-structure programs. Bookkeeping with a kill row is a physical claim.

Objection O-BH-21 (the ER=EPR advocate, on the bridge branch)**“Your late-time ‘bridge branch’ is ER=EPR wearing your notation.”**

Response. The kinship is acknowledged where the idea appears (§36), with the proposal credited (Maldacena and Susskind, 2013). The difference in register: ER=EPR is a conjectural identification of entanglement with geometry; the GFF branch switch is a ledger statement — which coarse-grained saddle dominates an explicitly defined free-energy functional (Eq. (44)) — whose turnover time is fixed by an audited counting premise (L-BH-Fresh). The reading of the dominant branch *as* bridge geometry is dictionary-grade and labeled so in the island dictionary (Appendix N).

Objection O-BH-22 (the anomaly connoisseur, on arithmetic versus anomalies)**“In string theory the entropy coefficient traces to central charges fixed by anomaly cancellation — deep consistency. Your $4\pi/16\pi$ is arithmetic. Why should arithmetic beat anomalies?”**

Response. They are the same species of statement at different addresses: a normalization fixed by a consistency condition upstream of the entropy computation. The conical 4π is the Gauss–Bonnet normalization of distributional curvature; the 16π is the Einstein–Hilbert normalization fixed by the Newtonian limit; the theorem’s content is that the entropy coefficient is their ratio and nothing else enters (§31.5). Anomaly consistency enters GG-Theory where it should — the GG-A9-SM field content that the species sum runs over — and that is exactly the open, sign-sensitive part (F-BH-9). The ratio is not deeper than anomalies; it is orthogonal to them, which is why it survives their bookkeeping.

Objection O-BH-23 (the AdS₃ practitioner, on the missing Cardy check)**“Every serious entropy framework checks itself against BTZ/Cardy. Where is yours?”**

Response. Not performed, and conceded in those words. The module’s arena is the four-dimensional, asymptotically flat, Standard-Model-containing case, and its AdS-facing content is the Hawking–Page notebook (Appendix J), a template rather than a check. A BTZ/Cardy audit of the GG-Theory normalization — does the conical ratio reproduce the Cardy count where both are computable? — is a legitimate outstanding cross-check we have not done. It is recorded here so the concession has an address.

69 The islands and replica-wormhole theorist**Objection O-BH-24 (the Page-curve deriver, on translation versus addition)****“Your two-branch race is our min-ext, your turnover is our Page time, your dictionary table admits it. What do you add beyond a translation with constitutional decoration?”**

Response. Three items, each with an address. *Source:* why generalized entropy exists at all — the ledger $GFF = F_{GEP} + F_{CEP}$ is derived from the open-system constitution rather than posited as a prescription (§33). *Premise audit:* the islands literature’s random-state/typicality input is consumed here as an explicit, flagged condition with an expiry mechanism (L-BH-Fresh; the

ETH import is named in Appendix K, Step 4) — silent in most island derivations. *Reach*: the exact $1/4$ for the generic, non-supersymmetric, flat-space black hole, which the AdS derivations do not deliver. The translation is real and credited; the additions are graded.

Objection O-BH-25 (the ensemble interpreter, on factorization)

“Replica wormholes compute averaged quantities; the factorization puzzle suggests semiclassical gravity is an ensemble theory. Your single deterministic bulk is on the wrong side of that lesson.”

Response. Flagged as a genuine open commitment, not answered away. The module’s position: the bulk is a single theory; ensemble-like behavior is the coarse-grained ledger’s, and the residual question — whether the boundary description’s late-time correlators factorize as a single theory’s must — is the module’s adopted open problem (§66, the curved-space existence residue; roadmap item (vi)). If the ensemble reading proves obligatory rather than effective, the deterministic-bulk interpretation takes the hit, and the register says so. This and O-BH-4 are the two commitments this Part flags rather than defends.

Objection O-BH-26 (the gravitational path integrator, on derivation depth)

“The QES prescription is derived from the gravitational path integral in AdS. Your ledger ‘derivation’ is a heuristic by comparison.”

Response. For derivation depth in AdS, conceded without qualification. The asymmetry runs the other way in scope: the path-integral derivation is anchored where the dual is known, and its extension to the real universe is itself heuristic; the GG-Theory chain is Derived-conditional with premises named, in the flat, Standard-Model-containing case. The QES equivalence statement is registered at its honest grade — a deck proposition, not a theorem (Appendix V) — and the two programs’ asymmetry (depth versus address) is stated as such in the comparison section (§64).

Objection O-BH-27 (the complexity theorist, on the missing resource)

“Python’s lunch, reconstruction complexity, the exponential cost of interior decoding — the modern account of the interior is complexity-theoretic. Your paper has a trace and no complexity anywhere.”

Response. Correct, and no complexity claim is made. The DGI restriction plays the module’s role of reconstruction obstruction at the information level — which algebras are accessible, not how hard the decoding is (§33). Whether the complexity-theoretic account is the GDOS restriction seen through a resource theory is an attractive bridge question we have not developed; it is noted here as unclaimed territory rather than annexed.

70 The canonical-quantization referees

Objection O-BH-2 (the loop-quantum-gravity theorist, on the Immirzi parameter)

“You have merely replaced the Barbero–Immirzi tuning with a tuning of your own.”

Response. No free parameter enters the coefficient. In loop quantum gravity the Barbero–

Immirzi parameter γ is fixed *by demanding* that the area-law coefficient come out $1/4$; it is a fit. Here $1/4 = 4\pi/16\pi$ is the ratio of the conical-deficit curvature normalization to the Einstein–Hilbert normalization (§31.5, Theorem 1); both numbers are fixed before the entropy is computed. There is nothing to tune.

Objection O-BH-28 (the isolated-horizon builder, on glibness)

“Isolated-horizon state counting is a first-principles derivation too, given gamma. Your one-line dismissal of a thirty-year program is glib.”

Response. The program’s depth is real and the respect is meant: the isolated-horizon framework’s boundary-degree-of-freedom counting is structurally the closest relative of the pixel-holonomy counting this paper names (§31.6). The asymmetry O-BH-2 states is narrow and remains: one framework fixes its free parameter by the entropy match; the other has both normalizations fixed upstream. If the Immirzi parameter is ever derived independently — as several LQG lines attempt — the comparison reopens on equal terms, and this Part will owe an update. A standing test, not a dismissal.

Objection O-BH-29 (the quantum-geometry theorist, on pixels)

“Your ‘horizon pixels with $U(1)$ holonomies’ is ’t Hooft’s ledger and LQG’s punctures with a new fiber. What is distinct?”

Response. The lineage is printed — ’t Hooft’s ledger is cited at the microstate section (§31.6) — and the distinct content is the *identity* of the fiber: the counting states are holonomies of the specific compact factor S^1_θ carrying $U(1)_{Q_i}$, the same fiber whose winding classes carry charge superselection elsewhere in the cohort, so the microstate identification is imported structure, not a new posit for the horizon. What is shared with LQG — area quantization granularity — and what is not — no spin networks, no Immirzi freedom — is tabulated in the comparison section (§64).

71 The emergent-gravity theorist

Objection O-BH-1 (the Sakharov lineage keeper, on the sign problem)

“This is just Sakharov induced gravity, which has a known sign/magnitude problem.”

Response. It is in the induced-gravity family, and we say so (§64). But the result that matters — the coefficient $1/4 = 4\pi/16\pi$ — is fixed by the conical replica (§31.5) *independently* of the species content, regulator, and sign. The Sakharov sign/magnitude difficulty lives entirely in the *value* of the induced G_N , i.e. the species sum, which we flag as open (Appendix R). Two things distinguish GG-Theory from generic induced gravity: GGEP makes G_N *fully* induced (there is no bare term, so the entanglement entropy is not a correction but the entropy), and the microstates are identified geometrically (the hidden-fiber holonomies). The coefficient is not hostage to the sign problem; the magnitude is, and is labeled accordingly.

Objection O-BH-7 (the same referee, pressing on closure)

“The species sum does not close, so G_N is unconstrained and the result is empty.”

Response. The closure is *structural*: the Kaluza–Klein tower sum reproduces the scale dimensional reduction that defines G_N (Appendix R), and G_N is a member of GU-5 — the geometric scale conversion factor, an input, not a quantity the entropy must independently predict. Given a consistent ultraviolet completion (required by AX-4), the entanglement entropy is $A/4G_N$ by construction; the precise leading numerical coefficient is what remains open. The $1/4$ does not depend on it. [Reconciliation with O-BH-1: “input” refers to GU-5 anchoring, not to exemption from closure — the sign-resolved sum must still reproduce the fully induced $1/G_N$, and F-BH-9 prices the failure.]

Objection O-BH-30 (the thermodynamic-gravity theorist, on Jacobson)

“Jacobson derived the Einstein equations from the Clausius relation in 1995. If gravity is already thermodynamics, your constitutional derivation is a longer road to a known destination.”

Response. The two roads run in opposite directions and meet: Jacobson assumes the entropy-area law locally and derives the field equations (Jacobson, 1995); this paper derives the entropy-area law (coefficient included) from the constitution and consumes the field equations’ normalization. Neither subsumes the other; together they close a loop that neither closes alone — which is precisely how the related-work section files the relation (§64). The GG-Theory-specific content is the part Jacobson’s argument cannot reach by construction: which microstates, which fiber, and why the coefficient is $4\pi/16\pi$.

Objection O-BH-31 (the Visser reader, on the sign that kills)

“Visser showed the induced Newton constant can come out with the wrong sign for realistic field contents. If the A9 spectrum’s sign-resolved sum misbehaves, everything downstream is rubble.”

Response. Not rubble — a fired kill row, which is the designed outcome: F-BH-9 registers exactly this failure (the sign-resolved sum fails to reproduce the induced $1/G_N$ — wrong sign or off closure), and the deck shows what it takes with it — the closure claim and O-BH-7’s answer — and what it cannot touch — the coefficient $1/4$, whose derivation nowhere uses the sum (Theorem 1; the insulation argument is O-BH-1’s). The module’s exposure is stated, bounded, and priced; that is what an internal falsifier is for (Table 9).

72 The gravitational-wave observer

Objection O-BH-32 (the LVK data analyst, on free-riding)

“GW250114 confirmed the area theorem — a classical result. Your experiments Part cites it as if it supported your quantum story.”

Response. The record section scores it as exactly what it is: confirmation of the classical half of the GSL story — the area theorem under merger — with the quantum half (entropy production

toward the KMS face) explicitly marked as not directly tested by any interferometer (§55). The pre-registered verdict discipline (§57) exists to prevent precisely the inflation this objection fears: each instrument row states what a result does and does not touch before the result arrives.

Objection O-BH-33 (the ringdown spectroscopist, on waveform silence)

“Name one waveform feature that distinguishes your framework from general relativity, or concede the GW sector tests nothing of yours.”

Response. The honest answer is a null with teeth: GG-Theory predicts clean general-relativity ringdown — no echoes, no QNM spectrum modification, no horizon-scale structure — so every post-merger anomaly search is a live kill attempt on F-BH-2/F-BH-3, permanently (§48; §58). A framework whose horizon sector predicts *silence* is falsifiable by noise, exactly as no-hair general relativity is; its positive discriminators live where the deviations are unsuppressed — the collider lock cluster, not the interferometers. The division is printed in the falsification catalog (Appendix X).

Objection O-BH-34 (the primordial-black-hole hunter, on ratios only)

“Your PBH row is ‘ratio structure only, anchor open.’ A prediction with a floating anchor is not a prediction.”

Response. Conceded as to the anchor, and pre-fenced as such since the charter: the absolute mass scale needs cosmology-side input the module does not own (§49; roadmap item (vii)). The ratio ladder is nonetheless falsifiable on its own terms: any two confirmed detections begin testing the spacing claim, and the registered kill (F-BH-4) fires on a well-populated mass function off the e^7 ladder — no anchor required either way. A two-point test with a declared missing one-point input is a partial prediction honestly labeled — the alternative of retrofitting an anchor is what the register forbids.

Objection O-BH-35 (the LISA forecaster, on the EMRI decade)

“EMRIs will map Kerr geometry to extraordinary precision in the 2030s. What number of yours do they test?”

Response. They sharpen the Kerr-null rows: the module predicts pure Kerr multipoles with deviations carrying the lock scale (the scale is P1’s; the suppression bookkeeping is Part VII’s, §46) — quantitatively negligible at EMRI sensitivity — so EMRI precision converts the Kerr-null from a test into a vice: any confirmed non-Kerr multipole at the 10^{-4} class retires the module, unconditionally, under the registered row F-BH-10 (Table 9). No positive EMRI signature is claimed; the milestone’s pre-registered verdict says so.

73 The horizon-imaging astronomer

Objection O-BH-36 (the EHT calibrator, on padding)

“Shadow diameters test the metric at ten gravitational radii, not horizon quantum structure. Their presence in your evidence table pads the case.”

Response. The table’s scoring column is the anti-padding device: the EHT rows are scored

as metric-level Kerr consistency and object-class existence, nothing more (§47; §55). Their live function is exclusionary — polarization-resolved epochs increasingly disfavor horizonless exotica, which sharpens the clean-horizon prediction this module shares with general relativity. Consistency rows priced as consistency are context, and the table prices them.

Objection O-BH-37 (the accretion astrophysicist, on plasma mud)

“Sgr A* variability, M87* jet loading, hot-spot orbits — everything EHT measures is filtered through messy plasma. Nothing clean about horizons survives that.”

Response. Largely conceded: the module quotes the imaging record for existence, scale, and Kerr consistency, and hangs no quantum claim on any plasma-filtered observable. The census entries (masses, shadows, entropy bookkeeping in §46) are displays, and the honest-accounting language throughout Part VIII marks which rows are astrophysics-limited. Where the objection bites — any future claim that imaging tests horizon microphysics — the refresh discipline (§59) is the standing guard against it.

Objection O-BH-38 (the observational pragmatist, on an unobservable temperature)

“A solar-mass hole radiates at sixty billionths of a kelvin into a 2.7-kelvin sky. Your entire thermal sector is untestable in principle, and you know it.”

Response. For direct astrophysical detection, yes — printed with the numbers in the reference card and never claimed otherwise. The thermal sector’s exposure runs on three other rails, each named in the falsification catalog (Appendix X): laboratory analog horizons for the spectrum’s structural universality (imported with attribution); the internal-consistency kills, where the same T_H underwrites the GSL identity and the species-sum closure (F-BH-9); and the PBH window, where a landed anchor would make final-burst thermality observational (O-BH-34’s concession applies). Untestable-in-principle is the wrong verdict for a quantity that three independent registers can still kill.

74 The quantum-information theorist

Objection O-BH-39 (the Hayden–Preskill reader, on fast mirrors)

“Old black holes reflect infalling information in scrambling time. Your Page-turnover narrative — information waits until half the entropy is out — ignores twenty years of fast-scrambling results.”

Response. The two statements live on the two sides of the same clause: Hayden–Preskill mirror behavior *presupposes* a hole already past its Page point (maximally entangled with collected radiation) — which in this module’s language is precisely a stale port (Hayden and Preskill, 2007; §36.3). The freshness lemma quantifies when the mirror switches on: the qubit-model benchmark shows the information return crossing at the Page point with the predicted growth ratios (Appendix U). No tension; one clause, read before and after its expiry.

Objection O-BH-40 (the modular-theory specialist, on dressed identities)

“Your ‘exact exchange-class identity’ is the standard first-law-of-entanglement / relative-

entropy identity that Casini used in 2008. You have dressed a known engine in cohort livery.”

Response. The engine is credited at every appearance — Casini’s bound and Wall’s monotonicity are the named machinery (§41.1; Appendix O) — and the identity itself is displayed as what it is. The cohort-specific additions are the custody decomposition (the arrow consumed from a named source rather than smuggled through a limit), the equality-point corollary C-BH-1 fusing the identity with the conditional-GGA crossover (§42.2) — a statement with no counterpart in the modular literature — and the grade fence separating exchange-class-exact from Wall-class-open. Livery over a credited engine, with new corollaries: that is the claim, exactly.

Objection O-BH-41 (the channel theorist, on the translation step)

“DPI is about channels; the horizon is a region of spacetime. The step where a causal boundary becomes a CPTP map is where all the physics hides, and you have named it rather than proven it.”

Response. Named, graded, and partially proven, in that order: the translation is the third link of Lemma 2.2 precisely so it can be audited separately. Its proven core is L-BH-Port — at a horizon the trace is forced by causal structure plus the axioms, so the channel is not an analogy but the observer’s literal description (§25.2, Appendix O, Step 1). Its open remainder — the interacting, wedge-algebraic form where tensor factorization fails — is the declared Wall-class ceiling (§41.3). The objection’s instinct is right: that is where the physics hides, which is why the register keeps a lamp on it.

Objection O-BH-42 (the error-correction builder, on the missing code)

“The islands community has explicit quantum error-correcting code models of bulk reconstruction. You have a partial trace and a slogan about ‘non-naive reconstruction.’”

Response. Declined rather than defended: no code model is constructed in this paper, and the DGI-constrained reconstruction of the accessible algebra (§33) is the information-level counterpart, not a substitute. The constitutional order of explanation is port first, code later — the trace structure is forced, and any code model is a representation of it. Building the explicit GG-Theory code — which subspaces, which recovery maps, at which grade — is a well-posed project the module leaves on the table, and this entry is its acknowledgment.

Objection O-BH-43 (the monogamy enforcer, on AMPS)

“AMPS forced a trilemma: give up equivalence, unitarity, or field theory at the horizon. Every framework must name its casualty. Name yours.”

Response. The casualty is the trilemma’s silent premise: that the early radiation and the interior partner are *independent* systems whose entanglements compete. In GDOS they are two dictionary readings of the same bulk modes through one port, and monogamy binds independent systems, not one system doubly described (§37). The price paid is real and printed: late-time interior reconstructions are DGI-restricted — the naive interior algebra is not independently accessible — which is the module’s version of ‘field theory at the horizon’ bending. No firewall, no complementarity gymnastics, one restriction, stated.

75 The quantum-foundations referee

Objection O-BH-4 (the Bell-theorem referee, on hidden variables)

“A deterministic six-dimensional bulk is a hidden-variable theory; Bell’s theorem forbids it.”

Response. This is the strongest live objection, and we mark it as a commitment rather than a settled result. The bulk is *local in six dimensions*; the apparent four-dimensional non-locality measured by Bell/CHSH is read as a projection artifact of the bulk-to-boundary map — the “holographic shadow” developed in (Arisaka, 2026, A8-QM) — not as bulk non-locality. No-signaling is preserved because the boundary CPTP channel cannot transmit superluminally. Whether this reading is fully consistent with all Bell-type constraints in the gravitational setting is, in our judgment, the question on which the deterministic-bulk interpretation will ultimately be tested; we do not claim it closed. The commitment now has a cohort address and a grade: the selection of the real category is L-Constitution-1 (candidate grade) in GG-A1-Constitution, with its own falsifier row F-Constitution-10 — a non-geometric CP phase, an unshared i , or a quaternionic Born residue would kill it (Arisaka, 2026, A1-Constitution) — so this objection stress-tests a named, graded, falsifiable claim rather than a metaphysical preference.

Objection O-BH-44 (the relational purist, on perspectival entropy)

“If entropy is the entropy of a channel image, thermodynamics becomes observer-relative and ‘the black hole’s entropy’ denotes nothing objective.”

Response. The module embraces the relativity and locates the invariant: operational entropies are port data — different lawful observers, different ledgers — while the bulk state and the custody rules that map it to any port are shared (Appendix O, overture). What rescues objectivity is that the horizon port is *forced*: every exterior observer class inherits the same trace structure, so $A/4G_N$ is observer-class-invariant even though ‘entropy’ is channel-relative. Perspectival at the dictionary level, absolute at the constitutional level — the same division the cohort uses for vacua and clocks.

Objection O-BH-45 (the Born-rule auditor, on the unfinished chain)

“Your cohort’s Born rule is ‘closed at the relaxation level’ with a named analytic residue. A black-hole paper whose thermality claims rest on Born-rule probabilities should carry that flag on page one, not in a sibling paper.”

Response. The dependency is real and the separation is stated where it matters: every thermality statement in this paper is a statement about reduced states and their KMS structure — trace-level facts, valid under the standard reading of the formalism independently of how the cohort ultimately grounds probability. If the field-grade representation step failed, the casualty would be the cohort’s *derivation* of the Born weights (Arisaka, 2026, A8-QM), not the channel identities used here, which would revert to standard-quantum-mechanics grade unchanged. The flag lives in the source ledger (Appendix S), which is this paper’s page one for consumed premises.

76 The historian and the philosopher

Objection O-BH-6 (the philosopher of physics, on reinterpretation)

“Is this a physical theory or a reinterpretation of known results?”

Response. Both, in honestly separated parts. Much of the information-paradox, Page-curve, and firewall content is reinterpretation, and we label it as such. The exact-1/4 induced-gravity derivation and its reach to generic black holes are new results. The pre-declared kill criteria (§67) make the framework falsifiable, not merely narrative. A reinterpretation that also yields a new derivation and pre-registered falsifiers is doing more than relabeling.

Objection O-BH-46 (the rhetoric critic, on calibrated honesty as armor)

“Grades, decks, ledgers, ‘honest verdicts’ — the apparatus of calibrated honesty reads as pre-emptive rhetoric: concede the adjectives, keep the nouns, and disarm referees by performing humility.”

Response. The apparatus has teeth only if it can bite its owner, so point at the teeth: every empirical concession is paired with a kill condition (the deck’s Exposure column, Table 24, keyed to Table 9); the register’s standing challenge is printed — find one claim above its table grade and it is a defect to report (Appendix V); and two objections in this very Part are logged as open commitments, not answered (O-BH-4, O-BH-25). Performed humility cannot survive an audit that is invited in writing. If a reader finds nouns without kill rows, the criticism stands and the defect is ours.

Objection O-BH-47 (the historian of the paradox, on announcement fatigue)

“The information paradox has been ‘solved’ every decade since 1976 — complementarity, then fuzzballs, then firewalls dissolved, then islands. Why should your resolution age better?”

Response. The paper declines the word: its paradox content is graded reinterpretation (O-BH-6’s concession), its new results are two theorems and a set of identities with declared premises, and its survival criterion is externalized — the falsifier timeline with dated instrument rows (§58) rather than community persuasion. Frameworks age badly when their claims outrun their registers; the design here is that the register, not the narrative, is the load-bearing structure. Whether that discipline holds is itself testable, which is the only answer to fatigue we know how to print.

77 The student

Objection O-BH-48 (the first-year reader, on where to start)

“The paper is nearly 180 pages and Part IV assumes three companion papers. Where does a student actually start?”

Response. The annotated reading guide (Appendix Y) is the door: six shelves from Bekenstein 1973 forward, then this paper re-read from Part IV. Inside the paper, the course order is Block

II of the appendices — Appendix H is the first calculation to reproduce by hand, and each chapter closes its Steps with the physical picture. The three questions to ask of any section — what is derived here, what is consumed from a companion, what is imported from the literature — are answered mechanically by the source ledger (Appendix S). Total prerequisite: one quantum field theory course and the will to check algebra.

Objection O-BH-49 (the skeptical student, on numerology)

“Why is $1/4 = 4\pi/16\pi$ not numerology? Any ratio of small multiples of pi can be arranged to give $1/4$.”

Response. Because neither number is chosen: the 4π is the Gauss–Bonnet normalization of the conical deficit’s distributional curvature — change it and the Euler characteristic of the two-sphere stops being 2 — and the 16π is the Einstein–Hilbert normalization — change it and Newton’s constant stops matching Cavendish (§31.5, worked in the replica computation). Numerology is the freedom to arrange; here the theorem’s entire content is the *absence* of freedom: both constants are pinned by physics that has nothing to do with entropy, and the coefficient is their quotient. O-BH-2 makes the same point against a framework that does tune.

Objection O-BH-50 (the confused student, on whether the hole is hot)

“The infalling observer crosses the horizon and measures nothing thermal. The exterior observer sees a furnace. Which is the black hole, hot or cold?”

Response. Both, lawfully, because temperature is port data, not bulk data: the exterior observer’s description is the OGGEF trace of the bulk state, and that trace is thermal (Eq. (18)); the infalling observer keeps both partners of each mode pair and her description is the pure squeezed state, temperature-free. Neither is wrong; they hold different pages of the same ledger — the dictionary owns the readings, the bulk owns the state. Appendix H, Step 6 and the de Sitter mirror case (Appendix Q) are the worked homes of this answer.

78 The cohort-internal referee

Objection O-BH-51 (the dependency auditor, on blast radius)

“C3 consumes T-P5-B1, L-GDOS-Fresh, L-Constitution-1, and the A2 arena. For each: if the upstream key is downgraded tomorrow, what exactly breaks here?”

Response. The inheritance table (Appendix T) is the map; here is the blast radius row by row. T-P5-B1 downgraded: T-BH-5’s clause (ii) reverts from licensed to prescribed — the continuation becomes the literature’s standing assumption, every number unchanged, the grade drops. L-GDOS-Fresh downgraded: the Page turnover keeps its counting derivation (Appendix K) and loses its mechanism reading; T-BH-2 reverts to named description. L-Constitution-1 (already candidate): consumed only at the funnel and O-BH-4; its failure kills the real-category grading, touching no entropy statement. The GG-A2-GG6 arena: constitutional — its failure is cohort-wide, not C3-specific. No single upstream downgrade reaches Theorem 1’s three named premises.

Objection O-BH-52 (the firewall inspector, on convenient separability)

“F-BH-9 could fire and you would shrug: ‘the closure claim dies, the 1/4 survives.’ Is the module’s fate really that separable, or is the firewall drawn where it is convenient?”

Response. The separation is structural, not rhetorical, and it is checkable: Theorem 1’s statement lists its three premises, and the species sum appears in none of them — the sum enters only the closure appendix (Appendix R), whose own honest verdict already states the exposure. The deck’s Exposure column (Table 24) shows the same partition from the other side: F-BH-9 retires L-BH-Tower’s closure reading and O-BH-7’s answer, and leaves T-BH-1’s row untouched because their premise sets are disjoint. A firewall drawn by premise bookkeeping can be audited by reading the premises; convenient ones cannot.

Objection O-BH-53 (the custody clerk, on the unregistered proposal)

“The real-bulk register displays a proposed entry that A8/P4 have not ruled on. Printing it before the home papers decide is claim inflation by typography.”

Response. The typography is the anti-inflation device: the entry is displayed under the word *proposed*, with custody stated — numbering belongs to the home papers — and the deck carries it at proposal grade with no theorem key (§29; Appendix V). The alternative — consuming the reading silently while awaiting registration — is the actual inflation risk, because the dependency would exist without an address. When the home thread rules, a micro-pass updates the display either way; the proposal document is already lodged with the home papers’ materials.

Summary. Of the fifty-three: forty-four we regard as answered, each with pointers into theorems, registers, or kill rows; seven are conceded asymmetries or declined claims, said in those words (O-BH-5 on explicit counting, O-BH-23 on the missing Cardy check, O-BH-26 on derivation depth in AdS, O-BH-27 on complexity, O-BH-34 on the floating PBH anchor, O-BH-37 on plasma-filtered observables, O-BH-42 on code models); and two are genuine open commitments that we flag rather than claim to have settled (O-BH-4 on Bell-type constraints, O-BH-25 on factorization). The seed scorecard of the seven-objection era is superseded by this count. We consider this an honest scorecard for a framework at this stage, and the standing challenge of the named-statement deck — nothing is claimed above its table’s register — extends to every response above.

Part XIV

References

Cited works, each with a relevance note linking it to the present paper

Unlike a conventional reference list, every entry below carries a short *Relevance* note: a sentence or two saying what the cited work contributes and where in the present paper it bears, closing with the place where it is first cited. The aim is to turn the bibliography from a bare list of sources into a navigable companion — a reader can move through the references and see, at a glance, which derivation, table, or section each one underwrites.

The notes are written for two readers at once. For the newcomer, they trace the intellectual lineage of the present derivation: from Schwarzschild’s 1916 solution and the golden-age theorems, through Bekenstein’s entropy, Hawking’s radiation, and Page’s curve, to the holographic era — state counting, AdS/CFT, and the islands program — and then to the twenty-two-paper GG-Theory cohort, in which the present paper reads the horizon as an OGGEp port, the strong-gravity member of the program’s consequence branch. For the specialist, they make the paper’s debts explicit and checkable — each note names the specific result it supports, so that no citation is merely decorative, and each ends with its first-citation site (a section or appendix link where the enclosing unit is labeled, otherwise the Part).

The GG-Theory cohort papers are cited throughout under the canonical keys and titles fixed by the cross-paper coherence map of the program; the present paper’s place in that cohort — the black-hole member of the consequence series, running on the GDOS boundary law of GG-A5-GDOS — is set out in Part II. The cohort papers are listed first, in cohort order; the external references follow, alphabetically by first author.

A. Papers by Katsushi Arisaka (the twenty-two-paper GG-Theory cohort)

K. Arisaka (2026). *The GG-Constitution: Four Axioms and Seven Postulates of the GG-Theory Program, From Which GG-6, BI-6, OGN-5, and GDOS Descend as Theorems.* Preprint, available at <https://preprints.arisaka-gg.org/a1-constitution/>. doi:10.5281/zenodo.22063867.

Relevance. The constitutional root of the twenty-two-paper cohort: the four axioms and seven postulates from which the bulk, the boundary law, and this paper’s premises descend. Cited for AX-1 (local causality: the horizon as a one-way causal boundary), AX-2 (GGEP: no bare four-dimensional Newton constant — the premise of Theorem 1), AX-3 (OGGEp: observation as the bulk-to-boundary trace whose geometrically forced instance is the horizon port), AX-4

(MSMF: the no-dial selection whose admissibility filter carries this paper’s ultraviolet-completion premise), and L-Constitution-1 with its falsifier F-Constitution-10 (the candidate-grade reality register consumed at the central-claim funnel and stress-tested by O-BH-4). First cited in the front matter.

K. Arisaka (2026). *GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory)*. Preprint, available at <https://preprints.arisaka-gg.org/a2-gg6/>. doi:10.5281/zenodo.22064298.

Relevance. The arena companion: Theorem 0 forces the spacetime–scale–phase arena from AX-1+AX-2, and Theorem 2 fixes the canonical scale-profile metric with no free parameters (k , R_θ , r_c all eliminated). Cited in Part II for the arena and at Step 1 of the entropy derivation, whose fiber integrals are exactly Theorem 2’s profile; the renormalization-group reading of the scale direction (the holographic-RG construction) also lives there. First cited in the front matter.

K. Arisaka (2026). *Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6)*. Preprint, available at <https://preprints.arisaka-gg.org/a3-bi6/>. doi:10.5281/zenodo.22064300.

Relevance. The anomaly-freedom proof of the BI-6 modified Bianchi identity of the six-dimensional bulk. Cited where the bulk gauge structure $SU(5) \times U(1)_B \times U(1)_{Q_i}$ with BI-6 enters: the field content over which the induced-gravity species sum of Appendix R runs is fixed by the anomaly-free bulk theory this companion establishes. First cited in the front matter.

K. Arisaka (2026). *From Euclid to GG-6: The Realization of Euclid’s Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics*. Preprint, available at <https://preprints.arisaka-gg.org/a4-euclid/>. doi:10.5281/zenodo.22064302.

Relevance. The historical-lineage member of the bulk trio: the axiomatic method of Euclid’s *Elements* carried into the six-dimensional reconstruction, with the GU-5 constants and the OGN-5 ledger in their historical setting. Cited with the constitutional papers as the foundation on which the present black-hole module stands. First cited in Part II.

K. Arisaka (2026). *GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT*. Preprint, available at <https://preprints.arisaka-gg.org/a5-gdos/>. doi:10.5281/zenodo.22064304.

Relevance. The open-system runtime this paper computes in. Cited for the exact reduction and channel structure T2–T4 (Nakajima–Zwanzig, Stinespring/CPTP, GKLS — the port machinery of Part IV), for Definition 5’s Port-Freshness Clause (L-GDOS-Fresh, consumed by Lemma 2.1 with its grade left at its A5 home), and for the boundary-gravity reduction T8.1–T8.2, whose scale-overlap Newton constant coincides by construction with Step 1’s induced $1/G_N$. First cited in Part II.

K. Arisaka (2026). *The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind.* Preprint, available at <https://preprints.arisaka-gg.org/a6-grip/>.

[doi:10.5281/zenodo.22064306](https://doi.org/10.5281/zenodo.22064306). Companion paper Arisaka-GG-A6-GRIP.

Relevance. Companion paper establishing the four geometric-invariance pillars (TGG, CCI, GFF, GRIP) of the bulk GG-6 composite gauge–gravity connection. Cited for the GRIP-In/GRIP-Out regime taxonomy (Hawking–Page as the one genuine GRIP-Out of black-hole physics; the Page time as a GRIP-In milestone), for the GFF ledger structure theorems behind the four laws, and for the pillar-forcing of GDOS. First cited in Part II.

K. Arisaka (2026). *The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS.* Preprint, available at <https://preprints.arisaka-gg.org/a7-qi4/>.

[doi:10.5281/zenodo.22064308](https://doi.org/10.5281/zenodo.22064308).

Relevance. The quantum-information pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL) that force GDOS from the information side, as GG-A6-GRIP’s geometric pillars force it from the geometric side. Cited at the central-claim funnel: the open-system law whose horizon instance this paper develops is pillar-forced, not chosen. First cited in §1.

K. Arisaka (2026). *Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary.* Preprint, available at <https://preprints.arisaka-gg.org/a8-qm/>.

[doi:10.5281/zenodo.22064310](https://doi.org/10.5281/zenodo.22064310).

Relevance. The first paper of the 4D-Physics Trio. Operationalizes T2 Operational Closure and T7 4D Quantum Mechanics from GDOS into a concept-by-concept reformulation of textbook quantum mechanics: the Born rule as a theorem of the Stinespring dilation (closed at the relaxation level; the field-grade representation step is the declared residue), EPR / Bell / CHSH / Aspect / Hensen / Giustina as evidence of the non-invertibility of the 6D-to-4D projection, Schrödinger’s cat as a GDOS boundary-decoherence statement, and the central thesis that Einstein was right in 6D and Bohr was right in 4D, both within their respective domains. First cited in §13.1.

K. Arisaka (2026). *The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry.* Preprint, available at <https://preprints.arisaka-gg.org/a9-sm/>.

[doi:10.5281/zenodo.22064312](https://doi.org/10.5281/zenodo.22064312).

Relevance. The Standard-Model companion: the 28 parameters as observations of the one geometry. Cited for the boundary field content — the zero-mode Standard-Model spectrum whose sign-resolved species sum, with the Kaluza–Klein tower, fixes the induced Newton constant (Appendix R) — and for the massless content entering the $c_{\log}$ prediction of §31.9. First cited in the front matter.

K. Arisaka (2026). *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry.* Preprint, available at <https://preprints.arisaka-gg.org/a10-cosmos/>.

[doi:10.5281/zenodo.22064314](https://doi.org/10.5281/zenodo.22064314).

Relevance. This companion paper derives the cosmological parameters $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ from the GG-Constitution and the 6D bulk. Cited here for the observer-vacuum reading consumed at the thermal sections, the curved-space existence residue its observer-vacuum section imports (the collapse factorization itself closed at model level as R14–R15 of GG-A8-QM/GG-P4-Real), adopted as this paper’s named open problem, the PBH-scarcity theorem bounding the dark-matter share, and the de Sitter two-GGA terminus of the way-station close. First cited in Part II.

Katsushi Arisaka (2026). *The Geometric Inevitability of Cosmos, Life, and Mind: The Self-Referential Closure of the GG-Theory Program*. Arisaka-GG-A11-Origin.

Preprint, available at <https://preprints.arisaka-gg.org/a11-origin/>.

[doi:10.5281/zenodo.22064319](https://doi.org/10.5281/zenodo.22064319). No result of the present paper depends on it.

Relevance. Companion paper supplying the Geometric Attractor vocabulary — the GFF Morse flow, basin descent, channel-mediated exit, and the No-Exit theorem of the Ground Geometric Attractor (GGA) — that this paper uses as its dynamical grammar: Hawking radiation is read as channel-mediated exit on the gravitational stratum, ringdown as the GFF Hessian at the attractor, and evaporation as descent toward the GGA. First cited in §15.

K. Arisaka (2026). *The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory*. Arisaka-GG-P1-KK preprint (Predictions Branch); <https://preprints.arisaka-gg.org/p1-kk/>.

[doi:10.5281/zenodo.22064321](https://doi.org/10.5281/zenodo.22064321).

Relevance. The lock-cluster companion: the non-spatial Kaluza–Klein resonance near 3.85 TeV. Cited where the compactification scale $\ell_{\text{lock}} = \hbar/(M_{\text{lock}}c)$ enters the suppression $(\ell_{\text{lock}}/r_s)^2$ of every observable deviation from Kerr — the origin of this paper’s null predictions — and as one of the external falsifiers by which the black-hole module can be killed from outside black-hole physics. First cited in §47.

K. Arisaka (2026). *Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF, and GDOS Morse Stability*. Arisaka-GG-P2-N3 preprint (Predictions Branch); <https://preprints.arisaka-gg.org/p2-n3/>.

[doi:10.5281/zenodo.22064323](https://doi.org/10.5281/zenodo.22064323).

Relevance. The three-generations companion: $N_g = 3$ conditionally derived (the three-leg pincer, at the paper’s own stated grade). Cited beside the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$: $N_g = 3$, not a ledger entry, fixes the generation count in the field content of the species sum — the ledger governs the content, never the $1/4$ itself (Theorem 1). First cited in Part III.

K. Arisaka (2026). *The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory*. Arisaka-GG-P3-Electron preprint (Predictions Branch); <https://preprints.arisaka-gg.org/p3-electron/>.

[doi:10.5281/zenodo.22064325](https://doi.org/10.5281/zenodo.22064325).

Relevance. The electron companion: the electron as the unique ground state of its stratum. Cited in the interlocking-falsifiability argument: the same runtime that reads the black hole as a

way-station attractor stabilizes the electron, so the black-hole module shares its falsifiers with the particle-physics branch. First cited in §61.7.

K. Arisaka (2026). *The Real Geometry Beneath the Complex Numbers of Physics: Positive Frequency is Holomorphy — from Cardano’s Cubics to Twistor Space, via the Sixth Dimension of GG-Theory.* Preprint, available at <https://preprints.arisaka-gg.org/p4-real/>. doi:10.5281/zenodo.22064327.

Relevance. The origin-of-*i* member paper of the Foundations Branch of the twenty-two-paper cohort: the two real sources of the derived complex structure (the compact phase circle and the causal clock), the holomorphy bridge, the infrared door, and the custody principle at full length. Its causal clock underlies the analytic structure of the present paper: the rotating anchor and the collapse channel of §27 are entries R14–R15 of the real-bulk series it develops, and its custody proposition is the pattern their proofs extend to five instances. Cited at roster level in §10.1. First cited in §10.1.

K. Arisaka (2026). *The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature’s Mirrors Break, and Why the Universe Is Made of Matter.* Preprint, available at <https://preprints.arisaka-gg.org/p5-cpt/>. doi:10.5281/zenodo.22064329.

Relevance. The CPT member paper of the Foundations Branch: C, P, and T as orientation reversals of the three real bulk factors, time defined by its four endowments, and the law-level arrow of the one-way clock — kin of the horizon thermodynamics derived here. Its Wick-quartet theorem T-P5-B1 proves the KMS thermal clock face on which this paper’s Hawking temperature runs. Cited at roster level in §10.1. First cited in §10.1.

K. Arisaka. *The Proton as the Ground Geometric Attractor of Matter: The Eleven Questions of the Proton, Closed and Fenced within GG-Theory,* Arisaka-GG-P6-Proton preprint. Preprint, available at <https://preprints.arisaka-gg.org/p6-proton/>. doi:10.5281/zenodo.22064331.

Relevance. The object this paper’s central classification is drawn against. Part VI argues that a black hole is a Geometric Attractor with an exit channel and therefore a way station rather than a GGA; the proton is the same ledger’s minimum on the matter stratum with no exit channel at all, and it is what a true ground state looks like when one is reached. The contrast is the content of the classification, and without a paper on the other side of it the way-station reading would be a distinction with only one example. It supplies no number to any derivation here. First cited in §42.2.

K. Arisaka. *Symmetry, Conservation, and Geometry: Seven Steps from Noether, Bianchi, to Hopkins–Singer, and the Ledger the Open-System Turn Pays Out,* Arisaka-GG-W1-Symmetry preprint. Preprint, available at <https://preprints.arisaka-gg.org/w1-symmetry/>. doi:10.5281/zenodo.22064340.

Relevance. The White-Paper Branch survey of how conservation moved from symmetry to cohomology, and the background to what the arrow section does without arguing for it. The

direction of the generalized second law is established here from non-invertibility and a custody chain rather than from any symmetry of the horizon, and a reader who finds a conservation statement with no symmetry behind it unfamiliar will find the century of argument there rather than in this Part. It supplies no link of the chain and no number. First cited in §41.1.

K. Arisaka. *The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology*, Arisaka-GG-W2-Holonomy preprint.

Preprint, available at <https://preprints.arisaka-gg.org/w2-holonomy/>.

doi:10.5281/zenodo.22064342.

Relevance. The classification that Step 4 of the entropy derivation lands in. The microstates counted by $A/4G_N$ are identified here as the $U(1)_{Q_i}$ Wilson-line holonomies on the compact phase circle, traced out at the horizon; holonomy protection rather than symmetry protection is precisely W2’s subject, and the sixteen problems it sorts include the one this paper closes. Cited at roster level in the cohort map, and named here because the microstate identification is an instance of its classification rather than an application of its results. First cited in §31.6.

K. Arisaka (2026). *The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6*. Arisaka-GG-C1-Kaxion preprint (Cosmology & Astrophysics Branch); <https://preprints.arisaka-gg.org/c1-kaxion/>.

doi:10.5281/zenodo.22064333.

Relevance. The cohort’s dark-matter companion (mass $8.1 \mu\text{eV}$ ($\pm 1.3\%$ at fixed f_K); the 1.95 GHz laboratory line inside the closed 1.59–2.13 GHz window). Cited in the primordial-black-hole section: the cohort’s dark-matter slot is saturated by the Kaxion, bounding any PBH contribution to a subdominant share. First cited in §49.

K. Arisaka (2026). *The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$* . Arisaka-GG-C2-Hubble preprint (Cosmology & Astrophysics Branch); <https://preprints.arisaka-gg.org/c2-hubble/>.

doi:10.5281/zenodo.22064336.

Relevance. Companion paper reading the Hubble tension on the cohort’s one derived expansion history: a single background rate that the microwave background and every ladder-free probe read, with the calibrated distance ladder the predicted outlier and its excess carried as a named calibration systematic. Cited in the de Sitter–horizon section: the black-hole analysis here and the C2 cosmological analysis are two evaluations of one bulk-to-boundary projection, sharing the $S = A/4G_N$ horizon entropy. First cited in Part VII.

B. Other references

B. P. Abbott et al. (2016). Observation of gravitational waves from a binary black hole merger. *Physical Review Letters* **116**, 061102.

Relevance. The first direct detection of gravitational waves (GW150914) from a binary black-hole merger. Cited as the observational anchor for merger-ringdown phenomenology; in GG-Theory

language the quasi-normal ringdown measures the Hessian of the GFF ledger at the Geometric Attractor. First cited in Part I.

S. L. Adler (1982). Einstein gravity as a symmetry-breaking effect in quantum field theory. *Reviews of Modern Physics* **54**, 729.

Relevance. Adler’s review of Einstein gravity as a symmetry-breaking (induced) effect in quantum field theory. Cited in the induced-gravity lineage and in the weighted species sum rule — the modern form of Sakharov’s self-consistency condition relating $\sum_i \chi_i$ to the Planck-to-cutoff ratio. First cited in §64.

Ahmed Almheiri, Netta Engelhardt, Donald Marolf, and Henry Maxfield (2019). The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole. *Journal of High Energy Physics* **12**, 063.

Relevance. Independently and simultaneously with Penington, derived the Page curve from the generalized entropy of an evaporating black hole, introducing the island contribution. Cited alongside Penington’s paper as the QES/island breakthrough whose ledger structure GG-Theory reproduces at the horizon stratum. First cited in Part I.

Ahmed Almheiri, Thomas Hartman, Juan Maldacena, Edgar Shaghoulian, and Amirhossein Tajdini (2020). Replica wormholes and the entropy of Hawking radiation. *Journal of High Energy Physics* **05**, 013.

Relevance. Identified replica wormholes — connected saddles of the gravitational path integral over replicated boundaries — as the mechanism behind the island formula. Cited in Part I; the replica/conical-deficit machinery of the induced-gravity entropy derivation here is the GG-Theory counterpart of this construction. First cited in Part I.

A. Almheiri, T. Hartman, J. Maldacena, E. Shaghoulian, and A. Tajdini (2021). The entropy of Hawking radiation. *Reviews of Modern Physics* **93**, 035002.

Relevance. The authoritative review of the islands program by its architects, including the frank discussion of its open ends (dimensional rigor, the ensemble puzzle) that Appendix L’s residue list follows. The review against which this paper’s comparison Part is calibrated. First cited in Part XII.

Ahmed Almheiri, Donald Marolf, Joseph Polchinski, and James Sully (2013). Black holes: complementarity or firewalls?. *Journal of High Energy Physics* **02**, 062.

Relevance. The firewall argument: a smooth horizon, exterior effective field theory, and unitarity cannot all hold once entanglement monogamy is imposed. Cited in Part I and in the related-work placement as the information-theoretic anchor that any post-2012 framework, including GG-Theory, must answer. First cited in Part I.

J. M. Bardeen, B. Carter, and S. W. Hawking (1973). The four laws of black hole mechanics. *Communications in Mathematical Physics* **31**, 161–170.

Relevance. Derived the four laws of black-hole mechanics as formal analogies with thermodynamics. Cited in Part I, in the first-law consistency check of the GG-Theory entropy for the (M, J, Q)

family, and in the section where each of the four laws is re-derived as a GDOS open-system statement rather than an analogy. First cited in Part I.

J. D. Bekenstein (1973). Black holes and entropy. *Physical Review D* **7**, 2333–2346.

Relevance. The original statement that black holes carry entropy proportional to horizon area. Cited in Part I’s golden-age survey, in the generalized-second-law discussion, and throughout the entropy derivation: in the GG-Theory reading the horizon is an OGGEp port, and Bekenstein’s area entropy is the entanglement of the traced bulk sector, with the exact coefficient $1/4$ fixed by induced gravity (Theorem 1) — the quantitative completion of the proposal. First cited in Part I.

N. D. Birrell and P. C. W. Davies (1982). *Quantum Fields in Curved Space*. Cambridge University Press, Cambridge.

Relevance. The classic textbook for quantum field theory on curved backgrounds — Bogoliubov transformations, the Unruh effect, and the Hawking computation. Appendix H’s derivation follows its chs. 3 and 8 most closely. First cited in Part XII.

L. Bombelli, R. K. Koul, J. Lee, and R. D. Sorkin (1986). A quantum source of entropy for black holes. *Physical Review D* **34**, 373.

Relevance. The first identification of the entanglement entropy of quantum fields across a horizon, including its ultraviolet-divergent area law. Cited as the origin of the entanglement reading of black-hole entropy that GG-Theory promotes from divergent correction to the exact $A/4G_N$ itself. First cited in §31.2.

C. G. Callan and F. Wilczek (1994). On geometric entropy. *Physics Letters B* **333**, 55.

Relevance. Introduced the geometric (conical-deficit/replica) definition of entanglement entropy in field theory. Cited in the heat-kernel matching step: the equality of the entropy divergence and the induced $1/4G_N$ holds because both are the same Seeley–DeWitt coefficient a_1 of the same modes. First cited in §31.4.

S. Carlip (2000). Logarithmic corrections to black hole entropy from the Cardy formula. *Classical and Quantum Gravity* **17**, 4175.

Relevance. Derived logarithmic corrections to black-hole entropy from the Cardy formula. Cited in the $c_{\log}$ discriminator discussion and in the related-work placement as a sharp subleading test that any microscopic accounting, including the GG-Theory one, must face. First cited in §31.9.

S. M. Carroll (2004). *Spacetime and Geometry: An Introduction to General Relativity*. Addison-Wesley, San Francisco.

Relevance. The standard modern first-pass graduate text. Recommended in the reading course as the gentler route into the Schwarzschild derivation and causal structure of Appendix E before the heavier treatments of Wald and MTW. First cited in Part XII.

H. Casini (2008). Relative entropy and the Bekenstein bound. *Classical and Quantum Gravity* **25**, 205021.

Relevance. Recast the Bekenstein bound as positivity of a relative entropy, replacing heuristic entropy accounting with a channel-monotone quantity. Cited in the second-law-direction section

as the reading this paper’s exchange-class identity makes exact at the horizon: generalized-entropy production as spent relative entropy toward the thermal reference. First cited in §41.

S. Chandrasekhar (1983). *The Mathematical Theory of Black Holes*. Oxford University Press, Oxford.

Relevance. The complete classical theory of the Schwarzschild, Reissner–Nordström, and Kerr solutions, including the perturbation theory behind the quasi-normal ringdown spectra that Part VIII’s observational tests consume. The depth reference for Appendices E and P. First cited in Part XII.

P. C. W. Davies and S. A. Fulling (1976). Radiation from a moving mirror in two dimensional space-time: conformal anomaly. *Proceedings of the Royal Society of London A* **348**, 393–414.

Relevance. The moving-mirror model: an accelerating boundary with late-time exponential redshift radiates a thermal flux at the horizon temperature — the exactly solvable model form of gravitational collapse. T-BH-4’s class is this one, and its factorization content — universal late-time channel, decaying transient — is proved and verified on Davies–Fulling trajectories. First cited in §27.

J.-G. Demers, R. Lafrance, and R. C. Myers (1995). Black hole entropy without brick walls. *Physical Review D* **52**, 2245.

Relevance. Computed black-hole entropy without brick walls, confirming that the one-loop entropy matches the renormalization of G_N mode by mode. Cited in the Susskind–Uglum identity cluster underpinning the heat-kernel matching of the entropy derivation. First cited in §31.4.

Gia Dvali and Michele Redi (2008). Black Hole Bound on the Number of Species and Quantum Gravity at LHC. *Phys. Rev. D* **77**, 045027. doi: 10.1103/PhysRevD.77.045027.

Relevance. The species bound: N light species lower the gravitational cutoff to $\Lambda_{\text{sp}} = M_{\text{Pl}}/\sqrt{N}$. Cited in the species-puzzle discussion — why the entropy coefficient is $1/4$ and not $N/4$ — which the induced-gravity derivation resolves by making G_N itself carry the species sum. First cited in Part III.

A. Einstein (1915). Die Feldgleichungen der Gravitation, p. 844.

Relevance. Einstein’s 1915 paper introduced the field equations of general relativity, $G_{\mu\nu} = 8\pi G_N T_{\mu\nu}$, identifying gravity with the curvature of spacetime as a dynamical field obeying a variational principle. This is the foundational document of modern gravitational theory and the historical starting point for every attempt to unify gravity with the other forces. We cite it at the opening of Part I as the start of the century this paper retraces: the field equations whose horizons, thermodynamics, and information problem the GG-Theory reading addresses. First cited in Part I.

Albert Einstein and Nathan Rosen (1935). The particle problem in the general theory of relativity. *Physical Review* **48**, 73–77.

Relevance. Discovered that the analytically extended Schwarzschild geometry contains a bridge connecting two asymptotic regions. Cited in Part I’s historical survey as the first hint of the

nontrivial global structure later made explicit by the Kruskal extension and revived in the ER=EPR proposal. First cited in Part I.

Event Horizon Telescope Collaboration (2022). First Sagittarius A* Event Horizon Telescope results. *Astrophysical Journal Letters* **930**, L12.

Relevance. The first horizon-scale image of Sgr A*, completing the supermassive pair with M87*. Cited in the black-hole census: their entropies (10^{90} – $10^{97} k_B$) are the largest single-object entropies in the universe, and GG-Theory predicts their Kerr shadows at EHT precision. First cited in §46.

Event Horizon Telescope Collaboration, Kazunori Akiyama, et al. (2019). First M87 event horizon telescope results. I. The shadow of the supermassive black hole. *Astrophysical Journal Letters* **875**, L1.

Relevance. The first horizon-scale image of a black hole (M87*). Cited as the observational benchmark for horizon-scale geometry: GG-Theory predicts the Kerr shadow at EHT precision, with induced-gravity corrections suppressed far below observability. First cited in Part I.

David Finkelstein (1958). Past-future asymmetry of the gravitational field of a point particle. *Physical Review* **110**, 965–967.

Relevance. Introduced the ingoing-null coordinate system in which the Schwarzschild surface is manifestly regular — a one-way membrane rather than a singularity. Cited in Part I for establishing that the horizon is a coordinate artifact, a fact the GG-Theory boundary reading inherits. First cited in Part I.

V. P. Frolov and D. V. Fursaev (1997). Mechanism of generation of black hole entropy in Sakharov’s induced gravity. *Physical Review D* **56**, 2212.

Relevance. Worked out the explicit mechanism by which black-hole entropy is generated in Sakharov-style induced gravity. Cited with Jacobson’s paper as the induced-gravity lineage whose program the GG-Theory derivation completes for the generic black hole. First cited in §31.2.

V. P. Frolov and I. D. Novikov (1998). *Black Hole Physics: Basic Concepts and New Developments*. Kluwer Academic, Dordrecht.

Relevance. The comprehensive monograph bridging classical black-hole theory and the quantum developments. Recommended for the astrophysical and Kerr–Newman chapters of the reading course (Appendix P). First cited in Part XII.

G. W. Gibbons and S. W. Hawking (1977). Cosmological event horizons, thermodynamics, and particle creation. *Physical Review D* **15**, 2738.

Relevance. Established that the de Sitter cosmological horizon carries a temperature and an entropy $S_{\text{dS}} = A_H/4G_N$. Cited in the de Sitter section, where the induced-gravity replica derivation applies verbatim on the Euclidean four-sphere — in GG-Theory the de Sitter entropy is the same theorem as the black-hole one. First cited in §52.2.

G. W. Gibbons and S. W. Hawking (1977). Action integrals and partition functions in quantum gravity. *Physical Review D* **15**, 2752.

Relevance. The boundary-term prescription that makes the Euclidean gravitational action finite and the saddle-point free energy well defined — the computation worked term by term in the Euclidean chapter of the appendices (Step 2), whose output $I_E = \beta M/2$ feeds the F_{GEP} piece of the operational ledger. First cited in Appendix I.

D. Harlow (2016). Jerusalem lectures on black holes and quantum information. *Reviews of Modern Physics* **88**, 015002.

Relevance. The standard modern lecture notes connecting quantum information theory to black-hole physics: channels, the Page argument, holographic entanglement entropy, and bulk reconstruction. The recommended toolkit companion for Appendices K, L, and O. First cited in Part XII.

S. W. Hawking (1975). Particle creation by black holes. *Communications in Mathematical Physics* **43**, 199–220.

Relevance. The discovery that black holes radiate thermally at $k_B T_H = \hbar\kappa/2\pi c$. Cited in Part I for the discovery and the paradox it created, in the Bogoliubov sections where the Planck spectrum is re-derived as the OGGE partial trace of a two-mode-squeezed bulk state, and in T-BH-4, whose collapse-channel factorization imports Hawking’s gravitating derivation with attribution. First cited in Part I.

S. W. Hawking (1976). Breakdown of predictability in gravitational collapse. *Physical Review D* **14**, 2460–2473.

Relevance. Hawking’s explicit argument that black-hole evaporation destroys information — pure states evolve to mixed states. Cited in Part I and in the framework-comparison table as the sharpest statement of the paradox; GG-Theory keeps the semiclassical calculation while removing its information-loss reading via the open-system boundary projection. First cited in Part I.

S. W. Hawking and D. N. Page (1983). Thermodynamics of black holes in anti-de Sitter space. *Communications in Mathematical Physics* **87**, 577.

Relevance. The first-order phase transition between thermal AdS and the large black hole — worked as a three-Step notebook in the appendices — and the cleanest classical template for a branch switch selected by a free-energy ledger, which is how this paper reads the Page turnover (T-BH-2). First cited in Appendix J.

S. W. Hawking and R. Penrose (1970). The singularities of gravitational collapse and cosmology. *Proceedings of the Royal Society A* **314**, 529–548.

Relevance. The general singularity theorems extending Penrose’s 1965 result to cosmological and generic collapse settings. Cited in Part I as the rigorous statement that classical general relativity predicts its own breakdown inside horizons. First cited in Part I.

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

Relevance. The black-holes-as-mirrors analysis of scrambling and information retrieval after the Page time. Cited in the related-work placement as one of the information-theoretic anchors of the modern paradox literature that the GG-Theory account engages. First cited in §36.3.

T. Jacobson (1994). Black hole entropy and induced gravity. *arXiv:gr-qc/9404039*.

Relevance. Jacobson’s proposal that black-hole entropy is entanglement entropy in induced gravity, with the divergence absorbed into G_N . Cited as the direct ancestor of the present derivation: GG-Theory supplies the missing premise (a fully induced G_N , with no bare term to split against) that turns the proposal into the exact $1/4$. First cited in §31.2.

T. Jacobson (1995). Thermodynamics of spacetime: the Einstein equation of state. *Physical Review Letters* **75**, 1260.

Relevance. Jacobson’s derivation of the Einstein equation as a thermodynamic equation of state from the Clausius relation on local horizons. Cited in the related-work placement as a complementary geometric reading of the same area law that the induced-gravity derivation makes exact. First cited in §64.

D. Kabat (1995). Black hole entropy and entropy of entanglement. *Nuclear Physics B* **453**, 281.

Relevance. Identified the contact-term subtlety in the entanglement entropy of gauge fields. Cited in the honest-accounting section as a known complication: the heat-kernel matching is clean for minimally coupled scalars but requires care for gauge fields and the graviton. First cited in §31.10.

B. S. Kay and R. M. Wald (1991). Theorems on the uniqueness and thermal properties of stationary, nonsingular, quasifree states on spacetimes with a bifurcate Killing horizon. *Physics Reports* **207**, 49–136.

Relevance. The rigorous uniqueness and thermality theorems on bifurcate Killing horizons — and the proof that the Kerr exterior admits no Hadamard state invariant under the horizon Killing flow — no Hartle–Hawking analogue. T-BH-3 takes the latter as the exact form of the ergoregion obstruction: the rotating-frame selection fails for a reason the literature has made precise, and the move of the anchor to the asymptotic frame is the framework’s placement of that fact. First cited in §27.

Roy P. Kerr (1963). Gravitational field of a spinning mass as an example of algebraically special metrics. *Physical Review Letters* **11**, 237–238.

Relevance. The exact solution for a rotating, uncharged black hole. Cited in Part I’s historical survey and again where the GG-Theory entropy $S = A/4G_N$ and the stationary first law are verified across the full Kerr family. First cited in Part I.

Martin D. Kruskal (1960). Maximal extension of Schwarzschild metric. *Physical Review* **119**, 1743–1745.

Relevance. Constructed the maximal analytic extension of the Schwarzschild solution, exhibiting the full four-region global geometry (two exterior regions, black-hole and white-hole interiors). Cited in Part I for the global causal structure on which all later discussion of horizons and interiors rests. First cited in Part I.

J. Maldacena (1998). The large N limit of superconformal field theories and supergravity. *Adv. Theor. Math. Phys.* **2**, 231.

Relevance. Maldacena’s AdS/CFT conjecture is the modern archetype for bulk–boundary duality and motivates the idea that geometry can encode quantum gauge dynamics. We cite it as a historical anchor for the GGEP/OGGEP posture: bulk geometry is not merely a classical background, but a structure from which boundary quantum physics and observables can be exported. First cited in Part I.

Juan Maldacena and Leonard Susskind (2013). Cool horizons for entangled black holes. *Fortschritte der Physik* **61**, 781–811.

Relevance. The ER=EPR proposal that entangled pairs are connected by non-traversable Einstein–Rosen bridges. Cited in Part I’s survey of post-firewall responses as a geometric reading of entanglement against which the GG-Theory bulk-to-boundary account is contrasted. First cited in Part I.

Samir D. Mathur (2009). The information paradox: a pedagogical introduction. *Classical and Quantum Gravity* **26**, 224001.

Relevance. Proved that no small correction to the Hawking calculation can restore unitarity — order-one departures at the horizon are required. Cited twice as the theorem that closes off the small-corrections escape route and thereby motivates the GG-Theory branch-switch reading of the Page turnover. First cited in Part III.

C. W. Misner, K. S. Thorne, and J. A. Wheeler (1973). *Gravitation*. W. H. Freeman, San Francisco.

Relevance. The encyclopedic classic of general relativity. Recommended in the Technical Appendices’ reading course as a primary text for the Schwarzschild geometry, the Kruskal extension, and the classical black-hole chapters (chs. 23–25 and 31–34), in whose conventions a generation of the cited literature was written. First cited in Part XII.

E. T. Newman, E. Couch, K. Chinnapared, A. Exton, A. Prakash, and R. Torrence (1965). Metric of a rotating, charged mass. *Journal of Mathematical Physics* **6**, 918.

Relevance. The Kerr–Newman metric for a rotating, charged mass — the most general stationary black hole of the (M, J, Q) family. Cited in the verification that the GG-Theory entropy and the stationary first law hold for every member of the family. First cited in §31.8.

J. Robert Oppenheimer and Hartland Snyder (1939). On continued gravitational contraction. *Physical Review* **56**, 455–459.

Relevance. The first detailed general-relativistic calculation of gravitational collapse (a pressureless dust ball), showing that horizon formation is the generic endpoint of collapse rather than a pathology. Cited in Part I as the result that turned the Schwarzschild solution into astrophysics. First cited in Part I.

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

Relevance. Established the Page curve: under unitary evaporation the radiation entropy must rise and then fall, turning over at the Page time. Cited throughout as the quantitative benchmark any resolution of the information paradox must reproduce; in GG-Theory the turnover is read as a GFF-ledger branch switch. First cited in Part I.

L. Parker and D. Toms (2009). *Quantum Field Theory in Curved Spacetime: Quantized Fields and Gravity*. Cambridge University Press, Cambridge.

Relevance. The modern treatment of effective action and heat-kernel methods in curved spacetime. The recommended companion for the induced-gravity species sum of Appendix R, whose Seeley–DeWitt technology it develops in full. First cited in Part XII.

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

Relevance. Derived the Page curve directly from semiclassical gravity via entanglement-wedge reconstruction and quantum extremal surfaces. Cited as one of the two 2019 breakthrough papers; GG-Theory recovers the island prescription as the GFF ledger at the horizon stratum, and benchmarking against these computations is flagged as an open task. First cited in Part I.

Roger Penrose (1965). Gravitational collapse and space-time singularities. *Physical Review Letters* **14**, 57–59.

Relevance. The first singularity theorem: under mild energy and causality assumptions, gravitational collapse generically produces incomplete geodesics. Cited in Part I; together with the Hawking–Penrose theorems it establishes that singularities — and hence horizons — are generic, not artifacts of symmetry. First cited in Part I.

E. Poisson (2004). *A Relativist’s Toolkit: The Mathematics of Black-Hole Mechanics*. Cambridge University Press, Cambridge.

Relevance. The standard modern source for null congruences, the Raychaudhuri equation, and horizon mechanics. Recommended alongside Appendices E–F, whose area-theorem derivation follows the toolkit’s route. First cited in Part XII.

H. Reissner and G. Nordström (1916). On the self-gravitation of the electric field according to Einstein’s theory. *Annalen der Physik* **50**, 106.

Relevance. The charged, spherically symmetric solution of the Einstein–Maxwell equations. Cited where the GG-Theory entropy $S = A/4G_N$ and the first law are verified across the full (M, J, Q) family of stationary black holes. First cited in §31.8.

Shinsei Ryu and Tadashi Takayanagi (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters* **96**, 181602.

Relevance. Proposed the holographic formula equating entanglement entropy with the area of a minimal bulk surface. Cited in Part I as the modern entanglement–geometry link whose lesson — entropy is geometry — the induced-gravity derivation of the exact $1/4$ makes literal for generic horizons. First cited in Part I.

A. D. Sakharov (1968). Vacuum quantum fluctuations in curved space and the theory of gravitation. *Soviet Physics Doklady* **12**, 1040.

Relevance. Sakharov’s induced-gravity hypothesis: the Einstein–Hilbert action is generated entirely by matter loops. Cited as the head of the lineage to which this paper’s entropy derivation belongs; GG-Theory promotes the hypothesis to a constitutional consequence of GGEP, and the weighted species sum rule is the modern form of his self-consistency condition. First cited in §31.2.

Karl Schwarzschild (1916). Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, 189–196.

Relevance. The first exact solution of the Einstein field equations: the static, spherically symmetric vacuum metric whose surface $r = 2G_N M$ became the black-hole horizon. Cited in Part I as the opening of the sixty-year arc — exact solution, horizon, thermodynamics, paradox — that the present paper closes in GG-Theory terms. First cited in Part I.

A. Sen (2013). Logarithmic corrections to Schwarzschild and other non-extremal black-hole entropy in different dimensions. *Journal of High Energy Physics* **2013**(04), 156.

Relevance. Computed logarithmic corrections to the entropy of Schwarzschild and other non-extremal black holes across dimensions. Cited in the $c_{\log}$ discussion: the log coefficient is a parameter-free discriminator between frameworks, and in the induced-gravity derivation it descends from the a_2 heat-kernel (conformal-anomaly) coefficient. First cited in §31.9.

S. N. Solodukhin (2011). Entanglement entropy of black holes. *Living Reviews in Relativity* **14**, 8.

Relevance. The standard review of the entanglement entropy of black holes, including the distributional (conical-deficit) curvature on the replica cover and the structure of logarithmic corrections. Cited for the conical-singularity identity used in the replica computation and for the $c_{\log}$ discriminator. First cited in §31.5.

H. Spohn (1978). Entropy production for quantum dynamical semigroups. *Journal of Mathematical Physics* **19**, 1227.

Relevance. The entropy-production theorem for quantum dynamical semigroups with a KMS fixed point: the production functional is non-negative and vanishes on the fixed point. Cited in the second-law-direction section: the exchange-channel identity there is the black-hole instance of Spohn’s theorem, with the produced entropy carried by horizon area through the induced first law. First cited in §41.

A. A. Starobinsky (1973). Amplification of waves during reflection from a rotating black hole. *Soviet Physics JETP* **37**, 28–32.

Relevance. The discovery of superradiance: waves scattered off a rotating black hole in the band $\omega < m\Omega_H$ return amplified, extracting rotational energy. T-BH-3 reads the phenomenon structurally: the amplification is a symplectic Gaussian pumping channel on the asymptotic frame’s derived complex structure, which it never touches. First cited in §27.

Andrew Strominger and Cumrun Vafa (1996). Microscopic origin of the Bekenstein-Hawking entropy. *Physics Letters B* **379**, 99–104.

Relevance. The first exact microscopic count of black-hole entropy: D-brane states reproducing $S = A/4G_N$ for extremal, supersymmetric black holes. Cited repeatedly as the decisive point of contrast — the GG-Theory induced-gravity count is exact for generic (astrophysical) black holes, while the string count is exact only in BPS-protected sectors — and as the benchmark for the open species-by-species comparison. First cited in Part III.

L. Susskind (1995). The world as a hologram. *Journal of Mathematical Physics* **36**(11), 6377–6396. <https://arxiv.org/abs/hep-th/9409089>.

Relevance. Susskind’s holographic principle paper is a key historical milestone leading to modern bulk–boundary dualities. We cite it to frame OGGEp as a constitution-level insistence that observation is a bulk→boundary export, generalized beyond correlators to full operational semantics. First cited in Part I.

L. Susskind and J. Lindesay (2005). *An Introduction to Black Holes, Information and the String Theory Revolution: The Holographic Universe*. World Scientific, Singapore.

Relevance. The most readable account of the information paradox, complementarity, and the holographic principle as the field understood them before the islands era. Recommended as the narrative spine for the reading course’s information chapters (Appendices G, K, and L). First cited in Part XII.

L. Susskind and J. Uglum (1994). Black hole entropy in canonical quantum gravity and superstring theory. *Physical Review D* **50**, 2700.

Relevance. Showed that the horizon entanglement entropy renormalizes Newton’s constant — the Susskind–Uglum identity equating the entropy divergence with the renormalization of $1/4G_N$. Cited repeatedly in the entropy derivation: in GG-Theory the identity is forced, because GGEP leaves no bare gravitational term for the divergence to hide in. First cited in §31.2.

Leonard Susskind, Larus Thorlacius, and John Uglum (1993). The stretched horizon and black hole complementarity. *Physical Review D* **48**, 3743–3761.

Relevance. The original black-hole complementarity proposal (stretched horizon; exterior and infalling descriptions both valid but never jointly held). Cited in the survey of pre-island responses to the paradox; GG-Theory replaces observer-complementarity with the 6D-bulk/4D-boundary split. First cited in Part VI.

G. ’t Hooft (1993). Dimensional reduction in quantum gravity. Essay dedicated to Abdus Salam; arXiv:gr-qc/9310026. <https://arxiv.org/abs/gr-qc/9310026>.

Relevance. ’t Hooft articulated early modern arguments for holography as a fundamental constraint on quantum gravity. We cite it (together with Susskind and Maldacena) to motivate why OGGEp treats the bulk→boundary map as definitional rather than optional. First cited in Part I.

W. G. Unruh (1974). Second quantization in the Kerr metric. *Physical Review D* **10**, 3194–3205.

Relevance. The quantum field theory of the Kerr background: the mode analysis in which the ergoregion’s negative co-rotating energies and the superradiant band appear at the field level, and the vacuum choices of the rotating geometry are first laid out. T-BH-3’s model statements are the structural skeleton of this analysis, with the vacuum multiplicity placed as port assignment relative to the asymptotic anchor. First cited in §27.

W. G. Unruh (1976). Notes on black-hole evaporation. *Physical Review D* **14**, 870.

Relevance. Unruh’s analysis of black-hole evaporation and the thermal response of accelerated detectors. Cited alongside Hawking’s derivation for the Bogoliubov-coefficient calculation that yields the Planck spectrum at $k_B T_H = \hbar\kappa/2\pi c$. First cited in §26.4.

M. Visser (2002). Sakharov’s induced gravity: a modern perspective. *Modern Physics Letters A* **17**, 977.

Relevance. The modern critical assessment of Sakharov’s induced gravity. Cited both for the lineage placement and for the honest caveat it documents: spin-dependent signs in the induced G_N are delicate, and for some spectra the naive sum even has the wrong sign — one reason the numerical species closure is flagged as open. First cited in §64.

R. M. Wald (1984). *General Relativity*. University of Chicago Press, Chicago.

Relevance. Wald’s textbook ‘General Relativity’ (1984) is the canonical graduate-level reference for differential geometry and general relativity, including the formal mathematical machinery of warped product metrics, Christoffel symbols, the Riemann/Ricci tensors, and the Einstein equations. We cite it as the standard reference for the curvature conventions used throughout, and for the Damour–Ruffini/Wald positive-frequency continuation prescription in the Bogoliubov derivation of the thermal spectrum. First cited in §26.4.

R. M. Wald (1993). Black hole entropy is the Noether charge. *Physical Review D* **48**, R3427.

Relevance. Wald’s Noether-charge formulation of black-hole entropy, valid for arbitrary diffeomorphism-invariant Lagrangians. Cited in the stationary first-law consistency check ($A/4G_N$ is precisely the Bekenstein–Hawking–Wald entropy the first law requires) and as a complementary geometric reading of the area law. First cited in §31.8.

R. M. Wald (1994). *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*. University of Chicago Press, Chicago.

Relevance. The mathematically careful formulation of quantum fields in curved spacetime and the definitive early statement of the generalized second law’s status. Recommended companion for Appendices G and I. First cited in Part XII.

A. C. Wall (2012). A proof of the generalized second law for rapidly changing fields and arbitrary horizon slices. *Physical Review D* **85**, 104049.

Relevance. The strongest available proof of the generalized second law: free and superrenormalizable fields, arbitrary horizon cuts, via the monotonicity of relative entropy under restriction and the half-sided modular structure of the horizon algebra. Cited as the declared ceiling of this paper’s second-law-direction result: what a full promotion of L-BH-Arrow beyond the exchange class must reach, registered open. First cited in §41.3.